Advanced
CCIE SERVICE PROVIDER
v3.0

www.MicronicsTraining.com

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

Paul Negron
CCIE #14856
SP

VOL-I

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Table of Content:

Subject	Page	Volume
Topology	4	Vol-I
EIGRP		
Lab 1 EIGRP Configuration	5	Vol-I
Lab 2 Advanced EIGRP Stub Configuration	76	Vol-I
Lab 3 EIGRP & Default-information	85	Vol-I
Lab 4 EIGRP Filtering	96	Vol-I
ISIS		
Lab 1 Basic ISIS configuration	105	Vol-I
Lab 2 Multi-Instance	117	Vol-I
Lab 3 Multi-Topology	125	Vol-I
Lab 4 Redistribution & Summarization	137	Vol-I
Lab 5 ISIS Efficiency Mechanism	147	Vol-I
Lab 6 Authentication	162	Vol-I
Lab 7 Advanced Configuration I	167	Vol-I
Lab 8 Advanced Configuration II	174	Vol-I
OSPF		
Lab 1 Advertising Networks	187	Vol-I
Lab 2 Optimization of OSPF & Adjusting Timers	190	Vol-I
Lab 3 OSPF Authentication	197	Vol-I
Lab 4 OSPF Cost	222	Vol-I
Lab 5 OSPF Summarization	227	Vol-I
Lab 6 Virtual-links and GRE Tunnels	234	Vol-I
Lab 7 OSPF Stub, T/Stub, and NSSAs	244	Vol-I
Lab 8 OSPF Filtering	254	Vol-I
Lab 9 Additional OSPF Filtering	280	Vol-I
Lab 10 Redirecting Traffic in OSPF	290	Vol-I
Lab 11 Database Overload Protection	295	Vol-I
Lab 12 OSPF Non-Broadcast Networks	300	Vol-I
Lab 13 OSPF Broadcast Networks	309	Vol-I
Lab 14 OSPF Point-to-Point Networks	313	Vol-I
Lab 15 OSPF Point-to-Multipoint Networks	317	Vol-I
Lab 16 OSPF Point-to-Multi Network – II	324	Vol-I
Lab 17 OSPF P-to-M Non-Broadcast Networks	331	Vol-I
Lab 18 OSPF and NBMA	337	Vol-I
Lab 19 Forward Address Suppression	345	Vol-I
Lab 20 OSPF NSSA No-redistribution & Injection of Default Routes	351	Vol-I
BGP		
Lab 1 Establishing Neighbor Adjacency	345	Vol-I
Lab 2 Route Reflectors	362	Vol-I

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Lab 3 Conditional Adv. & Back door	378	Vol-I
Lab 4 Route Dampening	393	Vol-I
Lab 5 Route Aggregation	402	Vol-I
Lab 6 The Community Attribute	422	Vol-I
Lab 7 BGP Cost Community	438	Vol-I
Lab 8 BGP & Load Balancing – I	447	Vol-I
Lab 9 BGP Load Balancing – II	470	Vol-I
Lab 10 BGP Unequal Cost Load Balancing	474	Vol-I
Lab 11 BGP Local Preference – I	482	Vol-I
Lab 12 BGP Local Preference – II	493	Vol-I
Lab 13 The AS-Path Attribute	501	Vol-I
Lab 14 The Weight Attribute	509	Vol-I
Lab 15 MED	516	Vol-I
Lab 16 Filtering Using ACLs & Prefix-lists	532	Vol-I
Lab 17 Regular Expressions	542	Vol-I
Lab 18 Advanced BGP Configurations	559	Vol-I
Lab 19 Administrative Distance	569	Vol-I
Lab 20 BGP Confederation	577	Vol-I
Lab 21 BGP Hiding Local AS Number	582	Vol-I
Lab 22 BGP Allowas-in	590	Vol-I

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Advanced
CCIE SERVICE PROVIDER
v3.0

www.MicronicsTraining.com

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

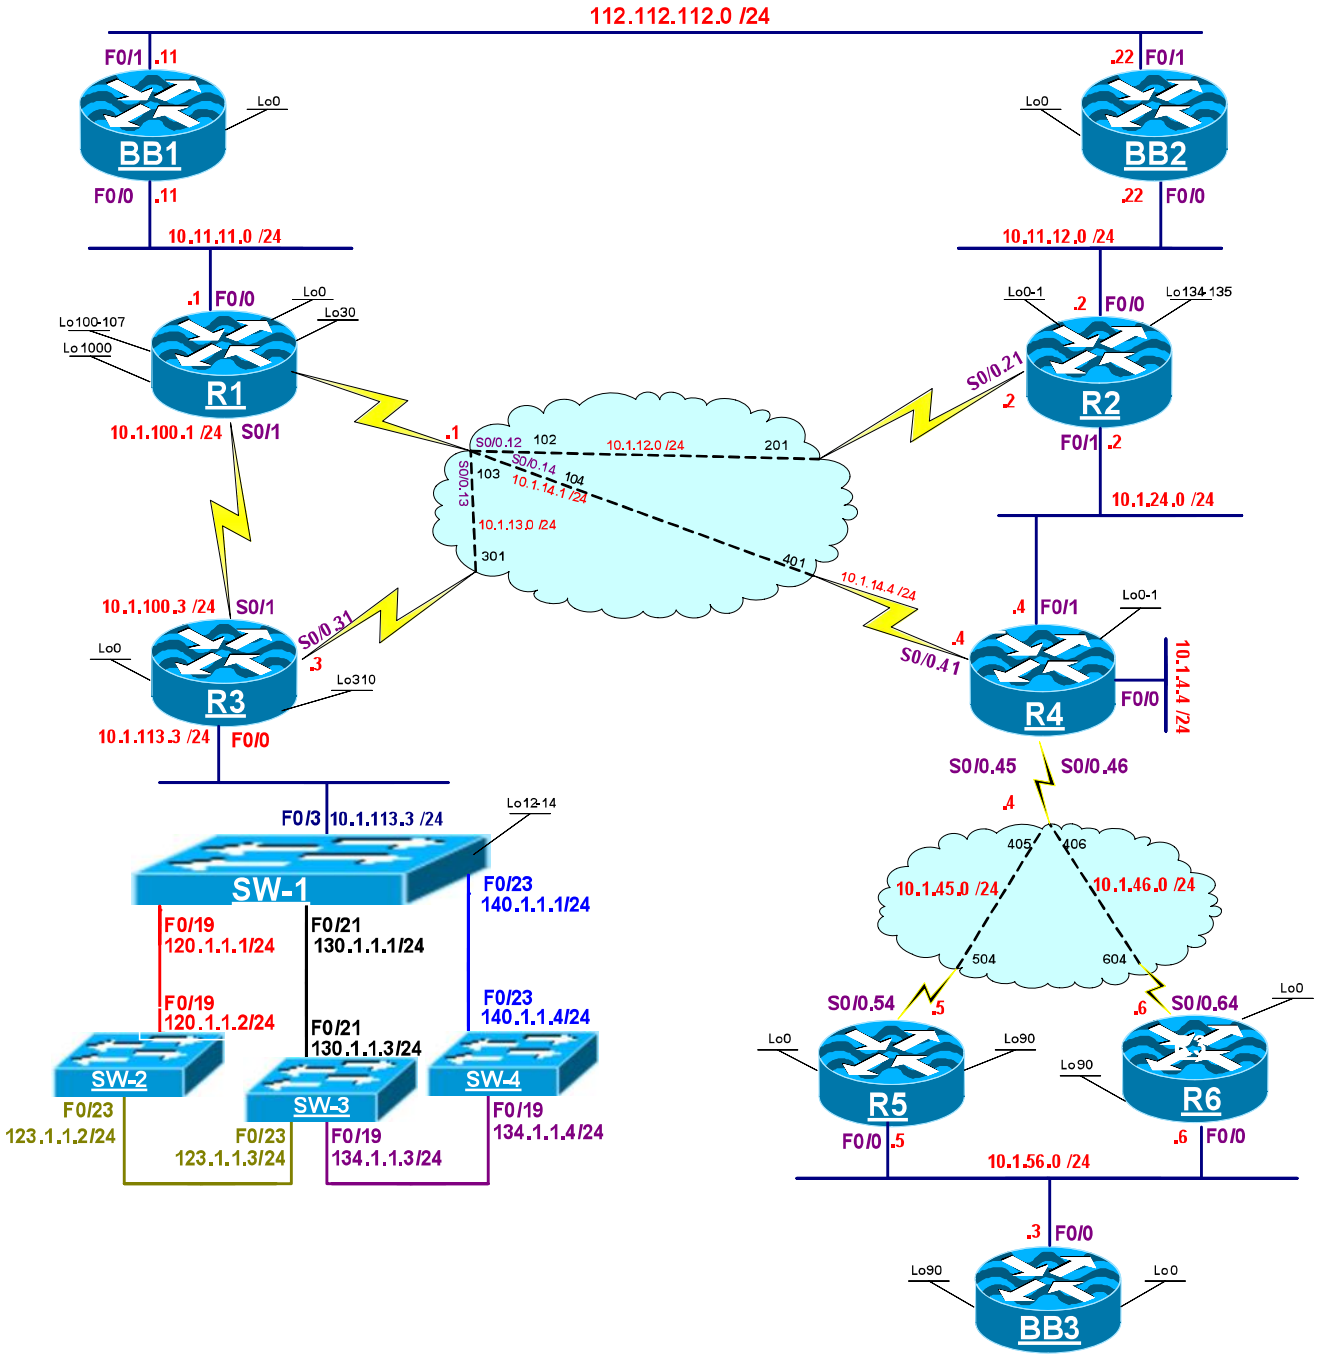
Paul Negron
CCIE #14856
SP

EIGRP

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Lab 1
Configuring EIGRP



©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Task 1

Configure the F0/0 interface of routers R5, R6 and BB3 in AS 100; these routers should be configured to advertise their Lo0 interface in this AS using the following policy:

BB3 and R6 should be configured to reach the existing and future internal routes advertised in this AS by going through R5, unless the prefix is advertised by the local router. DO NOT use PBR to accomplish this task. You should also disable the Classful behavior of these routers.

On R5

RIPv2, Eigrp and BGP are classless routing protocols and by default, they have a Classful nature/behavior, meaning that they automatically summarize the networks that they advertise to their major network boundary. This behavior can be disabled by using the "NO Auto-summary" router configuration command:

```
R5(config-if)#router eigrp 100
R5(config-router)#NO au
R5(config-router)#netw 5.5.5.5 0.0.0.0
R5(config-router)#netw 10.1.56.5 0.0.0.0
```

On R6

```
R6(config-if)#router eigrp 100
R6(config-router)#NO au
R6(config-router)#netw 6.6.6.6 0.0.0.0
R6(config-router)#netw 10.1.56.6 0.0.0.0
```

You should see the following console message:

%DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 10.1.56.5 (FastEthernet0/0) is up: new adjacency

On BB3

```
BB3(config-if)#router eigrp 100
BB3(config-router)#NO au
BB3(config-router)#netw 33.3.3.3 0.0.0.0
BB3(config-router)#netw 10.1.56.3 0.0.0.0
```

You should see the following two console messages stating that the local router has established two neighbor adjacency one with R5 and the second one with R6:

%DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 10.1.56.5 (FastEthernet0/0) is up: new adjacency
%DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 10.1.56.6 (FastEthernet0/0) is up: new adjacency

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the configuration:

On R5

```
R5#Show ip eigrp neighbors

IP-EIGRP neighbors for process 100
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   10.1.56.6               Fa0/0             10 00:00:09    5   200  0  40
0   10.1.56.3               Fa0/0             14 00:00:09    1   200  0  39

R5#sh ip eigrp interfaces

IP-EIGRP interfaces for process 100
                  Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface   Peers  Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Lo0           0        0/0        0       0/1           0           0
Fa0/0         2        0/0        3       0/1           50          0

R5#Show ip route eigrp

D    33.0.0.0/8 [90/156160] via 10.1.56.3, 00:03:34, FastEthernet0/0
D    6.0.0.0/8 [90/156160] via 10.1.56.6, 00:03:34, FastEthernet0/0
```

On R6

```
R6#Show ip eigrp neighbors

IP-EIGRP neighbors for process 100
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   10.1.56.5               Fa0/0             14 00:00:32    5   200  0  45
0   10.1.56.3               Fa0/0             11 00:00:28    1   200  0  39

R6#Show ip route eigrp

D    33.0.0.0/8 [90/156160] via 10.1.56.3, 00:10:23, FastEthernet0/0
D    5.0.0.0/8 [90/156160] via 10.1.56.5, 00:06:28, FastEthernet0/0
```

On BB3

```
BB3#Show ip eigrp neighbors

IP-EIGRP neighbors for process 100
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
H   Address          Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                     (sec)         (ms)       Cnt Num
1   10.1.56.5        Fa0/0             14 00:01:09    1   200  0  44
0   10.1.56.6        Fa0/0             14 00:01:28    1   200  0  40

BB3#Show ip route eigrp

D     5.0.0.0/8 [90/156160] via 10.1.56.5, 00:09:08, FastEthernet0/0
D     6.0.0.0/8 [90/156160] via 10.1.56.6, 00:13:04, FastEthernet0/0
```

In order for BB3 and R6 to go through R5 for existing and future internal routes, BB3 should have R5 as the next hop router to reach networks 6.0.0.0 /8 and 5.0.0.0 /8, and R6 should have R5 as the next hop to reach networks 5.0.0.0 /8 and 33.0.0.0 /8; To accomplish this task R6 and BB3 can NOT have a neighbor adjacency, or else they will exchange routes advertising themselves as the next hop for the routes they originated, therefore, R6 and BB3 should both be configured with a neighbor command pointing to R5; R5 should also be configured to have a neighbor command for R6 and another for BB3, as follows:

On R5

```
R5(config)#router eigrp 100
R5(config-router)#Neighbor 10.1.56.3 F0/0
R5(config-router)#Neighbor 10.1.56.6 F0/0
```

Note in order for two Eigrp routers to establish neighbor adjacency, they MUST both use Unicast or they must both use Multicast; you can NOT have one end using multicast and the other end using unicast, or if the local router has two Eigrp neighbors, it can NOT establish a Unicast neighbor adjacency to one and establish Multicast neighbor adjacency with the second one.

In the following configuration R6 is configured to establish neighbor adjacency with R5 using Unicast, because it is configured to use Unicast, it will drop adjacency to BB3, remember an Eigrp router can NOT Unicast to some neighbors and Multicast to the others.

When configuring Neighbor command you MUST specify the interface through which the Neighbor adjacency will be established.

On R6 and BB3

```
Rx(config)#router eigrp 100
Rx(config-router)#Neighbor 10.1.56.5 F0/0
```

To verify the configuration:

On R6

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R6#Show ip route eigrp

D     5.0.0.0/8 [90/156160] via 10.1.56.5, 00:06:52, FastEthernet0/0
```

On BB3

```
BB3#Show ip route eigrp

D     5.0.0.0/8 [90/156160] via 10.1.56.5, 00:04:31, FastEthernet0/0
```

On R5

```
R5#Show ip route Eigrp

D     33.0.0.0/8 [90/156160] via 10.1.56.3, 00:05:14, FastEthernet0/0
D     6.0.0.0/8 [90/156160] via 10.1.56.6, 00:08:02, FastEthernet0/0
```

Note R5 sees both routes in its routing table whereas, BB3 and R6 can ONLY see R5's routes. The reason for this behavior is "Split-Horizon" remember if R5 receives a route through an interface, it will NEVER advertise that route out of the interface through which it received the route.

The following command disables the split-horizon for Eigrp AS 100:

On R5

```
R5(config)#Int F0/0
R5(config-if)#NO IP split-horizon eigrp 100
```

Note you should see the following messages:

```
Neighbor 10.1.56.3 (FastEthernet0/0) is resync: split horizon changed
Neighbor 10.1.56.6 (FastEthernet0/0) is resync: split horizon changed
```

To test the configuration:

On R6

```
R6#Show ip route eigrp

D     33.0.0.0/8 [90/158720] via 10.1.56.5, 00:02:38, FastEthernet0/0
D     5.0.0.0/8 [90/156160] via 10.1.56.5, 00:15:28, FastEthernet0/0
```

On BB3

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
BB3#Show ip route eigrp

D    5.0.0.0/8 [90/156160] via 10.1.56.5, 00:13:10, FastEthernet0/0
D    6.0.0.0/8 [90/158720] via 10.1.56.5, 00:03:08, FastEthernet0/0
```

Task 2

Configure the frame-relay interface connecting R4 to R5 and R6; and the frame-relay interface of R5 and R6 in AS 200.

On R4

```
R4(config)#Router eigrp 200
R4(config-router)#No au
R4(config-router)#netw 10.1.45.4 0.0.0.0
R4(config-router)#netw 10.1.46.4 0.0.0.0
```

On R5

```
R5(config)#Router eigrp 200
R5(config-router)#No au
R5(config-router)#Netw 10.1.45.5 0.0.0.0
```

On R6

```
R6(config)#Router eigrp 200
R6(config-router)#No au
R6(config-router)#Netw 10.1.46.6 0.0.0.0
```

To verify the configuration:

On R4

```
R4#Show ip eigrp neighbors

IP-EIGRP neighbors for process 100
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   10.1.46.5               Se0/0.45          14 00:01:09   32   300  0  3
0   10.1.45.6               Se0/0.46          10 00:02:10   32   300  0  3
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 3

Configure a mutual redistribution on R5 and ensure that R4 and BB3 receive all routes.

Note when redistributing between two Eigrp routing protocols, it is NOT necessary to configure the "Bandwidth, Delay, Reliability, Load and MTU", unless otherwise stated.

On R5

```
R5(config)#Router Eigrp 100
R5(config-router)#redistribute eigrp 200

R5(config)#Router Eigrp 200
R5(config-router)#redistribute eigrp 100
```

To verify the configuration:

On BB3

```
BB3#Show ip route eigrp

D     5.0.0.0/8 [90/156160] via 10.1.56.5, 00:07:05, FastEthernet0/0
D     6.0.0.0/8 [90/158720] via 10.1.56.5, 00:07:05, FastEthernet0/0
      10.0.0.0/24 is subnetted, 3 subnets
D EX     10.1.46.0 [170/2684416] via 10.1.56.5, 00:02:23, FastEthernet0/0
D EX     10.1.45.0 [170/2172416] via 10.1.56.5, 00:02:23, FastEthernet0/0
```

On R4

```
R4#Show ip route eigrp

D EX 33.0.0.0/8 [170/2300416] via 10.1.45.5, 00:01:09, Serial0/0.45
D EX 5.0.0.0/8 [170/2297856] via 10.1.45.5, 00:01:09, Serial0/0.45
D EX 6.0.0.0/8 [170/2300416] via 10.1.45.5, 00:01:09, Serial0/0.45
      10.0.0.0/24 is subnetted, 6 subnets
D EX     10.1.56.0 [170/2172416] via 10.1.45.5, 00:01:09, Serial0/0.45
```

Task 4

Configure the frame-relay connection between R1, R2, R3 and R4 in AS 300.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1

```
R1(config)#Router eigrp 300
R1(config-router)#No au
R1(config-router)#Netw 10.1.12.1 0.0.0.0
R1(config-router)#Netw 10.1.13.1 0.0.0.0
R1(config-router)#Netw 10.1.14.1 0.0.0.0
```

On R2

```
R2(config)#Router Eigrp 300
R2(config-router)#No au
R2(config-router)#Netw 10.1.12.2 0.0.0.0
```

On R3

```
R3(config)#Router Eigrp 300
R3(config-router)#No au
R3(config-router)#Netw 10.1.13.3 0.0.0.0
```

On R4

```
R4(config)#Router eigrp 300
R4(config-router)#No au
R4(config-router)#Netw 10.1.14.4 0.0.0.0
```

To verify the configuration

On R1

```
R1#Show ip eigrp neighbors

IP-EIGRP neighbors for process 300
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
2   10.1.14.4               Se0/0.14          11 00:00:54   40   360  0  3
1   10.1.13.3               Se0/0.13          10 00:01:28   36   324  0  3
0   10.1.12.2               Se0/0.12          13 00:02:02   36   324  0  3
```

On R3

```
R3#Show ip eigrp neighbors

IP-EIGRP neighbors for process 300
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
                                          (sec)          (ms)        Cnt Num
0   10.1.13.1                 Se0/0.31      14 00:03:31 1278   5000  0   6
```

On R4

Neighbors in AS 200

```
R4#Show ip eigrp neighbors

IP-EIGRP neighbors for process 200
H   Address                 Interface       Hold Uptime    SRTT    RTO  Q  Seq
                                            (sec)          (ms)       Cnt Num
1   10.1.46.6               Se0/0.46          14 00:28:06   42    252  0  3
0   10.1.45.5               Se0/0.45          14 00:28:25   45    270  0  5

IP-EIGRP neighbors for process 300
H   Address                 Interface       Hold Uptime    SRTT    RTO  Q  Seq
                                            (sec)          (ms)       Cnt Num
0   10.1.14.1               Se0/0.41          11 00:02:35 1281   5000  0  9
```

Neighbors in AS 300

Task 5

Configure R1 to summarize and advertise its Lo100 – 107. The summary route should be advertised to R2, R3 and R4 based on the following policy:

1. R2 should receive the summary ONLY.
2. R3 should receive the summary plus Network 100.1.3.0 /24.
3. R4 should receive the summary plus all the specific routes.
4. Configure minimum number of "IP Summary-address" commands possible to accomplish this task.
5. R1 should NOT have a discard route for the summary in its routing table.

To summarize in any routing protocol, the specific routes must be known to that routing protocol. The first step is to verify that the Loopback interfaces exist:

On R1

```
R1#Sh run | Inc interface Loopback|ip address 100.

interface Loopback100
 ip address 100.1.0.1 255.255.255.0
interface Loopback101
 ip address 100.1.1.1 255.255.255.0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
interface Loopback102
 ip address 100.1.2.1 255.255.255.0
interface Loopback103
 ip address 100.1.3.1 255.255.255.0
interface Loopback104
 ip address 100.1.4.1 255.255.255.0
interface Loopback105
 ip address 100.1.5.1 255.255.255.0
interface Loopback106
 ip address 100.1.6.1 255.255.255.0
interface Loopback107
 ip address 100.1.7.1 255.255.255.0
```

In the second step, the specific routes are advertised in the routing protocol:

```
R1(config)#Router eigrp 300
R1(config-router)#Netw 100.1.0.1 0.0.0.0
R1(config-router)#Netw 100.1.1.1 0.0.0.0
R1(config-router)#Netw 100.1.2.1 0.0.0.0
R1(config-router)#Netw 100.1.3.1 0.0.0.0
R1(config-router)#Netw 100.1.4.1 0.0.0.0
R1(config-router)#Netw 100.1.5.1 0.0.0.0
R1(config-router)#Netw 100.1.6.1 0.0.0.0
R1(config-router)#Netw 100.1.7.1 0.0.0.0
```

Based on the requirements of this task, a single "summary-address" can be configured for R1's connection to R2; this should satisfy the first requirement.

On R1

```
R1(config)#Int S0/0.12
R1(config-subif)#IP summary-address eigrp 300 100.1.0.0 255.255.248.0
```

To verify the configuration:

On R2

```
R2#Show ip route eigrp 300

     100.0.0.0/21 is subnetted, 1 subnets
D       100.1.0.0 [90/2297856] via 10.1.12.1, 00:06:44, Serial0/0.21
        10.0.0.0/24 is subnetted, 5 subnets
D       10.1.14.0 [90/2681856] via 10.1.12.1, 00:17:38, Serial0/0.21
D       10.1.13.0 [90/2681856] via 10.1.12.1, 00:17:38, Serial0/0.21
```

To accomplish the second and the third requirements of this task, a "Leak-map" MUST be configured,

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

but since the "leak-map" option is NOT available under sub-interfaces, PPP is configured on DLCIs 103 and 104; Remember that the "leak-map" option is available under a virtual-template interface, as follows:

On R1

The IP address is removed so it can be assigned to the virtual-template:

```
R1(config)#Interface S0/0.13
R1(config-subif)#No ip address
```

An access-list is configured to permit the prefix that should be advertised with the summary route to R3:

```
R1(config)#Access-list 1 permit 100.1.3.0 0.0.0.255
```

A route-map is configured to reference the access-list created in the previous step:

```
R1(config)#Route-map TST3 permit 10
R1(config-route-map)#match ip addr 1
```

The virtual-template is configured, the IP address and the summary-address command is configured under the virtual-template:

```
R1(config)#Interface virtual-template 13
R1(config-if)#ip address 10.1.13.1 255.255.255.0
R1(config-if)#IP summary-add eigrp 300 100.1.0.0 255.255.248.0 leak-map TST3

R1(config)#Interface S0/0.13
R1(config-subif)#Frame inter 103 PPP Virtual-template 13
```

Since R1 is running PPP for its connection to R3, R3 should also be configured to run PPP on its connection to R1:

On R3

```
R3(config)#Interface S0/0.31
R3(config-subif)#No ip address

R3(config)#Interface Virtual-Template 31
R3(config-if)#Ip addr 10.1.13.3 255.255.255.0

R3(config)#Interface S0/0.31
R3(config-subif)#Frame inter 301 PPP Virtual-Template 31
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the configuration:

On R1

```
R1#Sh ip route | In D|/32

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       o - ODR, P - periodic downloaded static route
D       100.1.0.0/21 is a summary, 00:28:18, Null0
C       10.1.13.3/32 is directly connected, Virtual-Access2
```

The "D" route is the discard route; this route is injected into the routing table whenever summarization is performed.
The "10.1.13.3 /32" route is injected by IPCP, when ever PPP is configured on a link, a host route for the directly connected host is injected into the routing table for reachability.

```
R3#Show ip route | Inc D|/32

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       o - ODR, P - periodic downloaded static route
D       100.1.0.0/21 [90/2713600] via 10.1.13.1, 00:06:54, Virtual-Access2
D       100.1.3.0/24 [90/2713600] via 10.1.13.1, 00:06:54, Virtual-Access2
D       10.1.14.0/24 [90/4729856] via 10.1.13.1, 00:06:54, Virtual-Access2
C       10.1.13.1/32 is directly connected, Virtual-Access2
D       10.1.12.0/24 [90/4729856] via 10.1.13.1, 00:06:54, Virtual-Access2
```

The output displays the summary route (100.1.0.0/21) and the specific route (100.1.3.0/24). The (10.1.13.1/32) network was injected by PPP's configuration.

To configure the requirement for R4:

On R1

```
R1(config)#Interface S0/0.14
R1(config-subif)#No ip address

R1(config)#Interface Virtual-template 14
R1(config-if)#ip address 10.1.14.1 255.255.255.0
R1(config-if)#IP summary-add Eigrp 300 100.1.0.0 255.255.248.0 Leak-map TST4

R1(config)#Route-map TST4 permit 10

R1(config)#Interface S0/0.14
R1(config-subif)#Frame inter 104 PPP Virtual-template 14
```

On R4

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config)#Interface S0/0.41
R4(config-subif)#No ip address

R4(config)#Interface Virtual-Template 41
R4(config-if)#Ip addr 10.1.14.4 255.255.255.0

R4(config)#Interface S0/0.41
R4(config-subif)#Frame inter 401 PPP Virtual-Template 41
```

To verify the configuration:

On R1

```
R1#Show ip route | Inc D|/32

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       o - ODR, P - periodic downloaded static route
D        100.1.0.0/21 is a summary, 00:03:55, Null0
C        10.1.14.4/32 is directly connected, Virtual-Access3
C        10.1.13.3/32 is directly connected, Virtual-Access2
```

On R4

```
R4#Show ip route | Inc D|/32

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       o - ODR, P - periodic downloaded static route
D        100.1.5.0/24 [90/2713600] via 10.1.14.1, 00:01:42, Virtual-Access2
D        100.1.4.0/24 [90/2713600] via 10.1.14.1, 00:01:42, Virtual-Access2
D        100.1.7.0/24 [90/2713600] via 10.1.14.1, 00:01:42, Virtual-Access2
D        100.1.6.0/24 [90/2713600] via 10.1.14.1, 00:01:42, Virtual-Access2
D        100.1.1.0/24 [90/2713600] via 10.1.14.1, 00:01:42, Virtual-Access2
D        100.1.0.0/24 [90/2713600] via 10.1.14.1, 00:01:42, Virtual-Access2
D        100.1.0.0/21 [90/2713600] via 10.1.14.1, 00:01:42, Virtual-Access2
D        100.1.3.0/24 [90/2713600] via 10.1.14.1, 00:01:42, Virtual-Access2
D        100.1.2.0/24 [90/2713600] via 10.1.14.1, 00:01:42, Virtual-Access2
D EX 33.0.0.0/8 [170/2300416] via 10.1.45.5, 01:25:37, Serial0/0.45
D EX 5.0.0.0/8 [170/2297856] via 10.1.45.5, 01:25:37, Serial0/0.45
D EX 6.0.0.0/8 [170/2300416] via 10.1.45.5, 01:25:37, Serial0/0.45
C        10.1.14.1/32 is directly connected, Virtual-Access2
D        10.1.13.3/32 [90/5145600] via 10.1.14.1, 00:01:43, Virtual-Access2
D        10.1.13.0/24 [90/5145600] via 10.1.14.1, 00:01:43, Virtual-Access2
D        10.1.12.0/24 [90/4729856] via 10.1.14.1, 00:01:43, Virtual-Access2
D EX     10.1.56.0/24 [170/2172416] via 10.1.45.5, 01:25:39, Serial0/0.45
```

To remove the NULL0 route:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

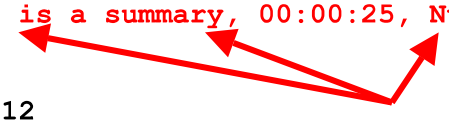
Note whenever summarization is performed in Eigrp, OSPF or BGP a discard route (The summary route to Null0) is created automatically, this route is injected into the routing table of the router that performed summarization and it's used as a forwarding loop prevention. To Remove the discard route in Eigrp, the summary route is configured with an administrative distance of 255, as follows:

On R1

```
R1#Show ip route | Inc Null0

D       100.1.0.0/21 is a summary, 00:00:25, Null0


R1(config)#int s0/0.12
R1(config-subif)#ip summary-address eigrp 300 100.1.0.0 255.255.248.0 255
```





The administrative dstance

To verify the configuration:

On R1

You should see the following console message after entering the above command:

%DUAL-5-NBRCHANGE: IP-EIGRP(0) 300: Neighbor 10.1.12.2 (Serial0/0.12) is resync: summary configured

You should clear the routing table for the changes to take effect:

```
R1#clear ip route *

R1#Show ip route | Inc Null0
R1#
```

Task 6

R1 – R4 should be configured to advertise their Lo0 interface in AS 300.

On R1

```
R1(config)#Router eigrp 300
R1(config-router)#Netw 1.1.1.1 0.0.0.0
```

On R2

```
R2(config)#Router eigrp 300
R2(config-router)#Netw 2.2.2.2 0.0.0.0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3

```
R3(config)#Router eigrp 300
R3(config-router)#Netw 3.3.3.3 0.0.0.0
```

On R4

```
R4(config)#Router eigrp 300
R4(config-router)#Netw 4.4.4.4 0.0.0.0
```

To verify the configuration:

On R1

```
R1#Show ip route eigrp 300

D    2.0.0.0/8 [90/2297856] via 10.1.12.2, 00:01:35, Serial0/0.12
D    3.0.0.0/8 [90/2713600] via 10.1.13.3, 00:01:13, Virtual-Access2
D    4.0.0.0/8 [90/2713600] via 10.1.14.4, 00:00:33, Virtual-Access3
```

On R2

```
R2#Show ip route eigrp 300

D    1.0.0.0/8 [90/2297856] via 10.1.12.1, 00:02:30, Serial0/0.21
     100.0.0.0/21 is subnetted, 1 subnets
D       100.1.0.0 [90/2297856] via 10.1.12.1, 00:12:41, Serial0/0.21
D    3.0.0.0/8 [90/4857856] via 10.1.12.1, 00:01:45, Serial0/0.21
D    4.0.0.0/8 [90/4857856] via 10.1.12.1, 00:01:06, Serial0/0.21
     10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
D       10.1.14.4/32 [90/4729856] via 10.1.12.1, 00:22:55, Serial0/0.21
D       10.1.14.0/24 [90/4729856] via 10.1.12.1, 00:22:55, Serial0/0.21
D       10.1.13.3/32 [90/4729856] via 10.1.12.1, 00:36:53, Serial0/0.21
D       10.1.13.0/24 [90/4729856] via 10.1.12.1, 00:36:53, Serial0/0.21
```

On R3

```
R3#Show ip route eigrp 300

D    1.0.0.0/8 [90/2713600] via 10.1.13.1, 00:03:05, Virtual-Access2
D    2.0.0.0/8 [90/4857856] via 10.1.13.1, 00:02:42, Virtual-Access2
     100.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
D       100.1.0.0/21 [90/2713600] via 10.1.13.1, 00:13:16, Virtual-Access2
D       100.1.3.0/24 [90/2713600] via 10.1.13.1, 00:37:25, Virtual-Access2
D    4.0.0.0/8 [90/5273600] via 10.1.13.1, 00:01:40, Virtual-Access2
     10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
D       10.1.14.4/32 [90/5145600] via 10.1.13.1, 00:23:30, Virtual-Access2
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
D       10.1.14.0/24 [90/5145600] via 10.1.13.1, 00:23:30, Virtual-Access2
D       10.1.12.0/24 [90/4729856] via 10.1.13.1, 00:37:25, Virtual-Access2
```

On R4

```
R4#Show ip route eigrp 300

D    1.0.0.0/8 [90/2713600] via 10.1.14.1, 00:03:26, Virtual-Access2
D    2.0.0.0/8 [90/4857856] via 10.1.14.1, 00:03:03, Virtual-Access2
     100.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
D       100.1.5.0/24 [90/2713600] via 10.1.14.1, 00:23:45, Virtual-Access2
D       100.1.4.0/24 [90/2713600] via 10.1.14.1, 00:23:45, Virtual-Access2
D       100.1.7.0/24 [90/2713600] via 10.1.14.1, 00:23:45, Virtual-Access2
D       100.1.6.0/24 [90/2713600] via 10.1.14.1, 00:23:45, Virtual-Access2
D       100.1.1.0/24 [90/2713600] via 10.1.14.1, 00:23:45, Virtual-Access2
D       100.1.0.0/24 [90/2713600] via 10.1.14.1, 00:23:45, Virtual-Access2
D       100.1.0.0/21 [90/2713600] via 10.1.14.1, 00:13:36, Virtual-Access2
D       100.1.3.0/24 [90/2713600] via 10.1.14.1, 00:23:45, Virtual-Access2
D       100.1.2.0/24 [90/2713600] via 10.1.14.1, 00:23:45, Virtual-Access2
D    3.0.0.0/8 [90/5273600] via 10.1.14.1, 00:02:40, Virtual-Access2
     10.0.0.0/8 is variably subnetted, 10 subnets, 2 masks
D       10.1.13.3/32 [90/5145600] via 10.1.14.1, 00:23:46, Virtual-Access2
D       10.1.13.0/24 [90/5145600] via 10.1.14.1, 00:23:46, Virtual-Access2
D       10.1.12.0/24 [90/4729856] via 10.1.14.1, 00:23:46, Virtual-Access2
```

Task 7

Configure a mutual redistribution between the two ASes on R4.

On R4

To verify the two ASes running on this router:

```
R4#Show ip protocol | Inc Routing Protocol

Routing Protocol is "eigrp 200"
Routing Protocol is "eigrp 300"
```

To redistribute between the two ASes:

```
R4(config)#router eigrp 200
R4(config-router)#redistribute eigrp 300

R4(config)#router eigrp 300
R4(config-router)#redistribute eigrp 200
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the configuration:

On R1

```
R1#Show ip route eigrp

D    2.0.0.0/8 [90/2297856] via 10.1.12.2, 00:11:15, Serial0/0.12
D EX 33.0.0.0/8 [170/4860416] via 10.1.14.4, 00:01:19, Virtual-Access3
D    3.0.0.0/8 [90/2713600] via 10.1.13.3, 00:10:52, Virtual-Access2
D    4.0.0.0/8 [90/2713600] via 10.1.14.4, 00:10:13, Virtual-Access3
D EX 5.0.0.0/8 [170/4857856] via 10.1.14.4, 00:01:19, Virtual-Access3
D EX 6.0.0.0/8 [170/4860416] via 10.1.14.4, 00:01:19, Virtual-Access3
     10.0.0.0/8 is variably subnetted, 10 subnets, 2 masks
D EX    10.1.46.0/24 [170/4729856] via 10.1.14.4, 00:01:19, Virtual-Access3
D EX    10.1.45.0/24 [170/4729856] via 10.1.14.4, 00:01:19, Virtual-Access3
D EX    10.1.56.0/24 [170/4732416] via 10.1.14.4, 00:01:19, Virtual-Access3
```

Task 8

Configure R2 to summarize its Loopback 134 and 135 and advertise a single summary in AS 300.

On R2

```
R2#Sh ip int br | Inc Loopback13[4|5]

Loopback134                102.1.34.2      YES manual up                    up
Loopback135                102.1.35.2      YES manual up                    up
```

To configure the summary:

On R2

```
R2(config)#Router eigrp 300
R2(config-router)#Netw 102.1.35.2 0.0.0.0
R2(config-router)#Netw 102.1.34.2 0.0.0.0
```

To summarize:

We can see that the first two octets of these two prefixes match (102.1.), the third octet is where they differ, if the third octet of these two prefixes are converted into binary we should see the following:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

	128	64	32	16	8	4	2	1
34	0	0	1	0	0	0	1	0
35	0	0	1	0	0	0	1	1
M	M	M	M	M	M	M	DM	

If the last bit of these two prefixes are compared, you will see that the last bit can be either a "0" or a "1" which covers all possibilities; therefore, the Don't care bit is the last bit; as a result, the summary route will be 102.1.34.0 0.0.1.0 or 255.255.254.0:

```
R2(config)#Int S0/0.21
R2(config-subif)#IP summary-address eigrp 300 102.1.34.0 255.255.254.0
```

To verify the configuration:

On R1

```
R1#Show ip rout eigrp 300 | Inc 102.

     102.0.0.0/23 is subnetted, 1 subnets
D       102.1.34.0 [90/2297856] via 10.1.12.2, 00:00:22, Serial0/0.12
```

To test the configuration:

On R1

```
R1#Ping 102.1.34.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 102.1.34.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/49/52 ms

R1#Ping 102.1.35.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 102.1.35.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/49/52 ms
```

Task 9

Configure Authentication on R1, R2 and R4 using the following policy:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

- The connection between R1 and R2 should be authenticated using “R1R2” as the password.
- The connection between R1 and R4 should be authenticated using “R1R4” as the password.

In the first step, R1 and R2 are configured for authentication; note the key chain names are locally significant:

On R1

```
R1(config)#Key chain R1R2
R1(config-keychain)#key 1
R1(config-keychain-key)#key-string R1R2

R1(config)#Int s0/0.12
R1(config-subif)#IP authentication key-chain eigrp 300 R1R2
R1(config-subif)#IP authentication mode eigrp 300 MD5
```

Note the adjacency is dropped:

To verify neighbor adjacency:

On R1

```
R1#Show ip eigrp neighbors

IP-EIGRP neighbors for process 300
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
2   10.1.14.4               Vi3               11 01:22:00   62   372  0  33
1   10.1.13.3               Vi2               14 01:35:58   56   336  0  33
```

On R2

```
R2(config)#Key chain R2R1
R2(config-keychain)#Key 1
R2(config-keychain-key)#Key-string R1R2

R2(config)#Int S0/0.21
R2(config-subif)#IP authentication key-chain eigrp 300 R2R1
R2(config-subif)#IP authentication mode eigrp 300 MD5
```

You should get the following console message:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
IP-EIGRP(0) 300: Neighbor 10.1.12.1 (Serial0/0.21) is up: new adjacency
```

To verify the configuration:

On R2

```
R2#Show ip eigrp neighbors

IP-EIGRP neighbors for process 300
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   10.1.12.1               Se0/0.21          11 00:00:27 1053  5000  0  109
```

In the last item of this task, authentication is configured between R1 and R4; note authentication is applied to the Virtual-template interface and NOT the sub-interface or the main interface:

On R1

```
R1(config)#Key chain R1R4
R1(config-keychain)#key 1
R1(config-keychain-key)#Key-string R1R4

R1(config)#int Virtual-template 14
R1(config-if)#IP authentication key-chain eigrp 300 R1R4
R1(config-if)#IP authentication mode eigrp 300 MD5
```

On R4

You have received the following console message on R4:

```
Neighbor 10.1.14.1 (Virtual-Access1) is down: Auth failure
```

```
R4(config)#Key chain R4R1
R4(config-keychain)#Key 1
R4(config-keychain-key)#Key-String R1R4

R4(config)#Int Virtual-template 41
R4(config-if)#IP Authentication key-chain eigrp 300 R4R1
R4(config-if)#IP authentication mode eigrp 300 MD5
```

You should get the following console message:

```
Neighbor 10.1.14.1 (Virtual-Access1) is up: new adjacency
```

To verify the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4

```
R4#Show ip eigrp neighbors

IP-EIGRP neighbors for process 200
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   10.1.46.6               Se0/0.46          14 02:57:10  160   960  0  16
0   10.1.45.5               Se0/0.45          10 02:57:29  203  1218  0  17

IP-EIGRP neighbors for process 300
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   10.1.14.1               Vi2               10 00:00:31 1118  5000  0  125
```

Task 10

Configure the S0/1 interface of R1 and R3 in Eigrp AS 500 using the following policy:

1. R1 should be configured to advertise its loopback 130 using a route-map such that a "Show ip route eigrp 500" on R3 resembles the following, DO NOT use an access-list or a prefix-list to accomplish this task:

   D EX 130.1.1.0 [170/2297856] via 10.1.100.1, 00:00:21, Serial0/1

2. R1 should also advertise its Loopback 1000 in this routing domain such that the network is injected as an external route, you should use an access-list and a route-map to accomplish this task.

3. R3 should be configured to advertise its loopback 310 such that a "Show ip route eigrp 300" on R1 resembles the following:

   D EX 130.3.0.0/16 [170/2297856] via 10.1.100.3, 00:00:44, Serial0/1

   R3 should use "Redistribute static" to accomplish this task. DO NOT configure a Static route to accomplish this task:

The neighbor adjacency of these two routers should be configured first:

On R1

```
R1(config)#router eigrp 500
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config-router)#No au
R1(config-router)#netw 10.1.100.1 0.0.0.0
```

On R3

```
R3(config)#router eigrp 500
R3(config-router)#No au
R3(config-router)#Netw 10.1.100.3 0.0.0.0
```

To verify the configuration:

On R3

```
R3#Show ip eigrp neighbors 500

IP-EIGRP neighbors for process 500
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   10.1.100.1              Se0/1             10 00:00:13    2   200  0  3
```

On R1

```
R1#Show ip eigrp neighbors 500

IP-EIGRP neighbors for process 500
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   10.1.100.3              Se0/1             10 00:01:00    1   200  0  3
```

Now that these two routers are neighbors we can proceed with the rest of the policies of this task:

On R1

```
R1(config)#Route-map TASK8 permit 10
R1(config-route-map)#match int lo130

R1(config-route-map)#Router eigrp 500
R1(config-router)#redistribute connected route-map TASK8
```

To verify the configuration:

On R3

```
R3#Show ip route eigrp 500
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
     130.1.0.0/24 is subnetted, 1 subnets
D EX    130.1.1.0 [170/2297856] via 10.1.100.1, 00:00:42, Serial0/1
```

To accomplish the second item:
An access-list is configured, BUT before configuring any access-list a show access-list is highly recommended so the access-list number of the existing access-list is NOT used in the new one:

On R1

```
R1#Show access-list

Standard IP access list 1
    10 permit 100.1.3.0, wildcard bits 0.0.0.255 (2 matches)

R1(config)#access-list 2 permit 130.100.100.0 0.0.0.255
```

Since we can NOT have multiple redistribute connected (The new one will override the existing one), the access-list is referenced in the same route-map:

```
R1(config)#route-map TASK8 permit 20
R1(config-route-map)#match ip addr 2
```

To see the running configuration:

On R1

```
route-map TASK8 permit 10
 match interface Loopback130
!
route-map TASK8 permit 20
 match ip address 2
```

(The output is modified to show the pertaining section)

To verify the configuration:

On R3

```
R3#Show ip route eigrp 500

     130.1.0.0/24 is subnetted, 1 subnets
D EX    130.1.1.0 [170/2297856] via 10.1.100.1, 00:09:38, Serial0/1
     130.100.0.0/24 is subnetted, 1 subnets

D EX    130.100.100.0 [170/2297856] via 10.1.100.1, 00:08:37, Serial0/1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To accomplish the last item:

This task can be tricky, but since a static route can NOT be used, the behavior of a command is used to accomplish this task. If the "IP default-network" command references a subnetted network, IOS will automatically inject a static route, let's test this:

On R3

```
R3(config)#IP default-network 130.3.3.0
```

To verify the configuration:

On R3

```
R3#Show run | Inc ip route

ip route 130.3.0.0 255.255.0.0 130.3.3.0

R3#Show ip route static

     130.3.0.0/16 is variably subnetted, 2 subnets, 2 masks
S       130.3.0.0/16 [1/0] via 130.3.3.0
```

Now that we see the default route automatically configured, we can use the redistribute static command to accomplish this task, as follows:

On R3

```
R3(config)#router eigrp 500
R3(config-router)#redistribute static
```

To verify the configuration:

On R1

```
R1#Show ip route eigrp 500

D EX 130.3.0.0/16 [170/2297856] via 10.1.100.3, 00:00:44, Serial0/1
```

Task 11

Configure R2 and R4's F0/1 and the F0/0 interface of R4 in AS 400.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2

```
R2(config)#router eigrp 400
R2(config-router)#NO au
R2(config-router)#netw 10.1.24.2 0.0.0.0
```

On R4

```
R4(config)#router eigrp 400
R4(config-router)#NO au
R4(config-router)#netw 10.1.24.4 0.0.0.0
R4(config-router)#netw 10.1.4.4 0.0.0.0
```

To verify the configuration:

On R2

```
R2#Show ip route eigrp 400

     10.0.0.0/8 is variably subnetted, 11 subnets, 2 masks
D       10.1.4.0/24 [90/30720] via 10.1.24.4, 00:01:45, FastEthernet0/1
```

Task 12

Redistribute AS 400 into the ASes configured on R2 and R4.

On R4

```
R4#Sh ip proto | Inc Routing Protocol

Routing Protocol is "eigrp 200"
Routing Protocol is "eigrp 300"
Routing Protocol is "eigrp 400"

R4(config)#router eigrp 200
R4(config-router)#redistribute eigrp 400

R4(config)#router eigrp 300

R4(config-router)#redistribute eigrp 400
```

On R2

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2#Sh ip protocol | Inc Routing Protocol

Routing Protocol is "eigrp 300"
Routing Protocol is "eigrp 400"

R2(config)#Router eigrp 300
R2(config-router)#redistribute eigrp 400
```

To verify the configuration:

On BB3

```
BB3#Show ip route eigrp | Inc EX

D EX    102.1.34.0 [170/5372416] via 10.1.56.5, 00:14:18, FastEthernet0/0
D EX 1.0.0.0/8 [170/4860416] via 10.1.56.5, 00:14:18, FastEthernet0/0
D EX 2.0.0.0/8 [170/5372416] via 10.1.56.5, 00:14:18, FastEthernet0/0
D EX    100.1.5.0/24 [170/4860416] via 10.1.56.5, 00:14:18, FastEthernet0/0
D EX    100.1.4.0/24 [170/4860416] via 10.1.56.5, 00:14:18, FastEthernet0/0
D EX    100.1.7.0/24 [170/4860416] via 10.1.56.5, 00:14:18, FastEthernet0/0
D EX    100.1.6.0/24 [170/4860416] via 10.1.56.5, 00:14:18, FastEthernet0/0
D EX    100.1.1.0/24 [170/4860416] via 10.1.56.5, 00:14:18, FastEthernet0/0
D EX    100.1.0.0/24 [170/4860416] via 10.1.56.5, 00:14:18, FastEthernet0/0
D EX    100.1.0.0/21 [170/4860416] via 10.1.56.5, 00:14:18, FastEthernet0/0
D EX    100.1.3.0/24 [170/4860416] via 10.1.56.5, 00:14:18, FastEthernet0/0
D EX    100.1.2.0/24 [170/4860416] via 10.1.56.5, 00:14:18, FastEthernet0/0
D EX 3.0.0.0/8 [170/7420416] via 10.1.56.5, 00:14:18, FastEthernet0/0
D EX 4.0.0.0/8 [170/2300416] via 10.1.56.5, 00:14:18, FastEthernet0/0
D EX    10.1.14.1/32 [170/4732416] via 10.1.56.5, 00:14:19, FastEthernet0/0
D EX    10.1.14.0/24 [170/4732416] via 10.1.56.5, 00:14:19, FastEthernet0/0
D EX    10.1.13.3/32 [170/7292416] via 10.1.56.5, 00:14:19, FastEthernet0/0
D EX    10.1.13.0/24 [170/7292416] via 10.1.56.5, 00:14:19, FastEthernet0/0
D EX    10.1.12.0/24 [170/5244416] via 10.1.56.5, 00:14:19, FastEthernet0/0
D EX    10.1.4.0/24 [170/2174976] via 10.1.56.5, 00:02:04, FastEthernet0/0
D EX    10.1.24.0/24 [170/2174976] via 10.1.56.5, 00:02:04, FastEthernet0/0
D EX    10.1.46.0/24 [170/2684416] via 10.1.56.5, 00:14:19, FastEthernet0/0
D EX    10.1.45.0/24 [170/2172416] via 10.1.56.5, 05:36:17, FastEthernet0/0
```

Task 13

Configure the F0/1 interface of BB1 and BB2 in AS 600 and advertise their Lo0 interface in this AS such that they appear as external routes.

On BB1

```
BB1(config)#router eigrp 600
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
BB1(config-router)#NO au
BB1(config-router)#Netw 112.112.112.11 0.0.0.0

BB1(config)#route-map TST permit 10
BB1(config-route-map)#match int lo0

BB1(config-route-map)#router eigrp 600
BB1(config-router)#redistribute connected route-map TST
```

On BB2

```
BB2(config)#router eigrp 600
BB2(config-router)#NO au
BB2(config-router)#netw 112.112.112.22 0.0.0.0

BB2(config)#route-map TST permit 10
BB2(config-route-map)#match int lo0

BB2(config-route-map)#router eigrp 600
BB2(config-router)#redistribute connected route-map TST
```

To verify the configuration:

On BB1

```
BB1#Show ip route eigrp

     112.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D EX    112.2.2.22/32
          [170/156160] via 112.112.112.22, 00:01:47, FastEthernet0/1
```

On BB2

```
BB2#Show ip route eigrp

     112.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D EX    112.1.1.11/32
          [170/156160] via 112.112.112.11, 00:02:53, FastEthernet0/1
```

Task 14

Configure the F0/0 interface of BB2, R2, BB1 and R1 in AS 600.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1

```
R1(config)#router eigrp 600
R1(config-router)#NO au
R1(config-router)#netw 10.11.11.1 0.0.0.0
```

On R2

```
R2(config)#router eigrp 600
R2(config-router)#NO au
R2(config-router)#Netw 10.11.12.2 0.0.0.0
```

On BB1

```
BB1(config)#router eigrp 600
BB1(config-router)#Netw 10.11.11.11 0.0.0.0
```

On BB2

```
BB2(config)#router eigrp 600
BB2(config-router)#Netw 10.11.12.22 0.0.0.0
```

To verify the configuration;

On R1

```
R1#Show ip route eigrp 600

      112.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D EX     112.2.2.22/32 [170/158720] via 10.11.11.11, 00:01:59, FastEthernet0/0
D        112.112.112.0/24 [90/30720] via 10.11.11.11, 00:01:59, FastEthernet0/0
D EX     112.1.1.11/32 [170/156160] via 10.11.11.11, 00:01:59, FastEthernet0/0
     10.0.0.0/8 is variably subnetted, 13 subnets, 2 masks
D        10.11.12.0/24 [90/33280] via 10.11.11.11, 00:00:59, FastEthernet0/0
```

On R2

```
R2#Show ip route eigrp 600

     112.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D EX     112.2.2.22/32 [170/156160] via 10.11.12.22, 00:01:44, FastEthernet0/0
D        112.112.112.0/24 [90/30720] via 10.11.12.22, 00:01:44, FastEthernet0/0
D EX     112.1.1.11/32 [170/158720] via 10.11.12.22, 00:01:44, FastEthernet0/0

     10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
D        10.11.11.0/24 [90/33280] via 10.11.12.22, 00:01:44, FastEthernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 15

Redistribute mutually between AS 600 and the Eigrp ASes running on R1, and R2.

On R1

```
R1#Sh ip protocol | Inc Routing Protocol

Routing Protocol is "eigrp 300"
Routing Protocol is "eigrp 500"
Routing Protocol is "eigrp 600"

R1(config)#Router eigrp 300
R1(config-router)#redistribute eigrp 600

R1(config)#Router eigrp 500
R1(config-router)#redistribute eigrp 600

R1(config)#Router eigrp 600
R1(config-router)#redistribute eigrp 500
R1(config-router)#redistribute eigrp 300
```

On R2

```
R2#Show ip protocol | Inc Routing Protocol

Routing Protocol is "eigrp 300"
Routing Protocol is "eigrp 400"
Routing Protocol is "eigrp 600"

R2(config)#Router eigrp 300
R2(config-router)#redistribute eigrp 600

R2(config)#Router eigrp 400
R2(config-router)#redistribute eigrp 600

R2(config)#router eigrp 600
R2(config-router)#redistribute eigrp 300
R2(config-router)#redistribute eigrp 400
```

To verify the configuration:

On BB3

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
BB3#Show ip route eigrp | Inc EX

102.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D EX    102.1.35.0/24 [170/4868096] via 10.1.56.5, 00:02:49, FastEthernet0/0
D EX    102.1.34.0/24 [170/4868096] via 10.1.56.5, 00:02:49, FastEthernet0/0
D EX    102.1.34.0/23 [170/5372416] via 10.1.56.5, 01:53:12, FastEthernet0/0
D EX 1.0.0.0/8 [170/4860416] via 10.1.56.5, 01:53:12, FastEthernet0/0
D EX 2.0.0.0/8 [170/5372416] via 10.1.56.5, 01:53:12, FastEthernet0/0
     100.0.0.0/24 is subnetted, 8 subnets
D EX    100.1.5.0 [170/4860416] via 10.1.56.5, 01:53:12, FastEthernet0/0
D EX    100.1.4.0 [170/4860416] via 10.1.56.5, 01:53:12, FastEthernet0/0
D EX    100.1.7.0 [170/4860416] via 10.1.56.5, 01:53:12, FastEthernet0/0
D EX    100.1.6.0 [170/4860416] via 10.1.56.5, 01:53:12, FastEthernet0/0
D EX    100.1.1.0 [170/4860416] via 10.1.56.5, 01:53:12, FastEthernet0/0
D EX    100.1.0.0 [170/4860416] via 10.1.56.5, 01:53:12, FastEthernet0/0
D EX    100.1.3.0 [170/4860416] via 10.1.56.5, 01:53:12, FastEthernet0/0
D EX    100.1.2.0 [170/4860416] via 10.1.56.5, 01:53:13, FastEthernet0/0
D EX 3.0.0.0/8 [170/7420416] via 10.1.56.5, 01:53:13, FastEthernet0/0
D EX 4.0.0.0/8 [170/2300416] via 10.1.56.5, 01:53:13, FastEthernet0/0
D    5.0.0.0/8 [90/156160] via 10.1.56.5, 07:27:10, FastEthernet0/0
D    6.0.0.0/8 [90/158720] via 10.1.56.5, 07:24:14, FastEthernet0/0
     112.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D EX    112.2.2.22/32 [170/2305536] via 10.1.56.5, 00:34:16, FastEthernet0/0
D EX    112.112.112.0/24
           [170/2180096] via 10.1.56.5, 00:34:16, FastEthernet0/0
D EX    112.1.1.11/32 [170/2308096] via 10.1.56.5, 00:34:18, FastEthernet0/0
D EX 130.3.0.0/16 [170/2822656] via 10.1.56.5, 00:04:14, FastEthernet0/0
     10.0.0.0/8 is variably subnetted, 13 subnets, 2 masks
D EX    10.1.14.1/32 [170/4732416] via 10.1.56.5, 01:53:14, FastEthernet0/0
D EX    10.1.14.0/24 [170/4732416] via 10.1.56.5, 01:53:14, FastEthernet0/0
D EX    10.1.13.3/32 [170/7292416] via 10.1.56.5, 01:53:14, FastEthernet0/0
D EX    10.1.13.0/24 [170/7292416] via 10.1.56.5, 01:53:14, FastEthernet0/0
D EX    10.1.12.0/24 [170/5244416] via 10.1.56.5, 01:53:14, FastEthernet0/0
D EX    10.11.11.0/24 [170/2182656] via 10.1.56.5, 00:34:18, FastEthernet0/0
D EX    10.11.12.0/24 [170/2177536] via 10.1.56.5, 00:34:18, FastEthernet0/0
D EX    10.1.4.0/24 [170/2174976] via 10.1.56.5, 01:40:59, FastEthernet0/0
D EX    10.1.24.0/24 [170/2174976] via 10.1.56.5, 01:40:59, FastEthernet0/0
D EX    10.1.46.0/24 [170/2684416] via 10.1.56.5, 01:53:14, FastEthernet0/0
D EX    10.1.45.0/24 [170/2172416] via 10.1.56.5, 07:15:13, FastEthernet0/0
D EX    10.1.100.0/24 [170/2694656] via 10.1.56.5, 00:04:15, FastEthernet0/0
```

Task 16

Configure the routers in AS 500 such that they log neighbor warning messages and repeat the warning messages every 10 minutes. You should disable logging of neighbor changes for this AS.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1 and R3

```
Rx(config)#Router eigrp 500
Rx(config-router)#Eigrp log-neighbor-warnings 600
Rx(config-router)#NO Eigrp log-neighbor-changes
```

To verify the configuration:

On R1

```
R1#Sh run | s router eigrp 500

router eigrp 500
 redistribute connected route-map TASK8-1
 redistribute eigrp 600
 network 10.1.100.1 0.0.0.0
 no auto-summary
 no eigrp log-neighbor-changes
 eigrp log-neighbor-warnings 600
```

Task 17

Configure BB1 and BB2 in AS 600 such that they NEVER exchange existing and/or future external Eigrp routes that they originate, DO NOT configure an Access-list, Prefix-list or a route-map to accomplish this task.

On BB1

```
BB1#Show ip route 112.2.2.22

Routing entry for 112.2.2.22/32
  Known via "eigrp 600", distance 170, metric 156160, type external
  Redistributing via eigrp 600
  Last update from 112.112.112.22 on FastEthernet0/1, 00:00:06 ago
  Routing Descriptor Blocks:
  * 112.112.112.22, from 112.112.112.22, 00:00:06 ago, via FastEthernet0/1
      Route metric is 156160, traffic share count is 1

      Total delay is 5100 microseconds, minimum bandwidth is 100000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
```

Note, BB1 can reach 112.2.2.22/32 via BB2 (112.112.112.22).

On BB1 and BB2

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
BBx(config)#router eigrp 600
BBx(config-router)#Eigrp router-id 11.11.11.11
```

To verify the configuration:

On BB1

```
BB1#Show ip route 112.2.2.22

Routing entry for 112.2.2.22/32
  Known via "eigrp 600", distance 170, metric 2302976, type external
  Redistributing via eigrp 600
  Last update from 10.11.11.1 on FastEthernet0/0, 00:02:29 ago
  Routing Descriptor Blocks:
  * 10.11.11.1, from 10.11.11.1, 00:02:29 ago, via FastEthernet0/0
      Route metric is 2302976, traffic share count is 1
      Total delay is 25200 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 4/255, Hops 3
```

On BB2

```
BB2#Sh ip route 112.1.1.11

Routing entry for 112.1.1.11/32
  Known via "eigrp 600", distance 170, metric 2302976, type external
  Redistributing via eigrp 600
  Last update from 10.11.12.2 on FastEthernet0/0, 00:05:00 ago
  Routing Descriptor Blocks:
  * 10.11.12.2, from 10.11.12.2, 00:05:00 ago, via FastEthernet0/0
      Route metric is 2302976, traffic share count is 1
      Total delay is 25200 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 3
```

When routes are redistributed into Eigrp, the external routes are tagged with the router-id of the router that originated the route/s.
This can be seen in the topology table:

On BB2

```
BB2#Sh ip eigrp top 112.2.2.22/32 | Inc 11.11.11.11

        Originating router is 11.11.11.11 (this system)
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The receiving router sees the router-id in the updates and if it matches its own router-id, it will reject the routes.

Task 18

Configure a frame-relay Point-to-point connection between R3 and R4. R3 should be configured with 10.1.34.3/24 and R4 should be configured with 10.1.34.4/24. These routers should advertise this IP address in Eigrp AS 300.

On R3

```
R3(config)#int s0/0.34 p
R3(config-subif)#ip addr 10.1.34.3 255.255.255.0
R3(config-subif)#frame-relay interface-dlci 304

R3(config)#router eigrp 300
R3(config-router)#netw 10.1.34.3 0.0.0.0
```

On R4

```
R4(config)#int s0/0.43 p
R4(config-subif)#ip addr 10.1.34.4 255.255.255.0
R4(config-subif)#frame-relay interface-dlci 403

R4(config-fr-dlci)#router eigrp 300
R4(config-router)#netw 10.1.34.4 0.0.0.0
```

To verify the configuration:

On R1

```
R1#Sh ip route eigrp 300

     102.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D       102.1.34.0/23 [90/2297856] via 10.1.12.2, 16:38:08, Serial0/0.12
D    2.0.0.0/8 [90/2297856] via 10.1.12.2, 16:38:08, Serial0/0.12

D EX 33.0.0.0/8 [170/4860416] via 10.1.14.4, 00:05:58, Virtual-Access3
D    3.0.0.0/8 [90/2713600] via 10.1.13.3, 00:05:59, Virtual-Access2
D    4.0.0.0/8 [90/2713600] via 10.1.14.4, 00:05:58, Virtual-Access3
D EX 5.0.0.0/8 [170/4857856] via 10.1.14.4, 00:05:58, Virtual-Access3
D EX 6.0.0.0/8 [170/4860416] via 10.1.14.4, 00:05:58, Virtual-Access3
     112.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D EX    112.2.2.22/32 [170/2300416] via 10.1.12.2, 00:05:58, Serial0/0.12
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
     10.0.0.0/8 is variably subnetted, 14 subnets, 2 masks
D EX    10.1.46.0/24 [170/4729856] via 10.1.14.4, 00:05:58, Virtual-Access3
D EX    10.1.45.0/24 [170/4729856] via 10.1.14.4, 00:05:58, Virtual-Access3
D       10.1.34.0/24 [90/4729856] via 10.1.14.4, 00:06:01, Virtual-Access3
                     [90/4729856] via 10.1.13.3, 00:06:01, Virtual-Access2
D EX    10.1.56.0/24 [170/4732416] via 10.1.14.4, 00:06:00, Virtual-Access3
```

Task 19

R1 should be configured to load balance in a ratio of 1:3 for network 10.1.34.0 /24.

When we see load balancing and ratios, we immediately think of offset-list. Remember that offset-list only affects the advertised distance and if we add to the advertised distance the prefix may NOT even meet the Feasibility Condition. Therefore, the delay value should be manipulated to achieve this ratio, the following steps outlines the formula and steps to accomplish this task:

The first step:
We should see the vector metric from R1's perspective:

On R1

```
R1#Show ip eigrp topology 10.1.34.0 255.255.255.0

IP-EIGRP (AS 300): Topology entry for 10.1.34.0/24
  State is Passive, Query origin flag is 1, 2 Successor(s), FD is 4729856
  Routing Descriptor Blocks:
  10.1.13.3 (Virtual-Access2), from 10.1.13.3, Send flag is 0x0
      Composite metric is (4729856/2169856), Route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 120000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
  10.1.14.4 (Virtual-Access3), from 10.1.14.4, Send flag is 0x0
      Composite metric is (4729856/2169856), Route is Internal
      Vector metric:

        Minimum bandwidth is 1544 Kbit
        Total delay is 120000 microseconds
        Reliability is 255/255
        Load is 5/255
        Minimum MTU is 1500
        Hop count is 1
(The rest of the output is omitted)
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Second step:
Lets check the routing table of R1 for this prefix:

On R1

```
R1#Show ip route 10.1.34.0

Routing entry for 10.1.34.0/24
  Known via "eigrp 300", distance 90, metric 4729856, type internal

  Redistributing via eigrp 300, eigrp 600
  Advertised by eigrp 600
  Last update from 10.1.14.4 on Virtual-Access3, 00:06:24 ago
  Routing Descriptor Blocks:
    10.1.14.4, from 10.1.14.4, 00:06:24 ago, via Virtual-Access3
      Route metric is 4729856, traffic share count is 1
      Total delay is 120000 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 5/255, Hops 1
  * 10.1.13.3, from 10.1.13.3, 00:06:24 ago, via Virtual-Access2
      Route metric is 4729856, traffic share count is 1
      Total delay is 120000 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
```

We can see that R1 is performing an equal cost load balancing using a ratio of 1:1. If we need to change the ratio to 1:3, we should take the composite metric and multiply it by 3. It means that if the composite metric of R4 is three times as much as the composite metric of R3 then, R1 can do load balancing using a ratio of 1:3:

```
4729856 * 3 = 14189568
```

How do we arrive to the above value?
We know that the formula that Eigrp uses is:

```
Composite value = 256 (BW + DLY)
```

We also know that offset-list will NOT work. Therefore, we must manipulate the delay value of the virtual-template on R1.
The composite value should equal to 256 multiplied by the result of BW + Delay + X, where "X" is what we need to add to the total delay value to get the desired composite metric value:

```
14189568 = 256 (BW + Delay + X)
```

To calculate the value for the BW.
We know that BW is 10,000,000 Kbps divided by the slowest BW along the path to the destination, in this case we see that in the output of the Vector metric "minimum bandwidth is 1544 Kbit" for 10.1.14.4. Therefore, it will be:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

10,000,000 / 1544 = 6476.68393 We should drop the fractions, so the result is 6476.

To calculate the delay value.
The delay value = The sum of the delay values of the links along the path to the destination divided by 10.

The delay value of S0/0.43 interface of R4 is 20,000 PLUS The delay value of Virtual-template 14 of R1 is 100,000 DIVIDED by 10:

100,000 + 20,000 = 120,000 / 10 = 12,000

Therefore:

14189568 = 256 (6476 + 12,000 + X) which means:

14189568 = 256 (18,476 + X)

This can be resolved by dividing "14,189,568" by 256:

14,189,568 / 256 = 55,428 which means:
55,428 = 18,476 + X If this formula is rearranged, we should see the following:

55,428 - 18,476 = X OR

36,952 = X

Therefore, IF the delay value from R1's perspective is RAISED by 36,952, the resulting composite metric will give us the ratio of 1:3, as follows:

In this step we need to add 36,952 to the delay value of the virtual-template 14 on R1, which means:

10,000 + 36,952 = 46,952

On R1

```
R1(config)#int virtual-template 14
R1(config-if)#Delay 46952
```

But remember that the router will NOT have the second best route in its routing table unless the "Variance" command in router configuration mode is configured:

On R1

```
R1(config)#router eigrp 300
R1(config-router)#Variance 3
```

To enforce the changes, let's clear the neighbor adjacency and the routing table:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1

R1#**Clear ip eigrp neighbors**

R1#**Clear ip route ***

To verify the configuration:

On R1

```
R1#Show ip eigrp topology 10.1.34.0 255.255.255.0

IP-EIGRP (AS 300): Topology entry for 10.1.34.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 4729856
  Routing Descriptor Blocks:
  10.1.13.3 (Virtual-Access2), from 10.1.13.3, Send flag is 0x0
      Composite metric is (4729856/2169856), Route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 120000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
  10.1.14.4 (Virtual-Access3), from 10.1.14.4, Send flag is 0x0
      Composite metric is (14189568/2169856), Route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 489520 microseconds
        Reliability is 255/255
        Load is 5/255
        Minimum MTU is 1500
        Hop count is 1
```

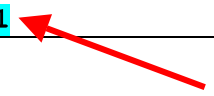
(The rest of the output is omitted)

On R1

```
R1#Show ip route 10.1.34.0

Routing entry for 10.1.34.0/24

  Known via "eigrp 300", distance 90, metric 4729856, type internal
  Redistributing via eigrp 300, eigrp 600
  Advertised by eigrp 600
  Last update from 10.1.14.4 on Virtual-Access3, 00:04:42 ago
  Routing Descriptor Blocks:
    10.1.14.4, from 10.1.14.4, 00:04:42 ago, via Virtual-Access3
      Route metric is 14189568, traffic share count is 1
```



©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

```
      Total delay is 489520 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 5/255, Hops 1
* 10.1.13.3, from 10.1.13.3, 00:04:42 ago, via Virtual-Access2
      Route metric is 4729856, traffic share count is 3
      Total delay is 120000 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
```



Task 20

Configure the routers in AS 500 such that a dead neighbor is detected within 3 seconds.

On R1 and R3

```
R3(config)#int s0/1
R3(config-if)#ip hello-interval eigrp 500 1
R3(config-if)#ip hold-time eigrp 500 3
```

Task 21

Routers in AS 400 should be configured to ONLY use Bandwidth when calculating their composite metric.

On R2

```
R2(config)#router eigrp 400
R2(config-router)#Metric weights 0 1 0 0 0 0
```

You should see the following message:

IP-EIGRP(0) 400: Neighbor 10.1.24.4 (FastEthernet0/1) is down: metric changed
IP-EIGRP(0) 400: Neighbor 10.1.24.4 (FastEthernet0/1) is down: K-value mismatch

On R4

```
R4(config)#router eigrp 400
R4(config-router)#Metric weights 0 1 0 0 0 0
```

You should see the following message:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP-EIGRP(0) 400: Neighbor 10.1.24.2 (FastEthernet0/1) is up: new adjacency

To verify the configuration:

On R4

```
R4#Show ip protocol | Inc EIGRP metric weight|Routing Protocol

Routing Protocol is "eigrp 300"
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0

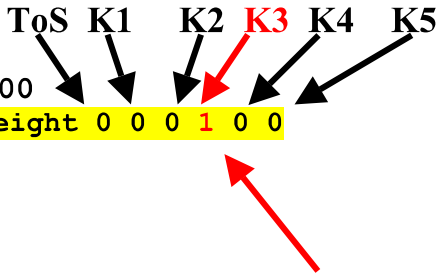
Routing Protocol is "eigrp 200"
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
Routing Protocol is "eigrp 400"
  EIGRP metric weight K1=1, K2=0, K3=0, K4=0, K5=0
```

Note the K value for AS 400 is changed.

Task 22

Routers in AS 500 should be configured to use Delay when calculating their composite metric.

On R1 and R3




```
R1(config)#Router eigrp 500
R1(config-router)#Metric weight 0 0 0 1 0 0
```

To verify the configuration:

On R1

```
R1#Show ip protocol | Inc EIGRP metric weight|Routing Protocol

Routing Protocol is "eigrp 300"
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0

Routing Protocol is "eigrp 600"
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
Routing Protocol is "eigrp 500"
  EIGRP metric weight K1=0, K2=0, K3=1, K4=0, K5=0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 23

BB3 should be configured to filter existing and future Eigrp routes that have a composite metric ranging 4,700,000 to 4,900,000.

To accomplish this task you must deny the routes that have a composite metric of 4,700,000 to 4,900,000 using a route-map, to deny the routes in the specified range, you must use the following formula to calculate the metric and the deviation value:

The two numbers are added and then, divided by 2:

4,900,000 + 4,700,000 = 9,600,000
9,600,000 / 2 = 4,800,000 → This gives us the metric value

To calculate the deviation value:
Subtract the start of the range number (The lower value or the "*from*" value, in this case 4,700,000) from the "*to*" value, in this case 4,900,000 and then, divide the result by two:

4,900,000 - 4,700,000 = 200,000
200,000 / 2 = 100,000

Therefore, the "match metric" command will have the following values:
Match metric 4,800,000 +- 100,000

To display the routing table of BB3:

On BB3

```
BB3#Show ip route eigrp

     102.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D EX    102.1.35.0/24 [170/4868096] via 10.1.56.5, 00:55:02, FastEthernet0/0
D EX    102.1.34.0/24 [170/4868096] via 10.1.56.5, 00:55:02, FastEthernet0/0
D EX    102.1.34.0/23 [170/5372416] via 10.1.56.5, 00:55:02, FastEthernet0/0
D EX 1.0.0.0/8 [170/4860416] via 10.1.56.5, 00:55:02, FastEthernet0/0
D EX 2.0.0.0/8 [170/5372416] via 10.1.56.5, 00:55:02, FastEthernet0/0
     100.0.0.0/24 is subnetted, 8 subnets
D EX    100.1.5.0 [170/4860416] via 10.1.56.5, 00:55:02, FastEthernet0/0

D EX    100.1.4.0 [170/4860416] via 10.1.56.5, 00:55:02, FastEthernet0/0
D EX    100.1.7.0 [170/4860416] via 10.1.56.5, 00:55:02, FastEthernet0/0
D EX    100.1.6.0 [170/4860416] via 10.1.56.5, 00:55:02, FastEthernet0/0
D EX    100.1.1.0 [170/4860416] via 10.1.56.5, 00:55:02, FastEthernet0/0
D EX    100.1.0.0 [170/4860416] via 10.1.56.5, 00:55:02, FastEthernet0/0
D EX    100.1.3.0 [170/4860416] via 10.1.56.5, 00:55:02, FastEthernet0/0
D EX    100.1.2.0 [170/4860416] via 10.1.56.5, 00:55:03, FastEthernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
D EX 3.0.0.0/8 [170/2812416] via 10.1.56.5, 00:55:03, FastEthernet0/0
D EX 4.0.0.0/8 [170/2300416] via 10.1.56.5, 00:55:03, FastEthernet0/0
D    5.0.0.0/8 [90/156160] via 10.1.56.5, 00:55:31, FastEthernet0/0
D    6.0.0.0/8 [90/158720] via 10.1.56.5, 00:55:31, FastEthernet0/0
     112.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D EX    112.2.2.22/32 [170/2305536] via 10.1.56.5, 00:42:07, FastEthernet0/0
D EX    112.112.112.0/24
           [170/2180096] via 10.1.56.5, 00:42:07, FastEthernet0/0
D EX    112.1.1.11/32 [170/4862976] via 10.1.56.5, 00:55:04, FastEthernet0/0
D EX 130.3.0.0/16 [170/2822656] via 10.1.56.5, 00:31:53, FastEthernet0/0
     10.0.0.0/8 is variably subnetted, 15 subnets, 2 masks
D EX    10.1.14.1/32 [170/4732416] via 10.1.56.5, 00:55:04, FastEthernet0/0
D EX    10.1.14.0/24 [170/4732416] via 10.1.56.5, 00:55:04, FastEthernet0/0
D EX    10.1.13.3/32 [170/7292416] via 10.1.56.5, 00:55:04, FastEthernet0/0
D EX    10.1.13.0/24 [170/5244416] via 10.1.56.5, 00:55:04, FastEthernet0/0
D EX    10.1.13.1/32 [170/5244416] via 10.1.56.5, 00:55:04, FastEthernet0/0
D EX    10.1.12.0/24 [170/5244416] via 10.1.56.5, 00:55:04, FastEthernet0/0
D EX    10.11.11.0/24 [170/2182656] via 10.1.56.5, 00:42:08, FastEthernet0/0
D EX    10.11.12.0/24 [170/2177536] via 10.1.56.5, 00:42:08, FastEthernet0/0
D EX    10.1.4.0/24 [170/2174976] via 10.1.56.5, 00:55:04, FastEthernet0/0
D EX    10.1.24.0/24 [170/2174976] via 10.1.56.5, 00:55:04, FastEthernet0/0
D EX    10.1.46.0/24 [170/2684416] via 10.1.56.5, 00:55:05, FastEthernet0/0
D EX    10.1.45.0/24 [170/2172416] via 10.1.56.5, 00:55:22, FastEthernet0/0
D EX    10.1.34.0/24 [170/2684416] via 10.1.56.5, 00:55:06, FastEthernet0/0
D EX    10.1.100.0/24 [170/2694656] via 10.1.56.5, 00:42:09, FastEthernet0/0
```

The routes that are highlighted are the ones that BB3 should deny:

On BB3

```
BB3(config)#Route-map TST deny 10
BB3(config-route-map)#match metric 4800000 +- 100000
BB3(config)#Route-map TST permit 90

BB3(config)#Router eigrp 100
BB3(config-router)#distribute-list route-map TST in F0/0
```

To verify the configuration:

On BB3

```
BB3#Show ip route eigrp

     102.0.0.0/23 is subnetted, 1 subnets
D EX    102.1.34.0 [170/5372416] via 10.1.56.5, 00:00:16, FastEthernet0/0
D EX 2.0.0.0/8 [170/5372416] via 10.1.56.5, 00:00:16, FastEthernet0/0
D EX 3.0.0.0/8 [170/2812416] via 10.1.56.5, 00:00:16, FastEthernet0/0
D EX 4.0.0.0/8 [170/2300416] via 10.1.56.5, 00:00:16, FastEthernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
D    5.0.0.0/8 [90/156160] via 10.1.56.5, 00:00:16, FastEthernet0/0
D    6.0.0.0/8 [90/158720] via 10.1.56.5, 00:00:16, FastEthernet0/0
     112.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
D EX    112.2.2.22/32 [170/2305536] via 10.1.56.5, 00:00:16, FastEthernet0/0
D EX    112.112.112.0/24
           [170/2180096] via 10.1.56.5, 00:00:16, FastEthernet0/0
D EX 130.3.0.0/16 [170/2822656] via 10.1.56.5, 00:00:16, FastEthernet0/0
     10.0.0.0/8 is variably subnetted, 13 subnets, 2 masks
D EX    10.1.13.3/32 [170/7292416] via 10.1.56.5, 00:00:16, FastEthernet0/0
D EX    10.1.13.0/24 [170/5244416] via 10.1.56.5, 00:00:16, FastEthernet0/0
D EX    10.1.13.1/32 [170/5244416] via 10.1.56.5, 00:00:17, FastEthernet0/0
D EX    10.1.12.0/24 [170/5244416] via 10.1.56.5, 00:00:17, FastEthernet0/0
D EX    10.11.11.0/24 [170/2182656] via 10.1.56.5, 00:00:17, FastEthernet0/0
D EX    10.11.12.0/24 [170/2177536] via 10.1.56.5, 00:00:17, FastEthernet0/0
D EX    10.1.4.0/24 [170/2174976] via 10.1.56.5, 00:00:17, FastEthernet0/0
D EX    10.1.24.0/24 [170/2174976] via 10.1.56.5, 00:00:17, FastEthernet0/0
D EX    10.1.46.0/24 [170/2684416] via 10.1.56.5, 00:00:17, FastEthernet0/0
D EX    10.1.45.0/24 [170/2172416] via 10.1.56.5, 00:00:17, FastEthernet0/0
D EX    10.1.34.0/24 [170/2684416] via 10.1.56.5, 00:00:19, FastEthernet0/0
D EX    10.1.100.0/24 [170/2694656] via 10.1.56.5, 00:00:19, FastEthernet0/0
```

Task 24

Configure BB2 such that it NEVER has external Eigrp routes in its routing table. DO NOT use any access-list, route-map or prefix-list to accomplish this task.

To see the existing routes in BB2's routing table:

On BB2

```
BB2#Show ip route eigrp | Inc D

D EX    102.1.35.0/24 [170/156160] via 10.11.12.2, 00:12:36, FastEthernet0/0
D EX    102.1.34.0/24 [170/156160] via 10.11.12.2, 00:12:36, FastEthernet0/0
D EX    102.1.34.0/23 [170/156160] via 10.11.12.2, 00:12:36, FastEthernet0/0
D EX 1.0.0.0/8 [170/158720] via 112.112.112.11, 00:12:36, FastEthernet0/1
D EX 2.0.0.0/8 [170/156160] via 10.11.12.2, 00:12:36, FastEthernet0/0
D EX    100.1.5.0 [170/158720] via 112.112.112.11, 02:07:20, FastEthernet0/1

D EX    100.1.4.0 [170/158720] via 112.112.112.11, 02:07:20, FastEthernet0/1
D EX    100.1.7.0 [170/158720] via 112.112.112.11, 02:07:20, FastEthernet0/1
D EX    100.1.6.0 [170/158720] via 112.112.112.11, 02:07:20, FastEthernet0/1
D EX    100.1.1.0 [170/158720] via 112.112.112.11, 02:07:20, FastEthernet0/1
D EX    100.1.0.0 [170/158720] via 112.112.112.11, 02:07:20, FastEthernet0/1
D EX    100.1.3.0 [170/158720] via 112.112.112.11, 02:07:20, FastEthernet0/1
D EX    100.1.2.0 [170/158720] via 112.112.112.11, 02:07:20, FastEthernet0/1
D EX 33.0.0.0/8 [170/5886976] via 10.11.12.2, 00:12:16, FastEthernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
D EX 3.0.0.0/8 [170/2718720] via 112.112.112.11, 00:12:37, FastEthernet0/1
D EX 4.0.0.0/8 [170/5372416] via 10.11.12.2, 00:12:37, FastEthernet0/0
D EX 5.0.0.0/8 [170/5884416] via 10.11.12.2, 00:12:37, FastEthernet0/0
D EX 6.0.0.0/8 [170/5886976] via 10.11.12.2, 00:12:37, FastEthernet0/0
D EX    112.1.1.11/32 [170/2302976] via 10.11.12.2, 00:12:37, FastEthernet0/0
D EX 130.3.0.0/16 [170/2302976] via 112.112.112.11, 01:22:06, FastEthernet0/1
D EX    10.1.14.4/32
D EX    10.1.14.0/24
D EX    10.1.13.3/32
D EX    10.1.13.0/24
D EX    10.1.12.0/24 [170/2172416] via 10.11.12.2, 00:12:39, FastEthernet0/0
D       10.11.11.0/24 [90/30720] via 112.112.112.11, 02:07:24, FastEthernet0/1
D EX    10.1.4.0/24 [170/33280] via 10.11.12.2, 00:12:39, FastEthernet0/0
D EX    10.1.24.0/24 [170/30720] via 10.11.12.2, 00:12:39, FastEthernet0/0
D EX    10.1.46.0/24 [170/5756416] via 10.11.12.2, 00:12:39, FastEthernet0/0
D EX    10.1.45.0/24 [170/5756416] via 10.11.12.2, 00:12:39, FastEthernet0/0
D EX    10.1.34.0/24 [170/5244416] via 10.11.12.2, 00:12:39, FastEthernet0/0
D EX    10.1.56.0/24 [170/5758976] via 10.11.12.2, 00:12:39, FastEthernet0/0
D EX    10.1.100.0/24
```

To configure the task:

On BB2

```
BB2(config)#Router eigrp 600
BB2(config-router)#Distance Eigrp 90 255
```

The above command assigns an administrative distance of 90 to all EIGRP internal routes and the second value assigns an administrative distance of 255 to the existing and future external routes, which means that the external routes will NOT be injected into the routing table.

To verify the configuration:

On BB2

```
BB2#Show ip route eigrp

     10.0.0.0/24 is subnetted, 2 subnets
D       10.11.11.0 [90/30720] via 112.112.112.11, 00:07:17, FastEthernet0/1
```

Task 25

Configure R2 such that Eigrp NEVER uses more than 25% of the frame-relay link's bandwidth.

On R2

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config)#Int s0/0.21
R2(config-subif)#IP bandwidth-percent eigrp 300 25
```

Task 26

Configure Eigrp AS 700 on SW1, SW2, SW3 and SW4 and advertise the following interfaces in this AS:

On SW1:	On SW2:	On SW3:	On SW4:
F0/19, F0/21 and F0/23	F0/19 and F0/23	F0/19, F0/21 and F0/23	F0/19 and F0/23

On SW1

```
SW1(config)#ip routing

SW1(config)#router eigrp 700
SW1(config-router)#NO au
SW1(config-router)#netw 120.1.1.1 0.0.0.0
SW1(config-router)#netw 130.1.1.1 0.0.0.0
SW1(config-router)#netw 140.1.1.1 0.0.0.0
```

On SW2

```
SW2(config)#IP routing

SW2(config)#Router eigrp 700
SW2(config-router)#NO au
SW2(config-router)#Netw 123.1.1.2 0.0.0.0
SW2(config-router)#Netw 120.1.1.2 0.0.0.0
```

On SW3

```
SW3(config)#IP routing

SW3(config)#Router eigrp 700
SW3(config-router)#NO au
SW3(config-router)#Netw 123.1.1.3 0.0.0.0
SW3(config-router)#Netw 134.1.1.3 0.0.0.0
SW3(config-router)#Netw 130.1.1.3 0.0.0.0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On SW4

```
SW4(config)#IP routing

SW4(config)#Router eigrp 700
SW4(config-router)#NO au
SW4(config-router)#Netw 134.1.1.4 0.0.0.0
SW4(config-router)#Netw 140.1.1.4 0.0.0.0
```

To verify the configuration:

On SW2

```
SW2#Show ip route eigrp

     140.1.0.0/24 is subnetted, 1 subnets
D       140.1.1.0 [90/30720] via 120.1.1.1, 00:04:07, FastEthernet0/19

     130.1.0.0/24 is subnetted, 1 subnets
D       130.1.1.0 [90/30720] via 123.1.1.3, 00:04:12, FastEthernet0/23
                  [90/30720] via 120.1.1.1, 00:04:12, FastEthernet0/19
     134.1.0.0/24 is subnetted, 1 subnets
   134.1.1.0 [90/30720] via 123.1.1.3, 00:04:08, FastEthernet0/23
```

On SW3

```
SW3#Show ip route eigrp

     140.1.0.0/24 is subnetted, 1 subnets
D       140.1.1.0 [90/30720] via 134.1.1.4, 00:06:52, FastEthernet0/19
                  [90/30720] via 130.1.1.1, 00:06:52, FastEthernet0/21
     120.0.0.0/24 is subnetted, 1 subnets
D       120.1.1.0 [90/30720] via 130.1.1.1, 00:06:47, FastEthernet0/21
                  [90/30720] via 123.1.1.2, 00:06:47, FastEthernet0/23
```

On SW4

```
SW4#Show ip route eigrp

     130.1.0.0/24 is subnetted, 1 subnets
D       130.1.1.0 [90/30720] via 140.1.1.1, 00:08:10, FastEthernet0/23

                  [90/30720] via 134.1.1.3, 00:08:10, FastEthernet0/19
     123.0.0.0/24 is subnetted, 1 subnets
D       123.1.1.0 [90/30720] via 134.1.1.3, 00:08:10, FastEthernet0/19
     120.0.0.0/24 is subnetted, 1 subnets
D       120.1.1.0 [90/30720] via 140.1.1.1, 00:08:10, FastEthernet0/23
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On SW1

```
SW1#Show ip route eigrp

     123.0.0.0/24 is subnetted, 1 subnets
D       123.1.1.0 [90/30720] via 130.1.1.3, 00:08:48, FastEthernet0/21
                  [90/30720] via 120.1.1.2, 00:08:48, FastEthernet0/19
     134.1.0.0/24 is subnetted, 1 subnets
D       134.1.1.0 [90/30720] via 140.1.1.4, 00:08:48, FastEthernet0/23
                  [90/30720] via 130.1.1.3, 00:08:48, FastEthernet0/21
```

Task 27

Configure SW1 to advertise its Lo12, Lo13 and Lo14 as candidate default routes.

This task calls for the "IP default-network" global config command; this command can ONLY be used when "ip routing" is enabled. When this command is configured, the router considers routes to the configured network for installation as the gateway of last resort; the receiving router will flag the route/s as a candidate default route ONLY if it has reachability to the specified network/s.

Step One:
In order to provide reachability to the rest of the switches, the three networks MUST be advertised in Eigrp:

On SW1

```
SW1(config)#Router eigrp 700
SW1(config-router)#Netw 12.1.1.1 0.0.0.0
SW1(config-router)#Netw 13.1.1.1 0.0.0.0
SW1(config-router)#Netw 14.1.1.1 0.0.0.0
```

Step Two:
Configure "IP Default-Network" for the three networks:

On SW1

```
SW1(config)#IP default-network 12.0.0.0
SW1(config)#IP default-network 13.0.0.0
SW1(config)#IP default-network 14.0.0.0
```

To verify the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On SW2

```
SW2#Sh ip route eigrp

     140.1.0.0/24 is subnetted, 1 subnets
D       140.1.1.0 [90/30720] via 120.1.1.1, 03:45:56, FastEthernet0/19
     130.1.0.0/24 is subnetted, 1 subnets
D       130.1.1.0 [90/30720] via 123.1.1.3, 03:46:01, FastEthernet0/23
                  [90/30720] via 120.1.1.1, 03:46:01, FastEthernet0/19
D*   12.0.0.0/8 [90/156160] via 120.1.1.1, 00:01:41, FastEthernet0/19
D*   13.0.0.0/8 [90/156160] via 120.1.1.1, 00:01:37, FastEthernet0/19
     134.1.0.0/24 is subnetted, 1 subnets
D       134.1.1.0 [90/30720] via 123.1.1.3, 03:45:57, FastEthernet0/23
D*   14.0.0.0/8 [90/156160] via 120.1.1.1, 00:01:34, FastEthernet0/19
```

On SW3

```
SW3#Sh ip route eigrp


     140.1.0.0/24 is subnetted, 1 subnets
D       140.1.1.0 [90/30720] via 134.1.1.4, 03:46:26, FastEthernet0/19
                  [90/30720] via 130.1.1.1, 03:46:26, FastEthernet0/21
D*   12.0.0.0/8 [90/156160] via 130.1.1.1, 00:03:20, FastEthernet0/21
D*   13.0.0.0/8 [90/156160] via 130.1.1.1, 00:03:17, FastEthernet0/21
D*   14.0.0.0/8 [90/156160] via 130.1.1.1, 00:03:12, FastEthernet0/21
     120.0.0.0/24 is subnetted, 1 subnets
D       120.1.1.0 [90/30720] via 130.1.1.1, 03:46:21, FastEthernet0/21
                  [90/30720] via 123.1.1.2, 03:46:21, FastEthernet0/23
```

On SW4

```
SW4#Show ip route eigrp

     130.1.0.0/24 is subnetted, 1 subnets
D       130.1.1.0 [90/30720] via 140.1.1.1, 03:47:02, FastEthernet0/23
                  [90/30720] via 134.1.1.3, 03:47:02, FastEthernet0/19
     123.0.0.0/24 is subnetted, 1 subnets
D       123.1.1.0 [90/30720] via 134.1.1.3, 03:47:02, FastEthernet0/19
D*   12.0.0.0/8 [90/156160] via 140.1.1.1, 00:04:00, FastEthernet0/23
D*   13.0.0.0/8 [90/156160] via 140.1.1.1, 00:03:57, FastEthernet0/23
D*   14.0.0.0/8 [90/156160] via 140.1.1.1, 00:03:53, FastEthernet0/23
     120.0.0.0/24 is subnetted, 1 subnets

D       120.1.1.0 [90/30720] via 140.1.1.1, 03:47:02, FastEthernet0/23
```

Task 28

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Configure a frame-relay Point-to-point connection between R2 and R3. R2 should be configured with 10.1.23.2 /24 and R3 should be configured with 10.1.23.3/24. These routers should advertise this IP address in Eigrp AS 300.

On R2

```
R2(config)#int S0/0.23 p
R2(config-subif)#ip addr 10.1.23.2 255.255.255.0
R2(config-subif)#frame inter 203

R2(config)#router eigrp 300
R2(config-router)#Netw 10.1.23.2 0.0.0.0
```

On R3

```
R3(config)#int S0/0.32 p
R3(config-subif)#ip addr 10.1.23.3 255.255.255.0
R3(config-subif)#frame inter 302

R3(config-subif)#router eigrp 300
R3(config-router)#Netw 10.1.23.3 0.0.0.0
```

To verify the configuration:

On R2

```
R2#Show ip eigrp neighbors

IP-EIGRP neighbors for process 300
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   10.1.23.3               Se0/0.23          12 00:03:32  176  1056  0  67
0   10.1.12.1               Se0/0.21          12 00:26:04   99   594  0  88

IP-EIGRP neighbors for process 400
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   10.1.24.4               Fa0/1             12 00:26:24  374  2244  0  7


IP-EIGRP neighbors for process 600
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   10.11.12.22             Fa0/0             13 00:24:49    1   200  0  16
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 29

Configure SW2, SW3 and SW4 based on the following policy:

1. SW2 should have 12.0.0.0 /8, 13.0.0.0 /8 and 14.0.0.0 /8 in its routing table, BUT it should ONLY use 12.0.0.0 /8 as the candidate default route. SW2 should be configured to accomplish this task.

2. SW3 should have 12.0.0.0 /8, 13.0.0.0 /8 and 14.0.0.0 /8 in its routing table, BUT it should ONLY use 13.0.0.0 /8 as the candidate default route. SW3 should be configured to accomplish this task.

3. SW4 should have 12.0.0.0 /8, 13.0.0.0 /8 and 14.0.0.0 /8 in its routing table, BUT it should use 14.0.0.0 /8 as the candidate default route. SW4 should be configured to accomplish this task.

To configure the first policy:

On SW2

Note SW2 has three candidate default routes in its routing table.

```
SW2#Show ip route eigrp

     140.1.0.0/24 is subnetted, 1 subnets
D       140.1.1.0 [90/30720] via 120.1.1.1, 00:00:04, FastEthernet0/19
     130.1.0.0/24 is subnetted, 1 subnets
D       130.1.1.0 [90/30720] via 123.1.1.3, 00:00:04, FastEthernet0/23
                  [90/30720] via 120.1.1.1, 00:00:04, FastEthernet0/19
D*    12.0.0.0/8 [90/156160] via 120.1.1.1, 00:00:04, FastEthernet0/19
D*    13.0.0.0/8 [90/156160] via 120.1.1.1, 00:00:04, FastEthernet0/19
     134.1.0.0/24 is subnetted, 1 subnets
D       134.1.1.0 [90/30720] via 123.1.1.3, 00:00:04, FastEthernet0/23
D*    14.0.0.0/8 [90/156160] via 120.1.1.1, 00:00:05, FastEthernet0/19
```

In order to accomplish this task, you must configure an access-list to identify the network that must be used as a candidate default route:

```
SW2(config)#Access-list 12 permit 12.0.0.0 0.255.255.255
```

Once the access-list is configured, you must configure the "Default-information allowed in" command to reference the access-list. Basically this command will ONLY allow the specified Network to be used as a candidate default route.

```
SW2(config)#Router eigrp 700
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
SW2(config-router)#default-information allowed in 12

SW2#Clear ip eigrp neighbors

SW2#Show ip route eigrp

     140.1.0.0/24 is subnetted, 1 subnets
D       140.1.1.0 [90/30720] via 120.1.1.1, 00:00:14, FastEthernet0/19
     130.1.0.0/24 is subnetted, 1 subnets
D       130.1.1.0 [90/30720] via 123.1.1.3, 00:00:14, FastEthernet0/23
                  [90/30720] via 120.1.1.1, 00:00:14, FastEthernet0/19
D*   12.0.0.0/8 [90/156160] via 120.1.1.1, 00:00:14, FastEthernet0/19
D    13.0.0.0/8 [90/156160] via 120.1.1.1, 00:00:14, FastEthernet0/19
     134.1.0.0/24 is subnetted, 1 subnets
D       134.1.1.0 [90/30720] via 123.1.1.3, 00:00:14, FastEthernet0/23
D    14.0.0.0/8 [90/156160] via 120.1.1.1, 00:00:14, FastEthernet0/19
```

To configure the second policy:

On SW3

```
SW3#Show ip route eigrp

     140.1.0.0/24 is subnetted, 1 subnets
D       140.1.1.0 [90/30720] via 134.1.1.4, 00:14:15, FastEthernet0/19
                  [90/30720] via 130.1.1.1, 00:14:15, FastEthernet0/21
D*   12.0.0.0/8 [90/156160] via 130.1.1.1, 00:14:15, FastEthernet0/21
D*   13.0.0.0/8 [90/156160] via 130.1.1.1, 00:14:15, FastEthernet0/21
D*   14.0.0.0/8 [90/156160] via 130.1.1.1, 00:14:15, FastEthernet0/21
     120.0.0.0/24 is subnetted, 1 subnets
D       120.1.1.0 [90/30720] via 130.1.1.1, 00:14:15, FastEthernet0/21
                  [90/30720] via 123.1.1.2, 00:14:16, FastEthernet0/23

SW3(config)#access-list 13 permit 13.0.0.0 0.255.255.255

SW3(config)#router eigrp 700
SW3(config-router)#Default-information allowed in 13

SW3#Clear ip eigrp neighbors

SW3#Show ip route eigrp

     140.1.0.0/24 is subnetted, 1 subnets
D       140.1.1.0 [90/30720] via 134.1.1.4, 00:00:13, FastEthernet0/19
                  [90/30720] via 130.1.1.1, 00:00:13, FastEthernet0/21
D    12.0.0.0/8 [90/156160] via 130.1.1.1, 00:00:13, FastEthernet0/21
D*   13.0.0.0/8 [90/156160] via 130.1.1.1, 00:00:13, FastEthernet0/21
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
D    14.0.0.0/8 [90/156160] via 130.1.1.1, 00:00:13, FastEthernet0/21
     120.0.0.0/24 is subnetted, 1 subnets
D       120.1.1.0 [90/30720] via 130.1.1.1, 00:00:13, FastEthernet0/21
                  [90/30720] via 123.1.1.2, 00:00:13, FastEthernet0/23
```

To configure the third policy:

On SW4

```
SW4#Sh ip route eigrp

     130.1.0.0/24 is subnetted, 1 subnets
D       130.1.1.0 [90/30720] via 140.1.1.1, 00:07:45, FastEthernet0/23
                  [90/30720] via 134.1.1.3, 00:07:45, FastEthernet0/19
     123.0.0.0/24 is subnetted, 1 subnets
D       123.1.1.0 [90/30720] via 134.1.1.3, 00:07:45, FastEthernet0/19
D*   12.0.0.0/8 [90/156160] via 140.1.1.1, 00:07:45, FastEthernet0/23
D*   13.0.0.0/8 [90/156160] via 140.1.1.1, 00:07:45, FastEthernet0/23
D*   14.0.0.0/8 [90/156160] via 140.1.1.1, 00:07:45, FastEthernet0/23
     120.0.0.0/24 is subnetted, 1 subnets

D       120.1.1.0 [90/30720] via 140.1.1.1, 00:07:46, FastEthernet0/23

SW4(config)#access-list 14 permit 14.0.0.0 0.255.255.255

SW4(config)#router eigrp 700
SW4(config-router)#default-information in 14

SW4#Clear ip eigrp neighbors
```

To verify the configuration:

On SW4

```
SW4#Show ip route eigrp

     130.1.0.0/24 is subnetted, 1 subnets
D       130.1.1.0 [90/30720] via 140.1.1.1, 00:01:07, FastEthernet0/23
                  [90/30720] via 134.1.1.3, 00:01:07, FastEthernet0/19

     123.0.0.0/24 is subnetted, 1 subnets
D       123.1.1.0 [90/30720] via 134.1.1.3, 00:01:07, FastEthernet0/19
D    12.0.0.0/8 [90/156160] via 140.1.1.1, 00:01:07, FastEthernet0/23
D    13.0.0.0/8 [90/156160] via 140.1.1.1, 00:01:07, FastEthernet0/23
D*   14.0.0.0/8 [90/156160] via 140.1.1.1, 00:01:07, FastEthernet0/23
     120.0.0.0/24 is subnetted, 1 subnets
D       120.1.1.0 [90/30720] via 140.1.1.1, 00:01:08, FastEthernet0/23
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 30

Configure the link between R3 and SW1 in AS 800 and redistribute the routes from AS 700 into this AS. R3 should NOT have any candidate default route in its routing table.

On SW1

```
SW1(config)#router eigrp 800
SW1(config-router)#NO au
SW1(config-router)#netw 10.1.113.113 0.0.0.0
SW1(config-router)#redistribute eigrp 700
```

On R3

```
R3(config)#router eigrp 800
R3(config-router)#NO au
R3(config-router)#netw 10.1.113.3 0.0.0.0
```

To verify the configuration:

On R3

```
R3#Show ip route eigrp 800

     140.1.0.0/24 is subnetted, 1 subnets
D EX    140.1.1.0 [170/30720] via 10.1.113.113, 00:01:43, FastEthernet0/0
     130.1.0.0/24 is subnetted, 1 subnets
D EX    130.1.1.0 [170/30720] via 10.1.113.113, 00:01:43, FastEthernet0/0
     123.0.0.0/24 is subnetted, 1 subnets
D EX    123.1.1.0 [170/33280] via 10.1.113.113, 00:01:43, FastEthernet0/0
D*EX 12.0.0.0/8 [170/156160] via 10.1.113.113, 00:01:43, FastEthernet0/0
D*EX 13.0.0.0/8 [170/156160] via 10.1.113.113, 00:01:43, FastEthernet0/0
     134.1.0.0/24 is subnetted, 1 subnets
D EX    134.1.1.0 [170/33280] via 10.1.113.113, 00:01:43, FastEthernet0/0
D*EX 14.0.0.0/8 [170/156160] via 10.1.113.113, 00:01:43, FastEthernet0/0
     120.0.0.0/24 is subnetted, 1 subnets
D EX    120.1.1.0 [170/30720] via 10.1.113.113, 00:01:43, FastEthernet0/0
```

Note R3 is using the three highlighted routes as candidate defaults, to stop this behavior:

On R3

```
R3(config)#Router eigrp 800
R3(config-router)#NO default-information in
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the configuration:

On R3

```
R3#Show ip route eigrp 800

     140.1.0.0/24 is subnetted, 1 subnets
D EX    140.1.1.0 [170/30720] via 10.1.113.113, 00:00:06, FastEthernet0/0
     130.1.0.0/24 is subnetted, 1 subnets
D EX    130.1.1.0 [170/30720] via 10.1.113.113, 00:00:06, FastEthernet0/0
     123.0.0.0/24 is subnetted, 1 subnets
D EX    123.1.1.0 [170/33280] via 10.1.113.113, 00:00:06, FastEthernet0/0
D EX 12.0.0.0/8 [170/156160] via 10.1.113.113, 00:00:06, FastEthernet0/0
D EX 13.0.0.0/8 [170/156160] via 10.1.113.113, 00:00:06, FastEthernet0/0
     134.1.0.0/24 is subnetted, 1 subnets
D EX    134.1.1.0 [170/33280] via 10.1.113.113, 00:00:06, FastEthernet0/0
D EX 14.0.0.0/8 [170/156160] via 10.1.113.113, 00:00:06, FastEthernet0/0
     120.0.0.0/24 is subnetted, 1 subnets
D EX    120.1.1.0 [170/30720] via 10.1.113.113, 00:00:06, FastEthernet0/0
```

Task 31

Configure injection of default routes on R3 based on the following policy:

1. For R3's neighbor/s in AS 300:
   R3 should inject a default route for every neighbor in this AS. You should use an interface configuration mode to accomplish this task.

2. For R3's neighbor/s in AS 500:
   R3 should be configured to advertise a default route plus networks 12.0.0.0 /8, 13.0.0.0 /8 and 14.0.0.0 /8 from AS 800.
3. For R3's neighbor/s in AS 800:
   R3 should inject a default route into this AS.

4. R1 should always prefer the default route through AS 300 and NOT any other AS.

To configure the first item:

On R3

R3 has three neighbors in AS 300 and they are reached via S0/0.32, S0/0.34 and virtual-template 31 interfaces, therefore, the "IP summary-address" interface configuration mode command must be configured on these three interfaces:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3

```
R3(config)#int S0/0.32
R3(config-subif)#ip summary-address eigrp 300 0.0.0.0 0.0.0.0

R3(config)#int s0/0.34
R3(config-subif)#ip summary-address eigrp 300 0.0.0.0 0.0.0.0

R3(config)#Int virtual-template 31
R3(config-if)#ip summary-address eigrp 300 0.0.0.0 0.0.0.0
```

To verify the configuration:

On R2

```
R2#Show ip route eigrp 300 | Inc 0.0.0.0/0

D*    0.0.0.0/0 [90/2297856] via 10.1.23.3, 00:17:23, Serial0/0.23
```

On R4

```
R4#Show ip route eigrp 300 | Inc 0.0.0.0/0

D*    0.0.0.0/0 [90/2297856] via 10.1.34.3, 00:18:36, Serial0/0.43
```

On R1

```
R1#Show ip route eigrp 300 | Inc 0.0.0.0/0

D*    0.0.0.0/0 [90/2713600] via 10.1.13.3, 00:00:24, Virtual-Access1
```

To configure the second item:

On R3

By default, if an "IP summary-address" is configured, then every other route is suppressed, unless a "leak-map" is configured. But in order to configure a leak-map to leak networks 12.0.0.0/8, 13.0.0.0/8

and 14.0.0.0/8, these three networks must be known by AS 500.

Since the routes are ONLY advertised in AS 700, they should be redistributed into AS 500, as follows:

On R3

```
R3(config)#access-list 1 permit 12.0.0.0 0.255.255.255
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config)#access-list 1 permit 13.0.0.0 0.255.255.255
R3(config)#access-list 1 permit 14.0.0.0 0.255.255.255

R3(config)#route-map intoAS500
R3(config-route-map)#match ip addr 1

R3(config)#router eigrp 500
R3(config-router)#redistribute Eigrp 800 route-map intoAS500
```

To verify the configuration:

On R1

```
R1#Show ip route eigrp 500

D EX 130.3.0.0/16 [170/640000] via 10.1.100.3, 00:02:51, Serial0/1
D EX 12.0.0.0/8 [170/642560] via 10.1.100.3, 00:08:51, Serial0/1
D EX 13.0.0.0/8 [170/642560] via 10.1.100.3, 00:08:51, Serial0/1
D EX 14.0.0.0/8 [170/642560] via 10.1.100.3, 00:08:51, Serial0/1
```

The highlighted routes are the ones coming from AS 800. Network 130.3.0.0 /16 was redistributed into this AS in Task 10. Therefore, another access-list is configured to include this network, as follows:

On R3

```
R3(config)#access-list 2 permit 12.0.0.0 0.255.255.255
R3(config)#access-list 2 permit 13.0.0.0 0.255.255.255
R3(config)#access-list 2 permit 14.0.0.0 0.255.255.255
R3(config)#access-list 2 permit 130.3.0.0 0.0.255.255

R3(config)#Route-map TST permit 10
R3(config-route-map)#match ip addr 2

R3(config-route-map)#int s0/1
R3(config-if)#ip summary-address eigrp 500 0.0.0.0 0.0.0.0 leak-map TST
```

To verify the configuration:

On R1

```
R1#Show ip route eigrp 500

D EX 130.3.0.0/16 [170/640000] via 10.1.100.3, 00:09:38, Serial0/1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
D EX 12.0.0.0/8 [170/642560] via 10.1.100.3, 00:15:39, Serial0/1
D EX 13.0.0.0/8 [170/642560] via 10.1.100.3, 00:15:39, Serial0/1
D EX 14.0.0.0/8 [170/642560] via 10.1.100.3, 00:15:39, Serial0/1
D*   0.0.0.0/0 [90/640000] via 10.1.100.3, 00:00:36, Serial0/1
```

To configure the third item:

On R3

```
R3(config)#int f0/0
R3(config-if)#ip summary-address eigrp 800 0.0.0.0 0.0.0.0
```

To verify the configuration:

On SW1

```
SW1#Show ip route eigrp 800

D*   0.0.0.0/0 [90/30720] via 10.1.113.3, 00:04:43, FastEthernet0/3
```

To configure the fourth item:

Let's check the routing table of R1:

On R1

```
R1#Show ip route eigrp | Inc 0.0.0.0/0

D*   0.0.0.0/0 [90/640000] via 10.1.100.3, 00:13:57, Serial0/1
```

Note R1 is getting the default route through its S0/1 interface which is in AS 500. This task's requirement states that R1 should always prefer the default route through AS 300 and NOT any other AS.

Is R1 receiving the default route through its Frame-relay connection to R3?

If R1 was receiving a default route through its Frame-relay connection, the default route would be in the Topology table, let's verify that information:

On R1

```
R1#Show ip eigrp topology 300 | S 0.0.0.0/0

P 0.0.0.0/0, 0 successors, FD is Inaccessible
        via 10.1.13.3 (2713600/128256), Virtual-Access1
        via 10.1.14.4 (14317568/2297856), Virtual-Access2
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
        via 10.1.12.2 (2809856/2297856), Serial0/0.12
```

The output of the above show command reveals that the default route was advertised to the local router by R3; the default route was injected with a cost of 2,713,600. If the routing table of R1 is checked again, you will see that R1's cost through AS 500 (S0/1) is 640,000, and that's the reason that R1 chooses the default route through AS 500 and NOT AS 300.
Administrative distance can be manipulated to influence R1's decision, as follows:

```
R1(config)#router eigrp 500
R1(config-router)#distance 91 10.1.100.3 0.0.0.0
```

To verify the configuration:

On R1

```
R1#Sh ip route eigrp | Inc 0.0.0.0/0

D*    0.0.0.0/0 [90/2713600] via 10.1.13.3, 00:00:12, Virtual-Access1
```

Task 32

R4 in AS 400 should ONLY receive routes from R2 and it should NOT advertise any route/s to any router in this AS.

On R2

```
R2#Sh ip rou ei | Inc 10.1.24.4

D       10.1.4.0/24 [90/25600] via 10.1.24.4, 00:00:44, FastEthernet0/1
```

On R4

```
R4(config)#router eigrp 400
R4(config-router)#eigrp stub receive-only
```

To verify the configuration:

On R2

```
R2#Sh ip rou ei | Inc 10.1.24.4
R2#
```

On R4

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4#Show ip rout ei | Inc 10.1.24.2

D EX    112.2.2.22/32 [170/25600] via 10.1.24.2, 00:02:34, FastEthernet0/1
D EX    112.112.112.0/24 [170/25600] via 10.1.24.2, 00:02:34, FastEthernet0/1
D EX    10.11.11.0/24 [170/25600] via 10.1.24.2, 00:02:34, FastEthernet0/1
D EX    10.11.12.0/24 [170/25600] via 10.1.24.2, 00:02:34, FastEthernet0/1
```

Task 33

SW2 should be configured to reach network 130.1.1.0 /24 through SW3, SW2 should ONLY use SW1 if SW3 is down.

On SW2

```
SW2#Show ip route eigrp

     140.1.0.0/24 is subnetted, 1 subnets
D       140.1.1.0 [90/30720] via 120.1.1.1, 03:43:07, FastEthernet0/19
     130.1.0.0/24 is subnetted, 1 subnets
D       130.1.1.0 [90/30720] via 123.1.1.3, 03:43:07, FastEthernet0/23
                  [90/30720] via 120.1.1.1, 03:43:07, FastEthernet0/19
D*   12.0.0.0/8 [90/156160] via 120.1.1.1, 01:02:33, FastEthernet0/19
D    13.0.0.0/8 [90/156160] via 120.1.1.1, 03:43:07, FastEthernet0/19
     134.1.0.0/24 is subnetted, 1 subnets
D       134.1.1.0 [90/30720] via 123.1.1.3, 03:43:07, FastEthernet0/23
D    14.0.0.0/8 [90/156160] via 120.1.1.1, 03:43:08, FastEthernet0/19
```

Note SW2 has two paths in its routing table. In order to configure SW2 so it uses SW3 as the primary path to Network 130.1.1.0 /24 and use SW1 as the backup, the administrative distance can be manipulated:

On SW2

```
SW2(config)#access-list 1 permit 130.1.1.0 0.0.0.255

SW2(config)#router eigrp 700
SW2(config-router)#Distance 91 120.1.1.1 0.0.0.0 1
```

To verify the configuration:

On SW2

```
SW2#Show ip route eigrp
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
     140.1.0.0/24 is subnetted, 1 subnets
D       140.1.1.0 [90/33280] via 123.1.1.3, 00:01:15, FastEthernet0/23
     130.1.0.0/24 is subnetted, 1 subnets
D       130.1.1.0 [90/30720] via 123.1.1.3, 00:01:15, FastEthernet0/23
D    12.0.0.0/8 [90/158720] via 123.1.1.3, 00:01:15, FastEthernet0/23
D    13.0.0.0/8 [90/158720] via 123.1.1.3, 00:01:15, FastEthernet0/23
     134.1.0.0/24 is subnetted, 1 subnets
D       134.1.1.0 [90/30720] via 123.1.1.3, 00:01:15, FastEthernet0/23
D    14.0.0.0/8 [90/158720] via 123.1.1.3, 00:01:15, FastEthernet0/23
```

To test the configuration:

On SW3

```
SW3(config)#Int f0/23
SW3(config-if)#Shut

SW2#Show ip route eigrp

     140.1.0.0/24 is subnetted, 1 subnets
D       140.1.1.0 [90/30720] via 120.1.1.1, 00:00:41, FastEthernet0/19
     130.1.0.0/24 is subnetted, 1 subnets
D       130.1.1.0 [91/30720] via 120.1.1.1, 00:00:22, FastEthernet0/19
D*   12.0.0.0/8 [90/156160] via 120.1.1.1, 00:00:41, FastEthernet0/19
D    13.0.0.0/8 [90/156160] via 120.1.1.1, 00:00:41, FastEthernet0/19
     134.1.0.0/24 is subnetted, 1 subnets
D       134.1.1.0 [90/33280] via 120.1.1.1, 00:00:22, FastEthernet0/19
D    14.0.0.0/8 [90/156160] via 120.1.1.1, 00:00:41, FastEthernet0/19
```

On SW3

```
SW3(config)#Int f0/23
SW3(config-if)#No shut
```

Task 34

Configure BB3 to filter network 10.11.11.0 /24. You should use the existing “distribute-list route-map” to accomplish this task.

On BB3

```
BB3#Show ip route eigrp

     102.0.0.0/23 is subnetted, 1 subnets
D EX    102.1.34.0 [170/5372416] via 10.1.56.5, 00:00:05, FastEthernet0/0
D EX 2.0.0.0/8 [170/5372416] via 10.1.56.5, 00:00:05, FastEthernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
D EX 4.0.0.0/8 [170/2300416] via 10.1.56.5, 00:00:05, FastEthernet0/0
D    5.0.0.0/8 [90/156160] via 10.1.56.5, 00:00:05, FastEthernet0/0
D    6.0.0.0/8 [90/158720] via 10.1.56.5, 00:00:05, FastEthernet0/0
     112.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
D EX    112.2.2.22/32 [170/2305536] via 10.1.56.5, 00:00:05, FastEthernet0/0
D EX    112.112.112.0/24
           [170/2180096] via 10.1.56.5, 00:00:05, FastEthernet0/0
     10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
D EX    10.1.13.3/32 [170/7292416] via 10.1.56.5, 00:00:05, FastEthernet0/0
D EX    10.1.13.0/24 [170/7292416] via 10.1.56.5, 00:00:05, FastEthernet0/0
D EX    10.1.12.0/24 [170/5244416] via 10.1.56.5, 00:00:05, FastEthernet0/0
D EX    10.11.11.0/24 [170/2182656] via 10.1.56.5, 00:00:05, FastEthernet0/0
D EX    10.11.12.0/24 [170/2177536] via 10.1.56.5, 00:00:06, FastEthernet0/0
D EX    10.1.4.0/24 [170/2174976] via 10.1.56.5, 00:00:06, FastEthernet0/0
D EX    10.1.24.0/24 [170/2174976] via 10.1.56.5, 00:00:06, FastEthernet0/0
D EX    10.1.23.0/24 [170/5756416] via 10.1.56.5, 00:00:06, FastEthernet0/0
D EX    10.1.46.0/24 [170/2684416] via 10.1.56.5, 00:00:06, FastEthernet0/0
D EX    10.1.45.0/24 [170/2172416] via 10.1.56.5, 00:00:06, FastEthernet0/0
D EX    10.1.34.0/24 [170/2684416] via 10.1.56.5, 00:00:06, FastEthernet0/0
D*EX 0.0.0.0/0 [170/2812416] via 10.1.56.5, 00:00:06, FastEthernet0/0
```

Since the existing route-map MUST be used to accomplish this task, our ONLY choice here is to use the "Continue" clause.

A route-map evaluates match clauses until a successful match occurs. After a match occurs, the route-map STOPS evaluating match clauses.

When a continue clause is configured, the route-map continues to evaluate and execute match clauses in the specified route-map entry after a successful match occurs.

```
BB3(config)#access-list 11 permit 10.11.11.0 0.0.0.255

BB3(config)#route-map TST deny 10
BB3(config-route-map)#continue 20

BB3(config)#route-map TST deny 20
BB3(config-route-map)#match ip addr 11
```

To see the complete route-map:

On BB3

```
BB3#Sh run

 route-map TST deny 10
 match metric 4800000 +- 100000
 continue 20
```


©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

```
!
route-map TST deny 20
 match ip address 11
!
route-map TST permit 90
```

(The output is modified to show the pertaining section ONLY)

To verify the configuration:

On BB3

```
BB3#Show ip route eigrp | Inc 10.11.11.0
BB3#
```

Task 35

Configure R1 to filter Network 2.0.0.0 /8. DO NOT use distribute-list to accomplish this task.

On R1

```
R1#Sh ip route eigrp 300

     102.0.0.0/23 is subnetted, 1 subnets
D       102.1.34.0 [90/2297856] via 10.1.12.2, 01:55:51, Serial0/0.12
D    2.0.0.0/8 [90/2297856] via 10.1.12.2, 01:55:51, Serial0/0.12
D EX 33.0.0.0/8 [170/14320128] via 10.1.14.4, 00:17:51, Virtual-Access2
D    4.0.0.0/8 [90/12173312] via 10.1.14.4, 01:55:41, Virtual-Access2
D EX 5.0.0.0/8 [170/14317568] via 10.1.14.4, 01:55:41, Virtual-Access2
D EX 6.0.0.0/8 [170/14320128] via 10.1.14.4, 01:55:41, Virtual-Access2
     112.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D EX    112.2.2.22/32 [170/2300416] via 10.1.12.2, 01:38:31, Serial0/0.12
     10.0.0.0/8 is variably subnetted, 15 subnets, 2 masks
D EX    10.1.4.0/24 [170/12047872] via 10.1.14.4, 01:38:35, Virtual-Access2
D EX    10.1.24.0/24 [170/2172416] via 10.1.12.2, 01:55:41, Serial0/0.12
D       10.1.23.0/24 [90/2681856] via 10.1.12.2, 01:55:51, Serial0/0.12

D EX    10.1.46.0/24 [170/14189568] via 10.1.14.4, 01:55:41, Virtual-Access2
D EX    10.1.45.0/24 [170/14189568] via 10.1.14.4, 01:55:41, Virtual-Access2
D       10.1.34.0/24 [90/14189568] via 10.1.14.4, 01:55:42, Virtual-Access2
D EX    10.1.56.0/24 [170/14192128] via 10.1.14.4, 01:55:42, Virtual-Access2
D*   0.0.0.0/0 [90/2713600] via 10.1.13.3, 01:49:23, Virtual-Access1
               [90/2809856] via 10.1.12.2, 01:49:23, Serial0/0.12
```

To configure the filter, the administrative distance can be used. A route with an administrative distance

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

of 255 will NEVER get injected into the routing table. Remember the administrative distance can ONLY be used in this manner because the route is internal, the administrative distance can NOT be manipulated in this manner for the external routes.

On R1

```
R1(config)#router eigrp 300
R1(config-router)#distance 255 10.1.12.2 0.0.0.0 3

R1(config)#access-list 3 permit 2.0.0.0 0.255.255.255
```

To verify the configuration:

On R1

```
R1#Show ip route eigrp 300

   102.0.0.0/23 is subnetted, 1 subnets
D       102.1.34.0 [90/2297856] via 10.1.12.2, 00:00:03, Serial0/0.12
D EX 33.0.0.0/8 [170/14320128] via 10.1.14.4, 00:00:04, Virtual-Access2
D    4.0.0.0/8 [90/12173312] via 10.1.14.4, 00:00:04, Virtual-Access2
D EX 5.0.0.0/8 [170/14317568] via 10.1.14.4, 00:00:04, Virtual-Access2
D EX 6.0.0.0/8 [170/14320128] via 10.1.14.4, 00:00:04, Virtual-Access2
     112.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D EX    112.2.2.22/32 [170/2300416] via 10.1.12.2, 00:00:03, Serial0/0.12
     10.0.0.0/8 is variably subnetted, 15 subnets, 2 masks
D EX    10.1.4.0/24 [170/12047872] via 10.1.14.4, 00:00:04, Virtual-Access2
D EX    10.1.24.0/24 [170/2172416] via 10.1.12.2, 00:00:03, Serial0/0.12
D       10.1.23.0/24 [90/2681856] via 10.1.12.2, 00:00:03, Serial0/0.12
D EX    10.1.46.0/24 [170/14189568] via 10.1.14.4, 00:00:04, Virtual-Access2
D EX    10.1.45.0/24 [170/14189568] via 10.1.14.4, 00:00:04, Virtual-Access2
D       10.1.34.0/24 [90/14189568] via 10.1.14.4, 00:00:04, Virtual-Access2
D EX    10.1.56.0/24 [170/14192128] via 10.1.14.4, 00:00:05, Virtual-Access2
D*   0.0.0.0/0 [90/2713600] via 10.1.13.3, 00:00:04, Virtual-Access1
               [90/2809856] via 10.1.12.2, 00:00:04, Serial0/0.12
```

Task 36

BB1 should be configured to filter any route/s that have a hop count of 4 or more. DO NOT use an access-list or a prefix-list to accomplish this task.

On BB1

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
BB1#Sh ip route eigrp

      102.0.0.0/23 is subnetted, 1 subnets
D EX     102.1.34.0 [170/2300416] via 10.11.11.1, 00:07:42, FastEthernet0/0
D EX 1.0.0.0/8 [170/156160] via 10.11.11.1, 00:07:42, FastEthernet0/0
      100.0.0.0/24 is subnetted, 8 subnets
D EX     100.1.5.0 [170/156160] via 10.11.11.1, 00:07:42, FastEthernet0/0
D EX     100.1.4.0 [170/156160] via 10.11.11.1, 00:07:42, FastEthernet0/0
D EX     100.1.7.0 [170/156160] via 10.11.11.1, 00:07:42, FastEthernet0/0
D EX     100.1.6.0 [170/156160] via 10.11.11.1, 00:07:42, FastEthernet0/0
D EX     100.1.1.0 [170/156160] via 10.11.11.1, 00:07:42, FastEthernet0/0
D EX     100.1.0.0 [170/156160] via 10.11.11.1, 00:07:42, FastEthernet0/0
D EX     100.1.3.0 [170/156160] via 10.11.11.1, 00:07:42, FastEthernet0/0
D EX     100.1.2.0 [170/156160] via 10.11.11.1, 00:07:42, FastEthernet0/0
D EX 33.0.0.0/8 [170/14322688] via 10.11.11.1, 00:00:50, FastEthernet0/0
D EX 4.0.0.0/8 [170/12175872] via 10.11.11.1, 00:07:42, FastEthernet0/0
D EX 5.0.0.0/8 [170/14320128] via 10.11.11.1, 00:00:50, FastEthernet0/0
D EX 6.0.0.0/8 [170/14322688] via 10.11.11.1, 00:00:51, FastEthernet0/0
      112.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D EX     112.2.2.22/32 [170/2302976] via 10.11.11.1, 00:07:44, FastEthernet0/0
D EX 130.3.0.0/16 [170/2300416] via 10.11.11.1, 00:07:44, FastEthernet0/0
      10.0.0.0/8 is variably subnetted, 14 subnets, 2 masks
D EX     10.1.14.0/24 [170/12047872] via 10.11.11.1, 00:07:44, FastEthernet0/0
D EX     10.1.13.3/32 [170/2588160] via 10.11.11.1, 00:07:44, FastEthernet0/0
D EX     10.1.13.0/24 [170/2588160] via 10.11.11.1, 00:07:44, FastEthernet0/0
D EX     10.1.12.0/24 [170/2172416] via 10.11.11.1, 00:07:45, FastEthernet0/0
D        10.11.12.0/24 [90/30720] via 112.112.112.22, 00:07:46, FastEthernet0/1
D EX     10.1.4.0/24 [170/12050432] via 10.11.11.1, 00:07:45, FastEthernet0/0
D EX     10.1.24.0/24 [170/2174976] via 10.11.11.1, 00:07:45, FastEthernet0/0
D EX     10.1.23.0/24 [170/2684416] via 10.11.11.1, 00:07:45, FastEthernet0/0
D EX     10.1.46.0/24 [170/14192128] via 10.11.11.1, 00:07:45, FastEthernet0/0
D EX     10.1.45.0/24 [170/14192128] via 10.11.11.1, 00:00:56, FastEthernet0/0
D EX     10.1.34.0/24 [170/14192128] via 10.11.11.1, 00:07:45, FastEthernet0/0
D EX     10.1.56.0/24 [170/14194688] via 10.11.11.1, 00:00:52, FastEthernet0/0
D EX     10.1.100.0/24 [170/2172416] via 10.11.11.1, 00:07:45, FastEthernet0/0
D*EX 0.0.0.0/0 [170/2716160] via 10.11.11.1, 00:07:45, FastEthernet0/0
```

The hop count of the highlighted networks is 4 hops away. The following show command verifies that:

```
BB1#Sh ip eigrp topology 33.0.0.0 255.0.0.0 | Inc Hop

        Hop count is 4

BB1#Sh ip eigrp topology 6.0.0.0 255.0.0.0 | Inc Hop

        Hop count is 4

BB1(config)#router eigrp 600
BB1(config-router)#metric maximum-hops 3
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the configuration:

On BB1

```
BB1#Show ip route eigrp | Inc 33.0.0.0
BB1#

BB1#Show ip route eigrp | Inc 6.0.0.0
BB1#
```

Task 37

Configure the appropriate routers/switches such that if SW1 loses its Successor and has no feasible successor it will NOT send a query to SW2, SW3 or SW4.

On SW2, SW3, and SW4

```
SWx(config)#router eigrp 700
SWx(config-router)#Eigrp stub connected
```

SW1 will NOT send queries to neighbors that are configured as stub. This can be verified by running "debug eigrp packets" or "Debug eigrp packets terse" and then, shutting down a loopback on SW1 and observing the output of the debug. The difference between the "Debug eigrp packet" and "Debug eigrp packet Terse" is that the later one will NOT show the hello packets in the output of the debug.

Task 38

Routers in AS 400 should wait one hour (after sending queries) before declaring a non responding neighbor as stuck in active, whereas, routers in AS 500 should wait indefinitely.

On R2 and R4

```
Rx(config)#router eigrp 400
Rx(config-router)#timers active-time 60
```

This command instructs the routers in AS 400 to wait for 60 minutes for query replies before considering the route as stuck in active.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1 and R3

```
Rx(config)#router eigrp 500
Rx(config-router)#timers active-time disable
```

This command instructs the routers in AS 500 to wait indefinitely for query replies.

Task 40

Configure Loopback 999 on BB3 using the following policy:
IP address: 200.200.200.3 /24
Tag: 999

On BB3

```
BB3(config)#int lo999
BB3(config-if)#ip addr 200.200.200.3 255.255.255.0

BB3(config)#route-map Task40 permit 10
BB3(config-route-map)#match inter lo999
BB3(config-route-map)#set tag 999
BB3(config)#route-map Task40 permit 90

BB3(config-route-map)#router eigrp 100
BB3(config-router)#redistribute connected route-map Task40 metric 1 1 1 1 1
```

To verify the configuration:

On BB3

```
BB3#Show ip eigrp topology 200.200.200.0 255.255.255.0 | Inc Admin

        Administrator tag is 999 (0x000003E7)
```

On R1

```
R1#Show ip route 200.200.200.0 | Inc Route tag

      Route tag 999
```

Task 41

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Configure Eigrp AS 900 on the F0/0 interface of R5, R6 and BB3; these routers should advertise their Loopback 90 interface in this AS.

On R5

```
R5(config)#Router eigrp 900
R5(config-router)#No auto
R5(config-router)#Netw 10.1.56.5 0.0.0.0
R5(config-router)#Netw 199.1.1.1 0.0.0.0
```

On R6

```
R6(config)#router eigrp 900
R6(config-router)#NO auto
R6(config-router)#Netw 199.1.1.1 0.0.0.0
R6(config-router)#Netw 10.1.56.6 0.0.0.0
```

On BB3

```
BB3(config)#Router eigrp 900
BB3(config-router)#NO au
BB3(config-router)#Netw 10.1.56.3 0.0.0.0
BB3(config-router)#Netw 203.3.3.3 0.0.0.0
```

To verify the configuration:

On R6

```
R6#Sh ip eigrp neighbors 900

IP-EIGRP neighbors for process 900
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   10.1.56.3               Fa0/0             12 00:00:20    5   200  0  10
0   10.1.56.5               Fa0/0             12 00:00:20    4   200  0  14
```

On R5

```
R5#Show ip eigrp neighbors 900

IP-EIGRP neighbors for process 900
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   10.1.56.6               Fa0/0             12 00:01:20    1   200  0  4
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
1   10.1.56.3               Fa0/0              11 00:04:03    3   200  0  8
```

On BB3

```
BB3#Show ip eigrp neighbors 900

IP-EIGRP neighbors for process 900
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   10.1.56.6               Fa0/0             12 00:01:39    1   200  0  5
0   10.1.56.5               Fa0/0             10 00:04:22    5   200  0  12
```

On BB3

```
BB3#Show ip route eigrp 900

D    199.1.1.0/24 [90/156160] via 10.1.56.6, 00:01:56, FastEthernet0/0
                  [90/156160] via 10.1.56.5, 00:01:56, FastEthernet0/0
```

Task 42

BB1 should be configured NOT to have any internal Eigrp Routes in its routing table.

On BB1

```
BB1#show ip route eigrp

     102.0.0.0/23 is subnetted, 1 subnets
D EX    102.1.34.0 [170/2300416] via 10.11.11.1, 19:13:36, FastEthernet0/0
D EX 1.0.0.0/8 [170/156160] via 10.11.11.1, 19:13:36, FastEthernet0/0
     100.0.0.0/24 is subnetted, 8 subnets
D EX    100.1.5.0 [170/156160] via 10.11.11.1, 19:13:36, FastEthernet0/0
D EX    100.1.4.0 [170/156160] via 10.11.11.1, 19:13:36, FastEthernet0/0
D EX    100.1.7.0 [170/156160] via 10.11.11.1, 19:13:36, FastEthernet0/0
D EX    100.1.6.0 [170/156160] via 10.11.11.1, 19:13:36, FastEthernet0/0
D EX    100.1.1.0 [170/156160] via 10.11.11.1, 19:13:36, FastEthernet0/0
D EX    100.1.0.0 [170/156160] via 10.11.11.1, 19:13:36, FastEthernet0/0

D EX    100.1.3.0 [170/156160] via 10.11.11.1, 19:13:36, FastEthernet0/0
D EX    100.1.2.0 [170/156160] via 10.11.11.1, 19:13:36, FastEthernet0/0
D EX 4.0.0.0/8 [170/12175872] via 10.11.11.1, 19:13:36, FastEthernet0/0
D EX 5.0.0.0/8 [170/14320128] via 10.11.11.1, 19:13:36, FastEthernet0/0
     112.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D EX    112.2.2.22/32 [170/2302976] via 10.11.11.1, 19:13:37, FastEthernet0/0
D EX 130.3.0.0/16 [170/2300416] via 10.11.11.1, 19:13:37, FastEthernet0/0
     10.0.0.0/8 is variably subnetted, 14 subnets, 2 masks
D EX    10.1.14.0/24 [170/12047872] via 10.11.11.1, 19:13:37, FastEthernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
D EX    10.1.13.3/32 [170/2588160] via 10.11.11.1, 19:13:37, FastEthernet0/0
D EX    10.1.13.0/24 [170/2588160] via 10.11.11.1, 19:13:37, FastEthernet0/0
D EX    10.1.12.0/24 [170/2172416] via 10.11.11.1, 19:13:37, FastEthernet0/0
D       10.11.12.0/24 [90/30720] via 112.112.112.22, 19:13:37, FastEthernet0/1
D EX    10.1.4.0/24 [170/12050432] via 10.11.11.1, 19:13:39, FastEthernet0/0
D EX    10.1.24.0/24 [170/2174976] via 10.11.11.1, 19:13:39, FastEthernet0/0
D EX    10.1.23.0/24 [170/2684416] via 10.11.11.1, 19:13:39, FastEthernet0/0
D EX    10.1.46.0/24 [170/14192128] via 10.11.11.1, 19:13:39, FastEthernet0/0
D EX    10.1.45.0/24 [170/14192128] via 10.11.11.1, 19:13:39, FastEthernet0/0
D EX    10.1.34.0/24 [170/14192128] via 10.11.11.1, 19:13:39, FastEthernet0/0
D EX    10.1.56.0/24 [170/14194688] via 10.11.11.1, 19:13:39, FastEthernet0/0
D EX    10.1.100.0/24 [170/2172416] via 10.11.11.1, 19:13:39, FastEthernet0/0
D*EX 0.0.0.0/0 [170/2716160] via 10.11.11.1, 19:13:39, FastEthernet0/0
```

On BB1

```
BB1(config)#Router eigrp 600
BB1(config-router)#Distance Eigrp 255 170
```

To verify the configuration:

On BB1

```
BB1#Show ip route eigrp

     102.0.0.0/23 is subnetted, 1 subnets
D EX    102.1.34.0 [170/2300416] via 10.11.11.1, 00:00:39, FastEthernet0/0
D EX 1.0.0.0/8 [170/156160] via 10.11.11.1, 00:00:39, FastEthernet0/0
     100.0.0.0/24 is subnetted, 8 subnets
D EX    100.1.5.0 [170/156160] via 10.11.11.1, 00:00:39, FastEthernet0/0
D EX    100.1.4.0 [170/156160] via 10.11.11.1, 00:00:39, FastEthernet0/0
D EX    100.1.7.0 [170/156160] via 10.11.11.1, 00:00:39, FastEthernet0/0
D EX    100.1.6.0 [170/156160] via 10.11.11.1, 00:00:39, FastEthernet0/0
D EX    100.1.1.0 [170/156160] via 10.11.11.1, 00:00:39, FastEthernet0/0
D EX    100.1.0.0 [170/156160] via 10.11.11.1, 00:00:39, FastEthernet0/0
D EX    100.1.3.0 [170/156160] via 10.11.11.1, 00:00:39, FastEthernet0/0
D EX    100.1.2.0 [170/156160] via 10.11.11.1, 00:00:39, FastEthernet0/0
D EX 4.0.0.0/8 [170/12175872] via 10.11.11.1, 00:00:39, FastEthernet0/0
D EX 5.0.0.0/8 [170/14320128] via 10.11.11.1, 00:00:39, FastEthernet0/0
     112.0.0.0/8 is variably subnetted, 3 subnets, 2 masks

D EX    112.2.2.22/32 [170/2302976] via 10.11.11.1, 00:00:40, FastEthernet0/0
D EX 130.3.0.0/16 [170/2300416] via 10.11.11.1, 00:00:40, FastEthernet0/0
     10.0.0.0/8 is variably subnetted, 14 subnets, 2 masks
D EX    10.1.14.0/24 [170/12047872] via 10.11.11.1, 00:00:40, FastEthernet0/0
D EX    10.1.13.3/32 [170/2588160] via 10.11.11.1, 00:00:40, FastEthernet0/0
D EX    10.1.13.0/24 [170/2588160] via 10.11.11.1, 00:00:40, FastEthernet0/0
D EX    10.1.12.0/24 [170/2172416] via 10.11.11.1, 00:00:40, FastEthernet0/0
D EX    10.11.12.0/24 [170/2174976] via 10.11.11.1, 00:00:40, FastEthernet0/0
D EX    10.1.4.0/24 [170/12050432] via 10.11.11.1, 00:00:42, FastEthernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
D EX     10.1.24.0/24 [170/2174976] via 10.11.11.1, 00:00:42, FastEthernet0/0
D EX     10.1.23.0/24 [170/2684416] via 10.11.11.1, 00:00:42, FastEthernet0/0
D EX     10.1.46.0/24 [170/14192128] via 10.11.11.1, 00:00:42, FastEthernet0/0
D EX     10.1.45.0/24 [170/14192128] via 10.11.11.1, 00:00:42, FastEthernet0/0
D EX     10.1.34.0/24 [170/14192128] via 10.11.11.1, 00:00:42, FastEthernet0/0
D EX     10.1.56.0/24 [170/14194688] via 10.11.11.1, 00:00:42, FastEthernet0/0
D EX     10.1.100.0/24 [170/2172416] via 10.11.11.1, 00:00:42, FastEthernet0/0
D*EX 0.0.0.0/0 [170/2716160] via 10.11.11.1, 00:00:42, FastEthernet0/0
```

Task 43

Configure R1 to filter existing and future routes from any AS that have a tag value of 999.

On R1

To see the prefixes that are tagged with 999:

```
R1#Show ip Eigrp Topology | Inc AS|999

IP-EIGRP Topology Table for AS(300)/ID(130.100.100.1)
P 200.200.200.0/24, 1 successors, FD is 2572534528, tag is 999

IP-EIGRP Topology Table for AS(500)/ID(130.100.100.1)
IP-EIGRP Topology Table for AS(600)/ID(130.100.100.1)
P 200.200.200.0/24, 1 successors, FD is 2572534528, tag is 999

R1#Show ip route 200.200.200.0 | Inc Route tag

      Route tag 999
```

The output of the above show commands reveal that R1 received Network 200.200.200.1 /24 from two different ASes and its tagged with 999.

The following command displays all the ASes configured on R1:

```
R1#Sh ip proto | Inc Routing Protocol

Routing Protocol is "eigrp 300"
Routing Protocol is "eigrp 500"
Routing Protocol is "eigrp 600"
```

To filter existing and future routes that are tagged with 999:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1

```
R1(config)#Route-map Task43 deny 10
R1(config-route-map)#match tag 999
R1(config)#Route-map Task43 permit 90

R1(config)#Router eigrp 300
R1(config-router)#Distribute-list Route-map Task43 in

R1(config)#Router eigrp 500
R1(config-router)#Distribute-list Route-map Task43 in

R1(config)#Router eigrp 600
R1(config-router)#Distribute-list Route-map Task43 in
```

On BB1

```
R1#Show ip route 200.200.200.0 | Inc Route tag
R1#

R1#Show ip Eigrp Topology | Inc AS|999

IP-EIGRP Topology Table for AS(300)/ID(130.100.100.1)
IP-EIGRP Topology Table for AS(500)/ID(130.100.100.1)
IP-EIGRP Topology Table for AS(600)/ID(130.100.100.1)
```

Task 44

Configure BB3 in AS 900 such that it filters network 199.1.1.0 /24 advertised by R5. DO NOT use a "route-map" or a "distance" command to accomplish this task.

On BB3

```
BB3#Show ip route eigrp 900


D    199.1.1.0/24 [90/156160] via 10.1.56.6, 01:08:32, FastEthernet0/0
                  [90/156160] via 10.1.56.5, 01:08:32, FastEthernet0/0
```

Since BB3 ONLY has a single interface through which it receives network 199.1.1.0 /24 from R5 and R6, and the task states that a distance command or a route-map can NOT be used to accomplish this task; an extended access-list can be used to match Network 199.1.1.0 /24 advertised by R5, extended access-list can be used in this manner in Eigrp, RIPv2 and BGP:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The router that originated the route

```
BB3(config)#Access-list 100 deny ip host 10.1.56.5 host 199.1.1.0
BB3(config)#Access-list 100 permit ip any any
```

The network that needs to be filtered

```
BB3(config)#Router eigrp 900
BB3(config-router)#Distribute-list 100 in F0/0
```

To verify the configuration:

On BB3

```
BB3#Show ip route eigrp 900

D    199.1.1.0/24 [90/156160] via 10.1.56.6, 00:07:08, FastEthernet0/0
```

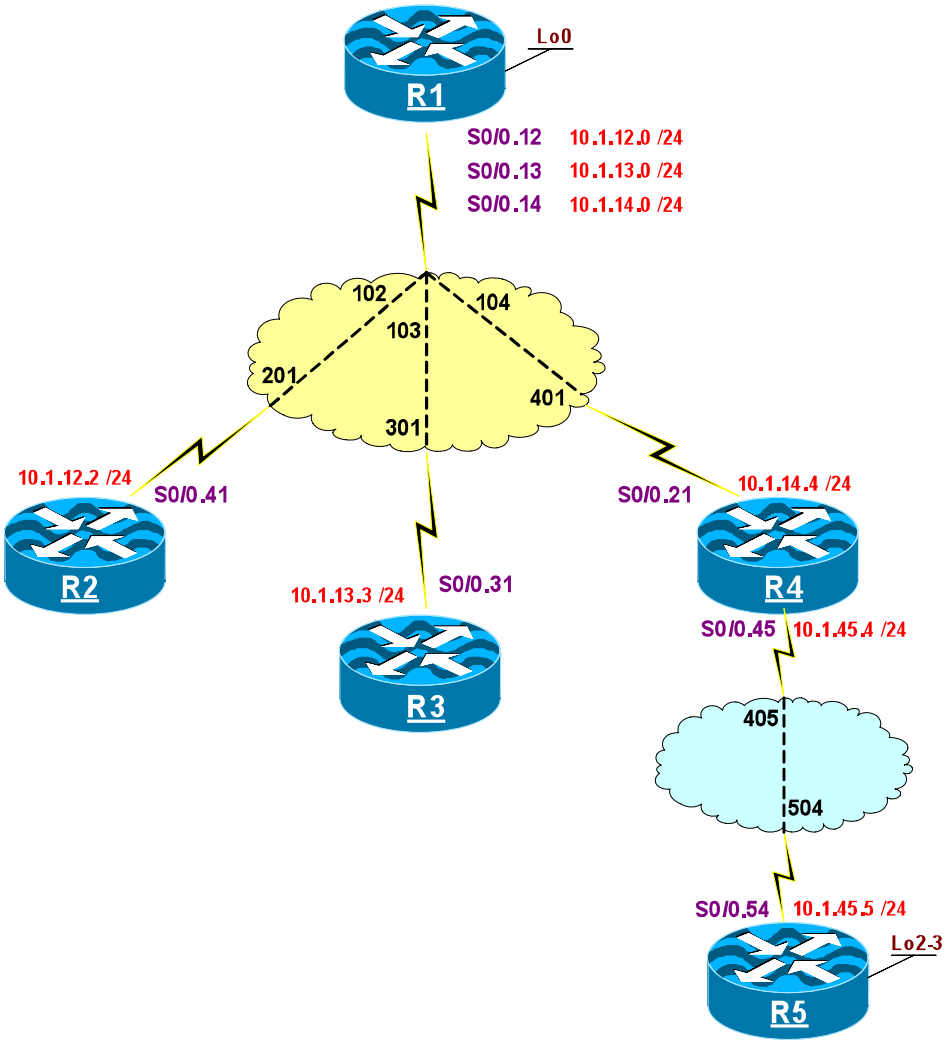
Task 45

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Lab 2
Advanced EIGRP Stub Configuration



Lab Setup:

- Configure all frame-relay connections in a point-to-point sub-interface manner.
- Use the IP addressing chart below for IP address assignment

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP Addressing:

Router	Interface / IP address	Connecting to router:
R1	**S0/0.12 – 10.1.12.1 /24** **S0/0.13 – 10.1.13.1 /24** **S0/0.14 – 10.1.14.1 /24** **Loopback 0 – 1.1.1.1 /24**	**R2** **R3** **R4**
R2	S0/0.21 – 10.1.12.2 /24	R1
R3	**S0/0.31 – 10.1.13.3 /24**	**R1**
R4	S0/0.41 – 10.1.14.4 /24 S0/0.45 – 10.1.45.4 /24	R1 R5
R5	**S0/0.54 – 10.1.45.5 /24** **Loopback 2 – 2.2.2.2 /24** **Loopback 3 – 3.3.3.3 /24**	**R4**

Task 1

Configure OSPF area 0 on the following routers/interfaces; ensure that the loopback interfaces are advertised with their correct mask:

Router	Interface
R4	S0/0.45
R5	**S0/0.54** **Loopback 2** **Loopback 3**

On R4:

```
R4(config)#router ospf 1
R4(config-router)#netw 10.1.45.4 0.0.0.0 area 0
```

On R5:

```
R5(config)#int lo2
R5(config-if)#ip ospf netw point-to-point

R5(config-if)#int lo3
R5(config-if)#ip ospf netw point-to-point

R5(config)#router ospf 1
R5(config-router)#netw 10.1.45.5 0.0.0.0 area 0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R5(config-router)#netw 2.2.2.2 0.0.0.0 area 0
R5(config-router)#netw 3.3.3.3 0.0.0.0 area 0
```

To verify the configuration:

On R4:

```
R4#show ip route ospf

     2.0.0.0/24 is subnetted, 1 subnets
O       2.2.2.0 [110/65] via 10.1.45.5, 00:01:18, Serial0/0.45
     3.0.0.0/24 is subnetted, 1 subnets
O       3.3.3.0 [110/65] via 10.1.45.5, 00:01:18, Serial0/0.45
```

Task 2

Configure EIGRP 100 on the following routers/interfaces; disable auto summarization.

Router	Interface
R1	S0/0.12 S0/0.13 S0/0.14 Loopback 0
R2	S0/0.21
R3	S0/0.31
R4	S0/0.41

On R1:

```
R1(config)#router eigrp 100
R1(config-router)#NO au
R1(config-router)#netw 10.1.12.1 0.0.0.0
R1(config-router)#netw 10.1.13.1 0.0.0.0
R1(config-router)#netw 10.1.14.1 0.0.0.0
R1(config-router)#netw 1.1.1.1 0.0.0.0
```

On R2:

```
R2(config)#router eigrp 100
R2(config-router)#NO au
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config-router)#netw 10.1.12.2 0.0.0.0
```

On R3:

```
R3(config)#router eigrp 100
R3(config-router)#NO au
R3(config-router)#netw 10.1.13.3 0.0.0.0
```

On R4:

```
R4(config)#router eigrp 100
R4(config-router)#NO au
R4(config-router)#netw 10.1.14.4 0.0.0.0
```

To verify the configuration:

On R4:

```
R4#show ip route eigrp

     1.0.0.0/24 is subnetted, 1 subnets
D       1.1.1.0 [90/2297856] via 10.1.14.1, 00:04:29, Serial0/0.41
     10.0.0.0/24 is subnetted, 4 subnets
D       10.1.13.0 [90/2681856] via 10.1.14.1, 00:14:24, Serial0/0.41
D       10.1.12.0 [90/2681856] via 10.1.14.1, 00:14:24, Serial0/0.41
```

On R3:

```
R3#show ip route eigrp

    1.0.0.0/24 is subnetted, 1 subnets
D       1.1.1.0 [90/2297856] via 10.1.13.1, 00:03:58, Serial0/0.31
     10.0.0.0/24 is subnetted, 3 subnets
D       10.1.14.0 [90/2681856] via 10.1.13.1, 00:14:48, Serial0/0.31
D       10.1.12.0 [90/2681856] via 10.1.13.1, 00:14:48, Serial0/0.31
```

On R2:

```
R2#show ip route eigrp

     1.0.0.0/24 is subnetted, 1 subnets
D       1.1.1.0 [90/2297856] via 10.1.12.1, 00:03:08, Serial0/0.21
     10.0.0.0/24 is subnetted, 3 subnets
D       10.1.14.0 [90/2681856] via 10.1.12.1, 00:15:16, Serial0/0.21
D       10.1.13.0 [90/2681856] via 10.1.12.1, 00:15:16, Serial0/0.21
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 3

Configure mutual redistribution between OSPF and EIGRP On R4:; use a metric of your choice.

On R4:

```
R4(config)#router ospf 1
R4(config-router)#redistribute eigrp 100 subnets

R4(config)#router eigrp 100
R4(config-router)#redistribute ospf 1 metric 1 1 1 1 1
```

To verify the configuration:

On R5:

```
R5#show ip route ospf | inc O

O E2    1.1.1.0 [110/20] via 10.1.45.4, 00:06:00, Serial0/0.54
O E2    10.1.14.0 [110/20] via 10.1.45.4, 00:07:39, Serial0/0.54
O E2    10.1.13.0 [110/20] via 10.1.45.4, 00:07:39, Serial0/0.54
O E2    10.1.12.0 [110/20] via 10.1.45.4, 00:07:39, Serial0/0.54
```

On R1:

```
R1#show ip route eigrp

     2.0.0.0/24 is subnetted, 1 subnets
D EX    2.2.2.0 [170/2560512256] via 10.1.14.4, 00:01:00, Serial0/0.14
     3.0.0.0/24 is subnetted, 1 subnets
D EX    3.3.3.0 [170/2560512256] via 10.1.14.4, 00:01:00, Serial0/0.14
     10.0.0.0/24 is subnetted, 4 subnets
D EX    10.1.45.0 [170/2560512256] via 10.1.14.4, 00:01:00, Serial0/0.14
```

On R2:

```
R2#show ip route eigrp

     1.0.0.0/24 is subnetted, 1 subnets
D       1.1.1.0    [90/2297856] via 10.1.12.1, 00:07:33, Serial0/0.21
     2.0.0.0/24 is subnetted, 1 subnets
D EX    2.2.2.0 [170/2561024256] via 10.1.12.1, 00:01:51, Serial0/0.21
     3.0.0.0/24 is subnetted, 1 subnets
D EX    3.3.3.0    [170/2561024256] via 10.1.12.1, 00:01:51, Serial0/0.21
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
     10.0.0.0/24 is subnetted, 4 subnets
D       10.1.14.0    [90/2681856] via 10.1.12.1, 00:19:41, Serial0/0.21
D       10.1.13.0    [90/2681856] via 10.1.12.1, 00:19:41, Serial0/0.21
D EX    10.1.45.0 [170/2561024256] via 10.1.12.1, 00:01:51, Serial0/0.21
```

On R3:

```
R3#show ip route eigrp

     1.0.0.0/24 is subnetted, 1 subnets
D       1.1.1.0 [90/2297856] via 10.1.13.1, 00:10:45, Serial0/0.31
     2.0.0.0/24 is subnetted, 1 subnets
D EX    2.2.2.0 [170/2561024256] via 10.1.13.1, 00:05:02, Serial0/0.31
     3.0.0.0/24 is subnetted, 1 subnets
D EX    3.3.3.0 [170/2561024256] via 10.1.13.1, 00:05:02, Serial0/0.31
     10.0.0.0/24 is subnetted, 4 subnets
D       10.1.14.0 [90/2681856] via 10.1.13.1, 00:21:34, Serial0/0.31
D       10.1.12.0 [90/2681856] via 10.1.13.1, 00:21:34, Serial0/0.31
D EX    10.1.45.0 [170/2561024256] via 10.1.13.1, 00:05:02, Serial0/0.31
```

Note: R2 and R3 have networks 2.2.2.0/24 and 3.3.3.0/24.

Task 4

Configure "EIGRP stub" on R1 such that it **only** advertises its directly connected interfaces that are advertised with a "network" command to its EIGRP neighbors.

On R1:

```
R1(config)#router eigrp 100
R1(config-router)#eigrp stub connected
```

To verify the configuration:

On R2:

```
R2#show ip route eigrp

    1.0.0.0/24 is subnetted, 1 subnets
D       1.1.1.0 [90/2297856] via 10.1.12.1, 00:02:01, Serial0/0.21
     10.0.0.0/24 is subnetted, 3 subnets
D       10.1.14.0 [90/2681856] via 10.1.12.1, 00:02:01, Serial0/0.21
D       10.1.13.0 [90/2681856] via 10.1.12.1, 00:02:01, Serial0/0.21
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3:

```
R3#show ip route eigrp

     1.0.0.0/24 is subnetted, 1 subnets
D       1.1.1.0 [90/2297856] via 10.1.13.1, 00:02:41, Serial0/0.31
     10.0.0.0/24 is subnetted, 3 subnets
D       10.1.14.0 [90/2681856] via 10.1.13.1, 00:02:41, Serial0/0.31
D       10.1.12.0 [90/2681856] via 10.1.13.1, 00:02:41, Serial0/0.31
```

On R1:

```
R1#show ip route eigrp

     2.0.0.0/24 is subnetted, 1 subnets
D EX    2.2.2.0 [170/2560512256] via 10.1.14.4, 00:03:25, Serial0/0.14
     3.0.0.0/24 is subnetted, 1 subnets
D EX    3.3.3.0 [170/2560512256] via 10.1.14.4, 00:03:25, Serial0/0.14
     10.0.0.0/24 is subnetted, 4 subnets
D EX    10.1.45.0 [170/2560512256] via 10.1.14.4, 00:03:25, Serial0/0.14
```

Note: R2 and R3 do not have the redistributed routes; this is because of R1's "stub connected" configuration.

Task 5

Configure R1 such that routers R2 and R3 have networks 2.2.2.0 /24 and 3.3.3.0 /24 in their routing table. **Do not** remove the "eigrp stub connected" configuration from R1 to accomplish this task.

To accomplish this task a "leak-map" is referenced in the "eigrp stub connected" command. The leak-map references a route-map called "TST", the route-map references an access-list. Any IP address/es that are permitted in the access-list is leaked along the connected networks.

On R1:

```
R1(config)#access-list 1 permit 2.2.2.0 0.0.0.255
R1(config)#access-list 1 permit 3.3.3.0 0.0.0.255

R1(config)#route-map TST permit 10
R1(config-route-map)#match ip addr 1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config)#router eigrp 100
R1(config-router)#eigrp stub connected leak-map TST
```

To verify the configuration:

On R2:

```
R2#Show ip route eigrp | inc EX

D EX     2.2.2.0 [170/2561024256] via 10.1.12.1, 00:04:29, Serial0/0.21
D EX     3.3.3.0 [170/2561024256] via 10.1.12.1, 00:04:29, Serial0/0.21
```

On R3:

```
R3#Show ip route eigrp | inc EX

D EX     2.2.2.0 [170/2561024256] via 10.1.13.1, 00:00:13, Serial0/0.31
D EX     3.3.3.0 [170/2561024256] via 10.1.13.1, 00:00:13, Serial0/0.31
```

Note: Both R2 and R3 have both networks (2.2.2.0 /24 and 3.3.3.0 /24) in their routing table.

Task 6

Re-configure R1 such that R2 gets network 2.2.2.0 /24 and R3 gets network 3.3.3.0 /24.
Do not remove the "eigrp stub connected" configuration from R1 to accomplish this task.

The access-list and the route-map TST should be removed before proceeding further:

On R1:

```
R1(config)#no access-list 1
R1(config)#no route-map TST
```

We should identify the two networks using two different access-lists, in this case access-list 2 identifies network 2.2.2.0 /24 and access-list 3 identifies network 3.3.3.0 /24:

```
R1(config)#access-list 2 permit 2.2.2.0 0.0.0.255
R1(config)#access-list 3 permit 3.3.3.0 0.0.0.255
```

The next step is to configure a new route-map as follows:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config)#route-map TST permit 10
R1(config-route-map)#match ip addr 2
R1(config-route-map)#match inter S0/0.12

R1(config-route-map)#route-map TST permit 20
R1(config-route-map)#match ip addr 3
R1(config-route-map)#match inter S0/0.13
```

To verify the configuration:

On R2:

```
R2#Show ip route eigrp | inc EX

D EX    2.2.2.0 [170/2561024256] via 10.1.12.1, 00:03:16, Serial0/0.21
```

On R3:

```
R3#Show ip route eigrp | inc EX

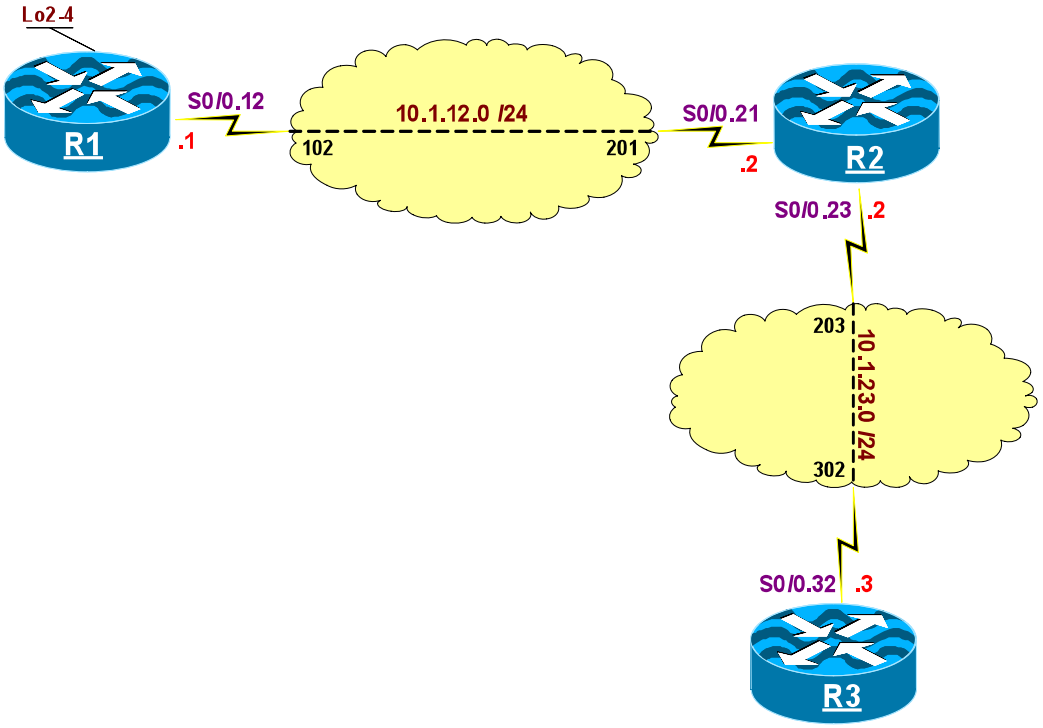
D EX    3.3.3.0 [170/2561024256] via 10.1.13.1, 00:03:43, Serial0/0.31
```

Task 7

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 3 – EIGRP & Default-Information



Lab Setup:

- Configure all frame-relay connections in a point-to-point manner.
- Use the IP addressing and DLCI chart below.

IP Addressing:

Router	Interface / IP address	DLCI assignment
R1	S0/0.12 = 10.1.12.1 /24 Loopback2 = 2.2.2.2 /8 Loopback3 = 3.3.3.3 /8 Loopback4 = 4.4.4.4 /8	102
R2	**S0/0.21 = 10.1.12.2 /24** **S0/0.23 = 10.1.23.2 /24**	**201** **203**
R3	S0/0.32 = 10.1.23.3 /24	302

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 1

Configure EIGRP on all routers and advertise their directly connected interfaces in AS 100. R1 should **not** advertise network 4.0.0.0 /8 in this routing protocol.

On R1:

```
R1(config)#router eigrp 100
R1(config-router)#No au
R1(config-router)#netw 2.2.2.2 0.0.0.0
R1(config-router)#netw 3.3.3.3 0.0.0.0
R1(config-router)#netw 10.1.12.1 0.0.0.0
```

On R2:

```
R2(config)#router eigrp 100
R2(config-router)#No au
R2(config-router)#network 10.1.12.2 0.0.0.0
R2(config-router)#network 10.1.23.2 0.0.0.0
```

On R3:

```
R3(config)#router eigrp 100
R3(config-router)#No au
R3(config-router)#network 10.1.23.3 0.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip route eigrp

     10.0.0.0/24 is subnetted, 2 subnets
D       10.1.23.0 [90/2681856] via 10.1.12.2, 00:01:55, Serial0/0.12
```

On R2:

```
R2#Show ip route eigrp

D    2.0.0.0/8 [90/2297856] via 10.1.12.1, 00:02:21, Serial0/0.21
D    3.0.0.0/8 [90/2297856] via 10.1.12.1, 00:02:21, Serial0/0.21
```

On R3:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3#Show ip route eigrp

D    2.0.0.0/8 [90/2809856] via 10.1.23.2, 00:01:57, Serial0/0.32
D    3.0.0.0/8 [90/2809856] via 10.1.23.2, 00:01:57, Serial0/0.32
     10.0.0.0/24 is subnetted, 2 subnets
D       10.1.12.0 [90/2681856] via 10.1.23.2, 00:01:57, Serial0/0.32
```

Task 2

Configure R1 such that R2 and R3 use network 2.0.0.0 /8 as a candidate default.

On R1:

```
R1(config)#ip default-network 2.0.0.0
```

To verify the configuration:

On R1:

```
R1#Sh ip route 2.0.0.0

Routing entry for 2.0.0.0/8
  Known via "connected", distance 0, metric 0 (connected, via interface),
candidate default path
  Redistributing via eigrp 100
  Routing Descriptor Blocks:
  * directly connected, via Loopback2
      Route metric is 0, traffic share count is 1

R1#Show ip route | b Gate

C*   2.0.0.0/8 is directly connected, Loopback2
C    3.0.0.0/8 is directly connected, Loopback3
C    4.0.0.0/8 is directly connected, Loopback4
     10.0.0.0/24 is subnetted, 2 subnets
C       10.1.12.0 is directly connected, Serial0/0.12
D       10.1.23.0 [90/2681856] via 10.1.12.2, 00:06:11, Serial0/0.12
```

On R2:

```
R2#Show ip route eigrp
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
D*   2.0.0.0/8 [90/2297856] via 10.1.12.1, 00:00:41, Serial0/1.21
D     3.0.0.0/8 [90/2297856] via 10.1.12.1, 00:07:55, Serial0/1.21
```

Note: the asterisk reveals that R2 is using that network as the candidate default.

On R3:

```
R3#Sh ip route eigrp

D*   2.0.0.0/8 [90/2809856] via 10.1.23.2, 00:02:56, Serial0/0.32
D    3.0.0.0/8 [90/2809856] via 10.1.23.2, 00:05:48, Serial0/0.32
     10.0.0.0/24 is subnetted, 2 subnets
D       10.1.12.0 [90/2681856] via 10.1.23.2, 00:06:46, Serial0/0.32
```

To test the configuration On R2:

On R2:

```
R2#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/26/60 ms
```

To test the configuration On R3:

On R3:

```
R3#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/32/92 ms
```

Task 3

Configure R2 such that R3 does **not** use network 2.0.0.0 /8 as the candidate default, R3 should still have network 2.0.0.0 /8 in its routing table.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2:

```
R2(config)#router eigrp 100
R2(config-router)#No default-information allowed out

R2#Clear ip eigrp neighbor
```

Note: The "no default-information allowed out" disables the redistribution of default route. Meaning that R3 will no longer use network 2.0.0.0 /8 as its candidate default, but it will still have that network in its routing table.

To verify the configuration:

On R2:

```
R2#Show ip route eigrp

D*    2.0.0.0/8 [90/2297856] via 10.1.12.1, 00:00:25, Serial0/1.21
D     3.0.0.0/8 [90/2297856] via 10.1.12.1, 00:00:25, Serial0/1.21
```

On R3:

```
R3#Show ip route eigrp

D     2.0.0.0/8 [90/2809856] via 10.1.23.2, 00:01:04, Serial0/1.32
D     3.0.0.0/8 [90/2809856] via 10.1.23.2, 00:01:04, Serial0/1.32
      10.0.0.0/24 is subnetted, 2 subnets
D        10.1.12.0 [90/2681856] via 10.1.23.2, 00:01:06, Serial0/1.32
```

To test the configuration:

On R3:

```
R3#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 4

Remove the command that was entered in the previous task on R2, and configure R3 to accomplish the same task, if this configuration is performed correctly, R3 should **not** use network 2.0.0.0 /8 as the candidate default.

To remove the command from R2:

On R2:

```
R2(config)#router eigrp 100
R2(config-router)#default-information allowed out

R2#Clear ip eigrp neighbor
```

To verify R2's configuration:

On R2:

```
R2#Sh run | s router eigrp

router eigrp 100
 network 10.1.12.2 0.0.0.0
 network 10.1.23.2 0.0.0.0
 no auto-summary
```

To verify the configuration on R3

On R3:

```
R3#Show ip route eigrp

D*    2.0.0.0/8 [90/2809856] via 10.1.23.2, 00:00:58, Serial0/1.32
D     3.0.0.0/8  [90/2809856] via 10.1.23.2, 00:00:58, Serial0/1.32
      10.0.0.0/24 is subnetted, 2 subnets
D        10.1.12.0 [90/2681856] via 10.1.23.2, 00:01:00, Serial0/1.32
```

Note: R3 uses network 2.0.0.0 /8 as candidate default. To configure R3 to accomplish the same task:

On R3:

```
R3(config)#router eigrp 100
R3(config-router)#No default-information allowed in
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3#Clear ip eigrp neighbor
```

Note: From R3's perspective it should disable the redistribution of the default route inbound, therefore, the direction of the command is configured inbound.

Note: If R2 enters the "no default-information allowed in", then R3 will not see network 2.0.0.0 as a candidate default either.

To test and verify the configuration:

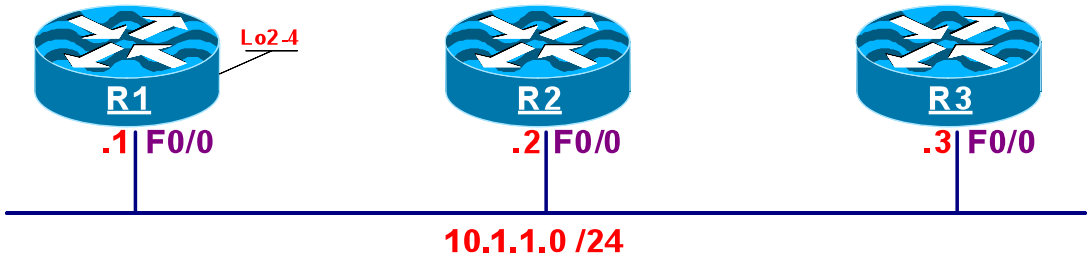
On R3:

```
R3#Show ip route eigrp

D    2.0.0.0/8 [90/2809856] via 10.1.23.2, 00:01:17, Serial0/0.32
D    3.0.0.0/8 [90/2809856] via 10.1.23.2, 00:01:17, Serial0/0.32
     10.0.0.0/24 is subnetted, 2 subnets
D       10.1.12.0 [90/2681856] via 10.1.23.2, 00:01:17, Serial0/0.32
```

Task 5

Reconfigure the routers based on the following topology and IP addressing:



You can copy and paste the initial configuration for this task from the initial configuration file provided.

Lab Setup:

- Configure the F0/0 interfaces of R1, R2 and R3 in VLAN 100.
- Use the IP addressing chart below.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP Addressing:

Router	Interface / IP address
R1	F0/0 = 10.1.1.1 /24 Loopback2 = 2.2.2.2 /8 Loopback3 = 3.3.3.3 /8 Loopback4 = 4.4.4.4 /8
R2	**F0/0 = 10.1.1.2 /24**
R3	F0/0 = 10.1.1.3 /24

Task 6

Configure EIGRP 100 on the routers and advertise their directly connected interfaces in AS 100. R1 should **not** advertise network 4.0.0.0 /8 in this routing protocol.

On R1:

```
R1(config-if)#router eigrp 100
R1(config-router)#No au
R1(config-router)#network 2.0.0.0
R1(config-router)#network 3.0.0.0
R1(config-router)#network 10.1.1.1 0.0.0.0
```

On R2:

```
R2(config)#router eigrp 100
R2(config-router)#No au
R2(config-router)#network 10.1.1.2 0.0.0.0
```

On R3:

```
R3(config)#router eigrp 100
R3(config-router)#No au
R3(config-router)#network 10.1.1.3 0.0.0.0
```

To verify the configuration

On R2:

```
R2#Show ip route eigrp
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
D    2.0.0.0/8 [90/409600] via 10.1.1.1, 00:15:21, Ethernet0/0
D    3.0.0.0/8 [90/409600] via 10.1.1.1, 00:15:21, Ethernet0/0
```

On R3:

```
R3#Show ip route eigrp

D    2.0.0.0/8 [90/409600] via 10.1.1.1, 00:15:06, Ethernet0/0
D    3.0.0.0/8 [90/409600] via 10.1.1.1, 00:15:06, Ethernet0/0
```

Task 7

Configure R1 to advertise network 2.0.0.0 /8 and network 3.0.0.0 /8 as two candidate defaults in this routing domain.

On R1:

```
R1(config)#ip default-network 2.0.0.0
R1(config)#ip default-network 3.0.0.0
```

To verify the configuration:

On R2:

```
R2#Show ip route eigrp

D*   2.0.0.0/8 [90/156160] via 10.1.1.1, 00:00:12, FastEthernet0/0
D*   3.0.0.0/8 [90/156160] via 10.1.1.1, 00:00:08, FastEthernet0/0
```

On R3:

```
R3#Show ip route eigrp

D*   2.0.0.0/8 [90/156160] via 10.1.1.1, 00:00:38, FastEthernet0/0
D*   3.0.0.0/8 [90/156160] via 10.1.1.1, 00:00:34, FastEthernet0/0
```

Note: Both R2 and R3 use networks 2.0.0.0 /8 and 3.0.0.0 /8 as their candidate default.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 8

Configure R2 and R3 such that R2 uses network 2.0.0.0 /8 and R3 uses network 3.0.0.0 /8 as their candidate default.

On R2:

To configure this task, an access-list is configured to identify the network (network 2.0.0.0 /8 in this case), then, the access-list is referenced in the "default-information allowed in" command. Which tells the router that only the network that is permitted in the access-list should be used as candidate default.

```
R2(config)#access-list 2 permit 2.0.0.0

R2(config)#router eigrp 100
R2(config-router)#default-information allowed in 2

R2#Clear ip eigrp neigh

R2#Show ip route eigrp

D*    2.0.0.0/8 [90/156160] via 10.1.1.1, 00:00:11, FastEthernet0/0
D     3.0.0.0/8 [90/156160] via 10.1.1.1, 00:00:11, FastEthernet0/0
```

The following shows the configuration of R3:

On R3:

```
R3(config)#access-list 3 permit 3.0.0.0

R3(config)#router eigrp 100
R3(config-router)#default-information allowed in 3

R3#Clear ip eigrp neighbor
```

To verify the configuration:

On R3:

```
R3#Show ip route eigrp

D     2.0.0.0/8 [90/409600] via 10.1.1.1, 00:01:01, Ethernet0/0
D*    3.0.0.0/8 [90/409600] via 10.1.1.1, 00:01:01, Ethernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To test the configuration:

On R2:

```
R2#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/11/24 ms
```

On R3:

```
R3#Ping 4.4.4.4

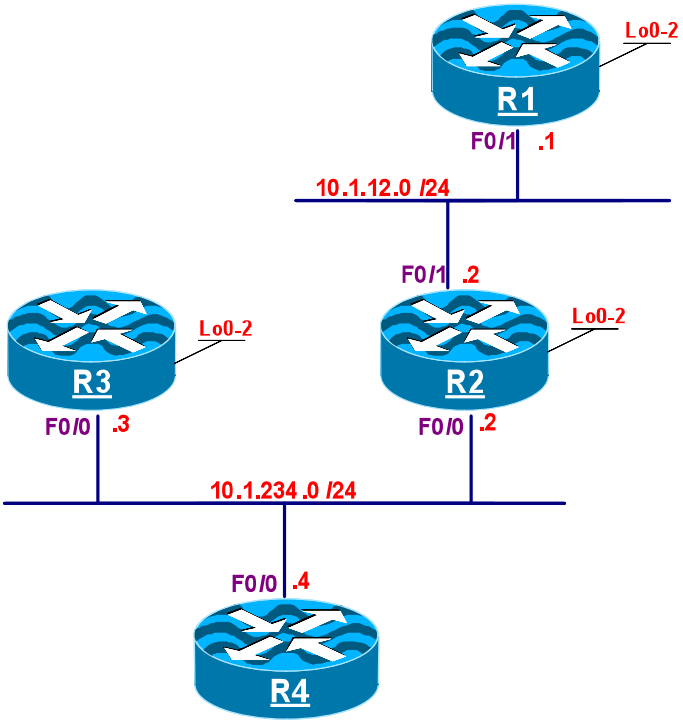
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/13/24 ms
```

Task 9

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 4 – EIGRP Filtering



Lab Setup:

- Configure the F0/1 interface of R1 and R2 in VLAN 12
- Configure the F0/0 interface of R2, R3 and R4 in VLAN 234
- Use the following IP addressing chart for IP assignment:

IP Addressing Chart:

Router	Interface / IP address
R1	F0/1 = 10.1.12.1 /24 Lo0 = 1.1.1.1 /8 Lo1 = 11.1.1.1 /8 Lo2 = 111.1.1.1 /8
R2	**F0/1 = 10.1.12.2 /24** **F0/0 = 10.1.234.2 /24** **Lo0 = 2.2.2.2 /8**

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

	Lo1 = 200.1.1.1 /24 **Lo2 = 200.2.2.2 /24**
R3	F0/0 = 10.1.234.3 /24 Lo0 = 3.3.3.3 /8 Lo1 = 200.1.1.1 /24 Lo2 = 200.2.2.2 /24
R4	**F0/0 = 10.1.234.4 /24**

Task 1

Configure EIGRP 100 on all routers and advertise their directly connected links. You should disable auto summarization on these routers.

On R1:

```
R1(config)#router eigrp 100
R1(config-router)#network 1.1.1.1 0.0.0.0
R1(config-router)#network 10.1.12.1 0.0.0.0
R1(config-router)#network 11.1.1.1 0.0.0.0
R1(config-router)#network 111.1.1.1 0.0.0.0
R1(config-router)#No auto-summary
```

On R2:

```
R2(config)#router eigrp 100
R2(config-router)#network 10.1.12.2 0.0.0.0
R2(config-router)#network 10.1.234.2 0.0.0.0
R2(config-router)#network 200.1.1.1 0.0.0.0
R2(config-router)#network 200.2.2.2 0.0.0.0
R2(config-router)#network 2.2.2.2 0.0.0.0
R2(config-router)#No auto-summary
```

On R3:

```
R3(config)#router eigrp 100
R3(config-router)#network 10.1.234.3 0.0.0.0
R3(config-router)#network 200.1.1.1 0.0.0.0
R3(config-router)#network 200.2.2.2 0.0.0.0
R3(config-router)#network 3.3.3.3 0.0.0.0
R3(config-router)#No auto-summary
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4:

```
R4(config)#router eigrp 100
R4(config-router)#network 10.1.234.4 0.0.0.0
R4(config-router)#No auto-summary
```

To verify the configuration:

On R4:

```
R4#Show ip route eigrp

D    1.0.0.0/8 [90/158720] via 10.1.234.2, 00:00:09, FastEthernet0/0
D    2.0.0.0/8 [90/156160] via 10.1.234.2, 00:00:09, FastEthernet0/0
D    3.0.0.0/8 [90/156160] via 10.1.234.3, 00:00:09, FastEthernet0/0
D    200.1.1.0/24 [90/156160] via 10.1.234.3, 00:00:09, FastEthernet0/0
                  [90/156160] via 10.1.234.2, 00:00:09, FastEthernet0/0
D    200.2.2.0/24 [90/156160] via 10.1.234.3, 00:00:09, FastEthernet0/0
                  [90/156160] via 10.1.234.2, 00:00:09, FastEthernet0/0
D    111.0.0.0/8 [90/158720] via 10.1.234.2, 00:00:09, FastEthernet0/0
     10.0.0.0/24 is subnetted, 2 subnets
D       10.1.12.0 [90/30720] via 10.1.234.2, 00:00:09, FastEthernet0/0
D    11.0.0.0/8 [90/158720] via 10.1.234.2, 00:00:09, FastEthernet0/0
```

Task 2

Configure R4 such that it filters existing (1.0.0.0/8, 11.0.0.0/8 and 111.0.0.0/8) and future network behind R1. **Do not** use "distribute-list" or "route-map" to accomplish this task

By default, EIGRP will discard routes that have a hop count of 101 or more, this behavior can be utilized to accomplish this task, as follows:

On R4:

```
R4#Sh ip eigrp topology 11.0.0.0 255.0.0.0 | inc Hop

        Hop count is 2

R4#Sh ip eigrp topology 1.0.0.0 255.0.0.0 | inc Hop

        Hop count is 2
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4#Sh ip eigrp topology 111.0.0.0 255.0.0.0 | inc Hop

        Hop count is 2
```

Note: The routes behind R1 have a hop count of 2, whereas, the other routes advertised in this topology have a hop count of 1:

```
R4#Sh ip eigrp topology 2.0.0.0 255.0.0.0 | inc Hop

    Hop count is 1

R4#Sh ip eigrp topology 3.0.0.0 255.0.0.0 | inc Hop

    Hop count is 1
```

Therefore, we should reject routes that have a hop count greater than 1:

On R4:

```
R4(config)#router eigrp 100
R4(config-router)#metric maximum-hops 1
```

Note: You should see the following console messages, this is because the policy for EIGRP is changed from 100 (default hop count) to only 1:

```
%DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 10.1.234.2 (FastEthernet0/0) is
down: Max hopcount changed
%DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 10.1.234.3 (FastEthernet0/0) is
down: Max hopcount changed
%DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 10.1.234.2 (FastEthernet0/0) is up:
new adjacency
%DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 10.1.234.3 (FastEthernet0/0) is up:
new adjacency
```

To verify the configuration:

On R4:

```
R4#Show ip route eigrp

D    2.0.0.0/8 [90/156160] via 10.1.234.2, 00:00:09, FastEthernet0/0
D    3.0.0.0/8 [90/156160] via 10.1.234.3, 00:00:09, FastEthernet0/0
D    200.1.1.0/24 [90/156160] via 10.1.234.3, 00:00:09, FastEthernet0/0
                  [90/156160] via 10.1.234.2, 00:00:09, FastEthernet0/0
D    200.2.2.0/24 [90/156160] via 10.1.234.3, 00:00:09, FastEthernet0/0
                  [90/156160] via 10.1.234.2, 00:00:09, FastEthernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
     10.0.0.0/24 is subnetted, 2 subnets
D       10.1.12.0 [90/30720] via 10.1.234.2, 00:00:09, FastEthernet0/0
```

Task 3

Configure R4 such that it **only** takes R2 to reach network 200.1.1.0 /24. R4 should **never** go through R3 to reach network 200.1.1.0/24. You should use an access-list to accomplish this task.

On R4:

Note: In this case an extended access-list can be used to filter a prefix from a given route-source; in the following extended access-list, the source address in the ACL references the advertising neighbor, whereas, the destination address in the ACL references the actual Network:

```
R4(config)#access-list 100 deny ip host 10.1.234.3 host 200.1.1.0
R4(config)#access-list 100 permit ip any any

R4(config)#router eigrp 100
R4(config-router)#distribute-list 100 in f0/0
```

To verify the configuration:

On R4:

```
R4#Sh ip route eigrp

D    2.0.0.0/8 [90/156160] via 10.1.234.2, 00:03:58, FastEthernet0/0
D    3.0.0.0/8 [90/156160] via 10.1.234.3, 00:03:58, FastEthernet0/0
D    200.1.1.0/24 [90/156160] via 10.1.234.2, 00:00:05, FastEthernet0/0
D    200.2.2.0/24 [90/156160] via 10.1.234.3, 00:03:58, FastEthernet0/0
                  [90/156160] via 10.1.234.2, 00:03:58, FastEthernet0/0
     10.0.0.0/24 is subnetted, 2 subnets
D       10.1.12.0 [90/30720] via 10.1.234.2, 00:03:58, FastEthernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 4

Re-configure the solution in task 3 with redundancy, such that if R2 is down, R4 can reach network 200.1.1.0/24 through R3.

In this case the distance can be manipulated to accomplish this task, as follows:

Let's remove the commands from the previous step:

```
R4(config)#router eigrp 100
R4(config-router)#No distribute-list 100 in F0/0

R4(config)#No access-list 100
```

Let's configure the new solution:

Step 1:
Configure an access-list to identify the network:

```
R4(config)#access-list 1 permit 200.1.1.0 0.0.0.255
```

Step 2:
Utilizing the distance command, the AD for network 200.1.1.0/24 through R3 is set higher than the default AD of 90:

```
R4(config-router)#distance 91 10.1.234.3 0.0.0.0 1
```

The above command sets the AD for network 200.1.1.0/24 through R3 to 91.

To verify the configuration:

On R4:

```
R4#Sh ip route eigrp

D     2.0.0.0/8 [90/156160] via 10.1.234.2, 00:01:58, FastEthernet0/0
D     3.0.0.0/8 [90/156160] via 10.1.234.3, 00:01:58, FastEthernet0/0
D     200.1.1.0/24 [90/156160] via 10.1.234.2, 00:01:58, FastEthernet0/0
D     200.2.2.0/24 [90/156160] via 10.1.234.3, 00:01:58, FastEthernet0/0
                   [90/156160] via 10.1.234.2, 00:01:58, FastEthernet0/0
      10.0.0.0/24 is subnetted, 2 subnets
D        10.1.12.0 [90/30720] via 10.1.234.2, 00:01:58, FastEthernet0/0
```

To test the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2:

```
R2(config)#int lo1
R2(config-if)#shut
```

To verify the test:

On R4:

```
R4#Sh ip route eigrp

D    2.0.0.0/8 [90/156160] via 10.1.234.2, 00:02:59, FastEthernet0/0
D    3.0.0.0/8 [90/156160] via 10.1.234.3, 00:02:59, FastEthernet0/0
D    200.1.1.0/24 [91/156160] via 10.1.234.3, 00:00:08, FastEthernet0/0
D    200.2.2.0/24 [90/156160] via 10.1.234.3, 00:02:59, FastEthernet0/0
                  [90/156160] via 10.1.234.2, 00:02:59, FastEthernet0/0
     10.0.0.0/24 is subnetted, 2 subnets
D       10.1.12.0 [90/30720] via 10.1.234.2, 00:02:59, FastEthernet0/0
```

Let's re-enable the Loopback 1 interface on R2:

On R2:

```
R2(config)#int lo1
R2(config-if)#No shut
```

Task 5

Filter network 2.0.0.0/8 On R4; **do not** use the "distribute-list" command to accomplish this task.

Once again, the distance command can be used to accomplish this task. The difference between the solution used in this task and the solution used in the previous task is that the AD is set to a value that is unreachable (255).

On R4:

```
R4(config)#access-list 2 permit 2.0.0.0 0.255.255.255

R4(config)#router eigrp 100
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config-router)#distance 255 10.1.234.2 0.0.0.0 2
```

To verify the configuration:

On R4:

```
R4#Show ip route eigrp

D    3.0.0.0/8 [90/156160] via 10.1.234.3, 00:00:07, FastEthernet0/0
D    200.1.1.0/24 [90/156160] via 10.1.234.2, 00:00:07, FastEthernet0/0
D    200.2.2.0/24 [90/156160] via 10.1.234.3, 00:00:07, FastEthernet0/0
                  [90/156160] via 10.1.234.2, 00:00:07, FastEthernet0/0
     10.0.0.0/24 is subnetted, 2 subnets
D       10.1.12.0 [90/30720] via 10.1.234.2, 00:00:07, FastEthernet0/0
```

Task 6

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Advanced
CCIE SERVICE PROVIDER
v3.0

www.MicronicsTraining.com

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

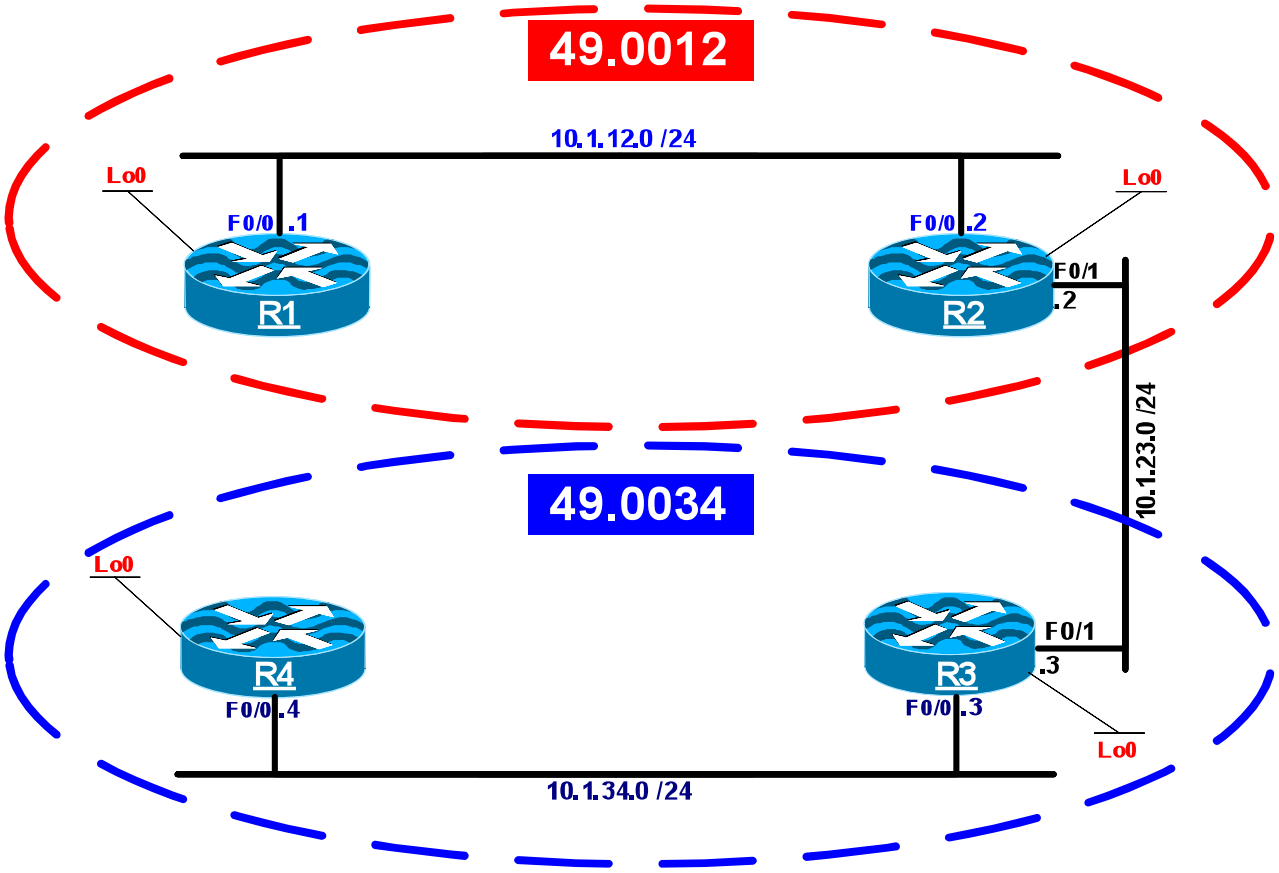
Paul Negron
CCIE #14856
SP

ISIS

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Lab 1
Basic ISIS Configuration



Lab Setup:

- R1 and R2's F0/0 interface should be configured in VLAN 12.
- R2 and R3's F0/1 interface should be configured in VLAN 23.
- R2 and R3's F0/0 interface should be configured in VLAN 34.

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

IP Addressing

Router	Interface	IP address
R1	Lo0 F0/0	1.1.1.1 /8 10.1.12.1 /24
R2	Lo0 F0/0 F0/1	2.2.2.2 /8 10.1.12.2 /24 10.1.23.2 /24
R3	Lo0 F0/0 F0/1	3.3.3.3 /8 10.1.34.3 /24 10.1.23.3 /24
R4	Lo0 F0/0	4.4.4.4 /8 10.1.34.4 /24

Task 1

Configure ISIS on all routers and advertise their directly connected networks. The loopback interfaces should be advertised as well. Do not apply any configuration commands to the loopback interfaces on R1 and R2. R1 and R2 should be in 49.0012, whereas, R3 and R4 should be configured in 49.0034 Domain/Area.

On R4

```
R4(config)#router isis
R4(config-router)#NET 49.0034.4444.0000.4444.00
```

ISIS requires a Network Entity Title (NET) address, in Hexadecimal format, to operate. The NET address is comprised of separate sections:

NSAP = Network Service Access Point defines the 20 Octets by standard, minimum of 8 bytes for Cisco Routers.

The Domain ID of "49" is used differently by Cisco. By standard, the Domain-Id is part of the Inter Domain Part (IDP), which aids in routing data between Domains. Since BGP does a better job of doing that, Cisco combines this value with Area-ID and it is interpreted as a part of the AREA.

The Area ID of "0034" defines an ISIS area. An area is where one ISIS speaker speaks to other ISIS speakers to accomplish routing. The AREA-ID must match between routers that perform Level 1 Routing. Two routers that do not have matching areas will perform inter-area routing. This type of routing also changes to be Level 2.

4444.0000.4444 = The System ID is comprised of 6 bytes that is fixed in size by a Cisco Router. It is used to help route traffic within an area and must be Globally unique within the Domain.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

00 = Network Selector (N-selector) address of "0" is also a fixed value by a Cisco Router and cannot be changed. There is no specific use by a Cisco Router. HP uses a value of "1".

```
R4(config)#Int f0/0
R4(config-if)#ip router isis

R4(config)#int Lo0
R4(config-if)#ip router isis
```

The "ip router isis" command is used under an interface that participates in the ISIS process. This command allows IP routing to occur in an ISIS environment.

On R3

We can also specify a particular instance of ISIS and allocate that instance to a particular interface.

```
R3(config)#router isis TST
R3(config-router)#NET 49.0034.3333.0000.3333.00

R3(config)#int f0/0
R3(config-if)#ip router isis TST

R3(config)#int Lo0
R3(config-if)#ip router isis TST

R3(config)#Int F0/1
R3(config-if)#ip router isis TST
```

On R2

```
R2(config)#router isis
R2(config-router)#NET 49.0012.2222.0000.2222.00

R2(config-router)#Int f0/0
R2(config-if)#ip router isis

R2(config)#int f0/1
R2(config-if)#ip router isis

R2(config)#router isis
R2(config-router)#passive-interface lo0
```

On R1

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config)#router isis
R1(config-router)#NET 49.0012.1111.0000.1111.00

R1(config)#Int f0/0
R1(config-if)#ip router isis

R1(config)#router isis
R1(config-router)#passive-interface Lo0
```

At least one interface must be participating in ISIS and must be configured as using ip to utilize the "passive" keyword on another interface. This states that this interface is not to send ANY ISIS updates out of the interface specified and allow it in the ISIS process at the same time.

Verify the configuration:

On R1

```
R1#Show isis neighbors

System Id      Type Interface   IP Address      State Holdtime Circuit Id
R2             L1   Fa0/0       10.1.12.2       UP    25       R1.01
R2             L2   Fa0/0       10.1.12.2       UP    26       R1.01
```

By default, the Routers automatically build L1/L2 (Level-1/Level-2) adjacencies with each other. The Circuit ID represents the Designated Intermediate System (DIS) router on the segment which functions like the DR (Designated Router) in OSPF.

On R2

```
R2#Show isis neighbors

System Id      Type Interface   IP Address      State Holdtime Circuit Id
R1             L1   Fa0/0       10.1.12.1       UP    9        R1.01
R1             L2   Fa0/0       10.1.12.1       UP    9        R1.01
R3             L2   Fa0/1       10.1.23.3       UP    8        R3.03
```

Unlike the L1/L2 adjacencies between routes in the same area, R2 and R3 DO NOT build a Level 1 adjacency due to the fact that they are NOT in the same area. Only a Level 2 adjacency will remain. It is possible to function as a DIS on one segment and not on a different one. The Circuit ID will change if this is the case. The links will also increment in the order they have been added. If F0/1 had been the 3rd interface added, the circuit would reflect this as well with an ID of R3.04 instead of R3.03.

On R4

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4#Show clns neighbors

System Id      Interface   SNPA            State  Holdtime  Type Protocol
R3             Fa0/0       0011.92a0.3a10  Up     27        L1L2 IS-IS
```

The Sub-network point of attachment (SNPA) reflects the encapsulation of the interface. In this case, the MAC-address of the F0/0 interface of R3. If this were a frame relay connection, it would reflect the "DLCI". If it were HDLC, it would be "HDLC". Also note the Type and Protocol field as it reflects both a Level-1 and Level-2 adjacencies.

On R3

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

  Hardware is Gt96k FE, address is 0011.92a0.3a10 (bia 0011.92a0.3a10)
```

Task 2

Remove all of the unneccesary peering information to optimize ISIS neighbor relationships without removing any part of the configuration.

Let's check the adjacencies established on R1 before the optimization:

On R1

```
R1#Show isis neighbors

System Id      Type Interface   IP Address      State Holdtime Circuit Id
R2             L1   Fa0/0       10.1.12.2       UP    26       R1.01
R2             L2   Fa0/0       10.1.12.2       UP    25       R1.01
```

Let's configure R1 to ONLY establish a L1 adjacency with R2 in area 49.0012.

```
R1(config)#router isis
R1(config-router)#is-type level-1
```

R1 will only build level-1 adjacency with any ISIS router that is connected to it in the same area; the above command configured in the router configuration mode affects ALL interfaces on that router.

This also prevents R2 from creating a Level-2 adjacency with R1. However, R2 should also be optimized for consistency and clarity, unless stated to use the least number of commands possible.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2

NOTE: In the previous task the same command revealed two adjacencies between R1 and R2, but since R1 can ONLY establish a L1 adjacency, R2 dropped the L2 adjacency with R1:

```
R2#Show isis neighbors

System Id      Type Interface   IP Address      State Holdtime Circuit Id
R1             L1   Fa0/0       10.1.12.1       UP    7        R1.01
R3             L2   Fa0/1       10.1.23.3       UP    9        R3.03
```

Let's configure R2 so that it ONLY establishes a L1 adjacency through its F0/0 and a L2 adjacency through its F0/1 interface:

```
R2(config)#int fa0/0
R2(config-if)#isis circuit-type level-1

R2(config)#int fa0/1
R2(config-if)#isis circuit-type level-2-only
```

The ISIS adjacency types can also be modified on an interface-by-interface basis. The command would be "isis circuit-type" to override the default Level 1 and level 2 state (L1/L2).

On R3

```
R3(config)#int f0/0
R3(config-if)#isis circuit-type level-1

R3(config)#int f0/1
R3(config-if)#isis circuit-type level-2-only
```

Let's configure R4 to establish a level-1 ONLY adjacency:

On R4

```
R4(config)#router isis
R4(config-router)#is-type level-1
```

To verify the configuration:

On R4

```
R4#Show isis neighbors
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
System Id      Type Interface   IP Address      State Holdtime Circuit Id
R3             L1   Fa0/0       10.1.34.3       UP    27       R4.01
```

Let's check R4's routing table:

On R4

```
R4#Show ip route isis

i L1 3.0.0.0/8 [115/20] via 10.1.34.3, FastEthernet0/0
i*L1 0.0.0.0/0 [115/10] via 10.1.34.3, FastEthernet0/0
```

The candidate default route behavior is similar to that of an OSPF stub area. The R4 router only etablishes a Level 1 adjacency with R3. R3 injects the default route to not allow R4 to see the Level-2 database and routing information. The default cost of ANY link in ISIS is "10".

To test the configuration:

On R4

```
R4#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

On R2

```
R2#Show ip route isis

i L1 1.0.0.0/8 [115/10] via 10.1.12.1, FastEthernet0/0
i L2 3.0.0.0/8 [115/20] via 10.1.23.3, FastEthernet0/1
i L2 4.0.0.0/8 [115/30] via 10.1.23.3, FastEthernet0/1
     10.0.0.0/24 is subnetted, 3 subnets
i L2    10.1.34.0 [115/20] via 10.1.23.3, FastEthernet0/1
```

Task 3

Change the metrics to generate and accept the more granular type of metric for ISIS. To remain backward compatible, the older type of metric should still be received on R3 and R4. Only the type "135" TLV should be accepted and sent on R1 and R2.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3

```
R3#Show clns protocol TST | B Generate

  Generate narrow metrics: level-1-2
  Accept narrow metrics:   level-1-2
  Generate wide metrics:   none
  Accept wide metrics:     none
```

ISIS routers will automatically Generate and Accept Narrow metrics. Narrow metrics are represented by 6 bits and can only reach a maximum value of "63".

```
R3(config)#router isis TST
R3(config-router)#metric-style wide transition

R3#Show clns protocol TST | B Generate

  Generate narrow metrics: none
  Accept narrow metrics:   level-1-2
  Generate wide metrics:   level-1-2
  Accept wide metrics:     level-1-2
```

Another type of metric-style reference would be called TLV-135, which would be the "metric-style wide" command. This type of metric allows for a 32 bit metric to be used. It was invented for the purpose of MPLS Traffic-Engineering but can also be used to provide granular metrics without running TE. As you can see, the "transition" mechanism allows the router to still process narrow metrics but NOT generate them.

Other TLV types include but certainly not limited to:

"1"- Areas

"2"- ISIS neighbors

"128"- Connected Prefixes = 4 byte metric, 4 byte prefix, 4 byte mask

"130"- External Prefixes = 4 byte metric, 4 byte prefix, 4 byte mask

On R4

Let's check the routing table of R4:

```
R4#Show ip route isis
R4#
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Let's change the metric-style and try again:

On R4

```
R4(config)#router isis
R4(config-router)#metric-style wide transition
```

The keyword "transition" can be used while some routers participate in Narrow and updated routers get deployed with wide metrics.

To verify the configuration:

On R4

```
R4#Show ip route isis

i L1 3.0.0.0/8 [115/20] via 10.1.34.3, FastEthernet0/0
i*L1 0.0.0.0/0 [115/10] via 10.1.34.3, FastEthernet0/0
```

On R1

```
R1(config)#router isis
R1(config-router)#metric-style wide
```

On R2

```
R2(config)#router isis
R2(config-router)#metric-style wide
```

To verify the configuration:

On R1

```
R1#Show ip route isis

i L1 2.0.0.0/8 [115/10] via 10.1.12.2, FastEthernet0/0
i*L1 0.0.0.0/0 [115/10] via 10.1.12.2, FastEthernet0/0
```

On R2

```
R2#Show ip route isis

i L1 1.0.0.0/8 [115/10] via 10.1.12.1, FastEthernet0/0
i L2 3.0.0.0/8 [115/20] via 10.1.23.3, FastEthernet0/1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
i L2 4.0.0.0/8 [115/30] via 10.1.23.3, FastEthernet0/1
     10.0.0.0/24 is subnetted, 3 subnets
i L2    10.1.34.0 [115/20] via 10.1.23.3, FastEthernet0/1
```

On R4

```
R4#Show clns prot | B Generate

  Generate narrow metrics: none
  Accept narrow metrics:   level-1-2
  Generate wide metrics:   level-1-2
  Accept wide metrics:     level-1-2
```

On R1

```
R1#Show clns prot | B Generate

  Generate narrow metrics: none
  Accept narrow metrics:   none
  Generate wide metrics:   level-1-2
  Accept wide metrics:     level-1-2
```

On R2

```
R2#Show clns prot | B Generate

  Generate narrow metrics: none
  Accept narrow metrics:   none
  Generate wide metrics:   level-1-2
  Accept wide metrics:     level-1-2
```

Task 4

Configure the hello timer from R1 to R2 to be one second and the holdtime to be four times the hello timer. The LSP database should be able to update every 10 minutes while holding onto the LSP for 20 minutes when it has not been updated.

To verify the current timer value:

On R1

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1#Debug isis adj
IS-IS Adjacency related packets debugging is on
R1#
*Mar  1 00:07:43.563: ISIS-Adj: Rec L1 IIH from 0030.949f.1ec1 (FastEthernet0/0), cir type L1, cir id
2222.0000.2222.01, length 1497
R1#
*Mar  1 00:07:46.235: ISIS-Adj: Rec L1 IIH from 0030.949f.1ec1 (FastEthernet0/0), cir type L1, cir id
2222.0000.2222.01, length 1497
R1#
*Mar  1 00:07:48.471: ISIS-Adj: Sending L1 LAN IIH on FastEthernet0/0, length 1497
*Mar  1 00:07:49.371: ISIS-Adj: Rec L1 IIH from 0030.949f.1ec1 (FastEthernet0/0), cir type L1, cir id
2222.0000.2222.01, length 1497
R1#
*Mar  1 00:07:52.335: ISIS-Adj: Rec L1 IIH from 0030.949f.1ec1 (FastEthernet0/0), cir type L1, cir id
2222.0000.2222.01, length 1497
R1#
*Mar  1 00:07:54.875: ISIS-Adj: Rec L1 IIH from 0030.949f.1ec1 (FastEthernet0/0), cir type L1, cir id
2222.0000.2222.01, length 1497
R1#
*Mar  1 00:07:56.111: ISIS-Adj: Sending L1 LAN IIH on FastEthernet0/0, length 1497

R1#u all
```

In the output above, R1 is sending hello's (approximately) every 10 seconds. R1 is receiving a hello every 3 seconds by R2 due to the fact that R2 is the Designated Intermediate System in the segment. (The DIS is the equivalent to a DR in OSPF). The DIS sends hellos on an ethernet segment every 3 seconds to ensure a neighbor will see the current DIS that has been elected. The DIS router will be reconfigured in a later lab.

Configure the Hello and Holdown timer:

On R1

```
R1(config)#int f0/0
R1(config-if)#isis hello-interval 1
R1(config-if)#isis hello-multiplier 4
```

The hello interval is set at "1" second and the dead-interval is at "4" seconds.

To verify the configuration:

On R1

```
R1#Debug isis adj-packets
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
IS-IS Adjacency related packets debugging is on

*Mar  1 00:15:25.947: ISIS-Adj: Sending L1 LAN IIH on FastEthernet0/0, length 1497
*Mar  1 00:15:26.787: ISIS-Adj: Sending L1 LAN IIH on FastEthernet0/0, length 1497

*Mar  1 00:15:27.187: ISIS-Adj: Rec L1 IIH from 0030.949f.1ec1 (FastEthernet0/0), cir type L1, cir id
2222.0000.2222.01, length 1497
*Mar  1 00:15:27.539: ISIS-Adj: Sending L1 LAN IIH on FastEthernet0/0, length 1497

*Mar  1 00:15:28.463: ISIS-Adj: Sending L1 LAN IIH on FastEthernet0/0, length 1497
*Mar  1 00:15:29.243: ISIS-Adj: Sending L1 LAN IIH on FastEthernet0/0, length 1497
```

Configure LSP Timers:

On ALL Routers:

```
Rx(config)#router isis
Rx(config-router)#max-lsp-lifetime 2400       (= 40 minutes)
Rx(config-router)#lsp-refresh-interval 600    (= 10 minutes)
```

Both timers are measured in seconds. By default, the lifetime of an LSP is 1200 seconds or 20 minutes, this command sets the maximum number of seconds that an LSP can stay in the LSDB without being refreshed, if exceeded, the LSP is dropped from the LSDB.

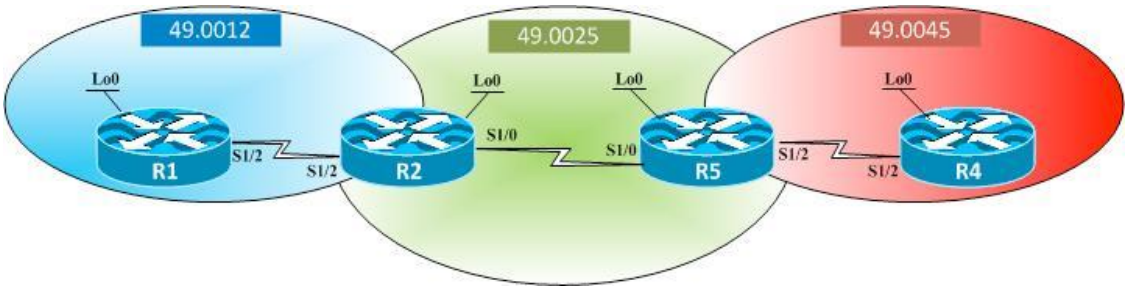
By default, the "LSP-Refresh-Interval" is set to 900 seconds or 15 minutes. This command sets the interval in which the LSPs are refreshed.

Task 5

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

LAB 2- Multi-Instance



Router	Interface	IPv4 address
R1	Lo0 S1/2	1.1.1.1/32 10.1.12.1/24
R2	Lo0 S1/2 S1/0	2.2.2.2 /32 10.1.12.2/24 10.1.25.2/ 24
R4	Lo0 S1/2	4.4.4.4/ 32 10.1.45.4/ 24
R5	Lo0 S1/0 S1/2	5.5.5.5/ 32 10.1.25.5/ 24 10.1.45.5/ 24

Lab Setup:

- The Connection between R1 and R2 should be in VLAN 12
- The Connection between R4 and R5 uses the default encapsulation

Task 1

Configure ISIS level-2 ipv4 routing in area 49.0012 and 49.0045. Use an instance name of "TST" for ISIS. The loopback interface on R1 and R4 should also be under the instance

On R1

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config)#router isis TST
R1(config-router)#net 49.0001.1111.0000.1111.00
R1(config-router)#is-type level-2-only
R1(config-router)#metric-style wide

R1(config-router)#int lo0
R1(config-if)#ip router isis TST

R1(config-if)#int S1/2
R1(config-subif)#ip router isis TST
```

On R2

```
R2(config)#router isis TST
R2(config-router)#net 49.0001.2222.0000.2222.00
R2(config-router)#is-type level-2-only
R2(config-router)#metric-style wide

R2(config-if)#int S1/2
R2(config-subif)#ip router isis TST
```

On R5

```
R1(config)#router isis TST
R1(config-router)#net 49.0045.5555.0000.5555.00
R1(config-router)#is-type level-2-only
R1(config-router)#metric-style wide

R1(config-if)#int s1/2
R1(config-subif)#ip router isis TST
```

On R4

```
R2(config)#router isis TST
R2(config-router)#net 49.0045.4444.0000.4444.00
R2(config-router)#is-type level-2-only
R2(config-router)#metric-style wide

R2(config-router)#int lo0
R2(config-if)#ip router isis TST

R2(config-if)#int S1/2
R2(config-subif)#ip router isis TST
```

Verifying Connectivity:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2

```
R2#Sh isis neighbors

System Id      Type Interface   IP Address      State Holdtime Circuit Id
R1             L2   Se1/2       10.1.12.1       UP    28       00

R2#Sh ip route isis | i L2

       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
i L2      1.1.1.1 [115/20] via 10.1.12.1, 00:01:54, Serial1/2

R2#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/23/64 ms
```

On R5

```
R5#Sh isis neighbors

System Id      Type Interface   IP Address      State Holdtime Circuit Id
R4             L2   Se1/2       10.1.45.4       UP    23       00

R5#Sh ip route isis | i L2

       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
i L2      4.4.4.4 [115/20] via 10.1.45.4, 00:01:45, Serial1/2

R5#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/16/44 ms
```

Task 2

Configure area 49.0025 in a way that does not allow communication between R1 and R4 without filtering of ANY kind.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2

```
R2(config)#router isis TST2
R2(config-router)#net 49.0025.2222.2525.2222.00
R2(config-router)#metric-style wide
R2(config-router)#is-type level-2-only

R2(config-router)#int lo0
R2(config-if)#ip router isis TST2

R2(config-if)#int S1/0
R2(config-if)#ip router isis TST2
```

ISIS Multi-Instance is the way to accomplish this task. Each ISIS Instance owns it's own resources and filtering tasks.

On R5

```
R5(config)#router isis TST2
R5(config-router)#net 49.0025.5555.2525.5555.00
R5(config-router)#is-type level-2-only
R5(config-router)#metric-style wide

R5(config-router)#int lo0
R5(config-if)#ip router isis TST2

R5(config-if)#int S1/0
R5(config-subif)#ip router isis TST2
```

The NET addresses MUST be unique when building another ISIS Instance. When building multiple NET addresses under one instance, the System ID portion would be the same.

Verifying Connectivity:

On R5

```
R5#Sh isis neighbors

Tag TST:
System Id      Type Interface   IP Address      State Holdtime Circuit Id
R4               L2   Se1/2       10.1.45.4       UP    22         00

Tag TST2:
System Id      Type Interface   IP Address      State Holdtime Circuit Id
R2             L2   Se1/0       10.1.25.2       UP    22         00
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R5#Sh ip route isis | i L2

       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
i L2     2.2.2.2 [115/20] via 10.1.25.2, 00:02:41, Serial1/0
i L2     4.4.4.4 [115/20] via 10.1.45.4, 00:12:04, Serial1/2
```

The routes on R1 are not visible as the lab Task required. It is the same result as if we were running 2 completely different routing protocols.

On R2

```
R2#Sh isis neighbors

Tag TST:
System Id      Type Interface   IP Address      State Holdtime Circuit Id
R1               L2   Se1/2        10.1.12.1       UP    25         00

Tag TST2:
System Id      Type Interface   IP Address      State Holdtime Circuit Id
R5               L2   Se1/0        10.1.25.5       UP    22         00

R2#Sh ip route isis | i L2

       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
i L2     1.1.1.1 [115/20] via 10.1.12.1, 00:16:28, Serial1/2
i L2     5.5.5.5 [115/20] via 10.1.25.5, 00:05:43, Serial1/0
```

The routes on R4 are not visible as the lab Task required.

Task 3

Configure the routers such that R1 and R4 have reachability to the loopback interfaces ONLY. No static routes.

On R2

```
R2(config)#access-list 1 permit 1.1.1.1

R2(config)#route-map TST-to-TST2 permit 10
R2(config-route-map)#match ip address 1

R2(config-route-map)#router isis TST2
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config-router)#redistribute isis TST ip route-map TST-to-TST2
R2(config-router)#end
```

Redistribution is required to have one routing instance speak to another routing instance on the same router. We distribute the route that is in the ISIS "TST" instance to redistributed into the TST2 instance. An access-list is used to single out the desired prefix.

On R5

```
R5(config)#access-list 1 permit 1.1.1.1

R5(config)#route-map TST2-to-TST permit 10
R5(config-route-map)#match ip address 1

R5(config-route-map)#router isis TST
R5(config-router)#redistribute isis TST2 ip route-map TST2-to-TST
R5(config-router)#end
```

To allow R4 to see the dired prefix we filtered out in the prior configuration, redistribution is equired from the TST2 instance to the TST instance between R5 and R4.

On R5

```
R5(config)#access-list 4 permit 4.4.4.4

R5(config)#route-map TST-to-TST2 permit 10
R5(config-route-map)#match ip address 4

R5(config-route-map)#router isis TST2
R5(config-router)#redistribute isis TST ip route-map TST-to-TST2
R5(config-router)#end
```

The same process must be repeated for the prefixes originating on R4.

On R2

```
R2(config)#access-list 4 permit 4.4.4.4

R2(config)#route-map TST2-to-TST permit 10
R2(config-route-map)#match ip address 4

R2(config-route-map)#router isis TST
R2(config-router)#redistribute isis TST2 ip route-map TST2-to-TST
R2(config-router)#end
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Verifying Connectivity:

On R2

```
R2#Sh ip route isis | i L2

       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
i L2      1.1.1.1 [115/20] via 10.1.12.1, 00:43:06, Serial1/2
i L2      4.4.4.4 [115/10] via 10.1.25.5, 00:02:20, Serial1/0
i L2      5.5.5.5 [115/20] via 10.1.25.5, 00:32:21, Serial1/0
```

On R5

```
R5#Sh ip route isis | i L2

       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
i L2      1.1.1.1 [115/10] via 10.1.25.2, 00:01:38, Serial1/0
i L2      2.2.2.2 [115/20] via 10.1.25.2, 00:33:49, Serial1/0
i L2      4.4.4.4 [115/20] via 10.1.45.4, 00:43:12, Serial1/2
```

On R4

```
R4#Sh ip route isis | i L2

      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
i L2      1.1.1.1 [115/10] via 10.1.45.5, 00:00:21, Serial1/2

R4#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

The pings that are generated from the physical interface should NOT work!

```
R4#Ping 1.1.1.1 so lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 4.4.4.4
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 12/18/28 ms
```

The ping generated from the Loopback Interface should work.

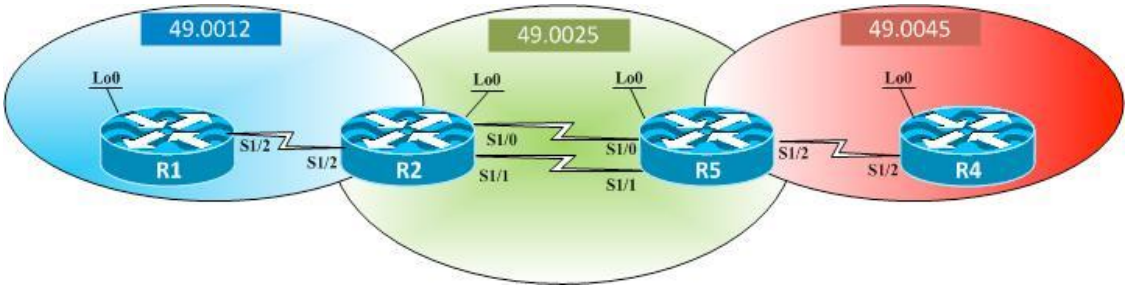
©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 4

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

LAB 3- Multi-topology



Router	Interface	IPv4 address	IPv6 Address
R1	Lo0 S1/2	1.1.1.1/32 10.1.12.1/ 24	1:1:1::1/128 10:1:12::1/64
R2	Lo0 S1/2 S1/0 S1/1	2.2.2.2 /32 10.1.12.2/ 24 10.1.25.2/ 24 10.1.52.2/ 24	2:2:2::2/128 10:1:12::2/64 10:1:25::2/64 10:1:52::2/64
R4	Lo0 S1/2	4.4.4.4/ 32 10.1.45.4/ 24	4:4:4::4/128 10:1:45::5/64
R5	Lo0 S1/0 S1/1 S1/2	5.5.5.5/ 32 10.1.25.5/ 24 10.1.52.5/ 24 10.1.45.5/ 24	5:5:5::5/128 10:1:25::5/64 10:1:52::5/64 10:1:45::5/64

Lab Setup:

- All Serial connections between the routers should use the default encapsulation

Task 1

Configure ISIS to support routing connectivity between ALL routers for ALL IPv4 prefixes. Only required adjacencies should be viewed configuring the fewest commands possible.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1

```
R1(config)#router isis
R1(config-router)#net 49.1245.1111.0000.1111.00
R1(config-router)#metric-style wide

R1(config-router)#int lo0
R1(config-if)#ip router isis

R1(config-if)#int s1/2
R1(config-subif)#ip router isis
```

On R2

```
R2(config)#router isis
R2(config-router)#net 49.1245.2222.0000.2222.00
R2(config-router)#metric-style wide
R2(config-router)#is-type level-2-only

R2(config-router)#int lo0
R2(config-if)#ip router isis

R2(config-if)#int s1/0
R2(config-subif)#ip router isis

R2(config-if)#int s1/1
R2(config-subif)#ip router isis

R2(config-if)#int s1/2
R2(config-subif)#ip router isis
```

The "is-type" command is only required on R2 and R5. The adjacency from on area to another only maintains the L2 adjacency.

On R5

```
R5(config)#router isis
R5(config-router)#net 49.1245.5555.0000.5555.00
R5(config-router)#metric-style wide
R5(config-router)#is-type level-2-only

R5(config-router)#int lo0
R5(config-if)#ip router isis

R5(config-if)#int s1/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R5(config-subif)#ip router isis

R5(config-if)#int s1/1
R5(config-subif)#ip router isis

R5(config-if)#int s1/2
R5(config-subif)#ip router isis
```

On R4

```
R4(config)#router isis
R4(config-router)#net 49.1245.4444.0000.4444.00
R4(config-router)#metric-style wide

R4(config-router)#int lo0
R4(config-if)#ip router isis

R4(config-if)#int s1/2
R4(config-subif)#ip router isis
```

Verifying Connectivity:

On R1

```
R1#Sh ip route isis | i L2

       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
i L2      2.2.2.2 [115/20] via 10.1.12.2, 00:01:46, Serial1/2
i L2      4.4.4.4 [115/40] via 10.1.12.2, 00:01:36, Serial1/2
i L2      5.5.5.5 [115/30] via 10.1.12.2, 00:01:36, Serial1/2
i L2      10.1.25.0/24 [115/20] via 10.1.12.2, 00:01:46, Serial1/2
i L2      10.1.45.0/24 [115/30] via 10.1.12.2, 00:01:36, Serial1/2
i L2      10.1.52.0/24 [115/20] via 10.1.12.2, 00:01:46, Serial1/2
```

On R2

```
R2#Sh ip route | b Gateway
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 1 subnets
i L2     1.1.1.1 [115/20] via 10.1.12.1, 00:10:34, Serial1/2
      2.0.0.0/32 is subnetted, 1 subnets
C        2.2.2.2 is directly connected, Loopback0
      4.0.0.0/32 is subnetted, 1 subnets
i L2     4.4.4.4 [115/30] via 10.1.52.5, 00:10:27, Serial1/1
                 [115/30] via 10.1.25.5, 00:10:27, Serial1/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
      5.0.0.0/32 is subnetted, 1 subnets
i L2     5.5.5.5 [115/20] via 10.1.52.5, 00:10:27, Serial1/1
                 [115/20] via 10.1.25.5, 00:10:27, Serial1/0
      10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
C        10.1.12.0/24 is directly connected, Serial1/2
L        10.1.12.2/32 is directly connected, Serial1/2
C        10.1.25.0/24 is directly connected, Serial1/0
L        10.1.25.2/32 is directly connected, Serial1/0
i L2     10.1.45.0/24 [115/20] via 10.1.52.5, 00:10:27, Serial1/1
                      [115/20] via 10.1.25.5, 00:10:27, Serial1/0
C        10.1.52.0/24 is directly connected, Serial1/1
L        10.1.52.2/32 is directly connected, Serial1/1
```

The IPv4 routes can choose one of 2 paths on the R2 and R5 routers. Cisco Express Forwarding (CEF) would load balance this traffic based on Source and Destination Prefix by default. Across the S1/0 and S1/1 interfaces.

On R5

```
R5#Sh ip route | b Gateway
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 1 subnets
i L2     1.1.1.1 [115/30] via 10.1.52.2, 00:09:26, Serial1/1
                 [115/30] via 10.1.25.2, 00:09:26, Serial1/0
      2.0.0.0/32 is subnetted, 1 subnets
i L2     2.2.2.2 [115/20] via 10.1.52.2, 00:09:26, Serial1/1
                 [115/20] via 10.1.25.2, 00:09:26, Serial1/0
      4.0.0.0/32 is subnetted, 1 subnets
i L2     4.4.4.4 [115/20] via 10.1.45.4, 00:09:26, Serial1/2
      5.0.0.0/32 is subnetted, 1 subnets
C        5.5.5.5 is directly connected, Loopback0
      10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
i L2     10.1.12.0/24 [115/20] via 10.1.52.2, 00:09:26, Serial1/1
                      [115/20] via 10.1.25.2, 00:09:26, Serial1/0
C        10.1.25.0/24 is directly connected, Serial1/0
L        10.1.25.5/32 is directly connected, Serial1/0
C        10.1.45.0/24 is directly connected, Serial1/2
L        10.1.45.5/32 is directly connected, Serial1/2
C        10.1.52.0/24 is directly connected, Serial1/1
L        10.1.52.5/32 is directly connected, Serial1/1
```

On R4

```
R4#Sh ip route | i L2

       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
i L2     1.1.1.1 [115/40] via 10.1.45.5, 00:06:00, Serial1/2
i L2     2.2.2.2 [115/30] via 10.1.45.5, 00:06:00, Serial1/2
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
i L2      5.5.5.5 [115/20] via 10.1.45.5, 00:06:00, Serial1/2
i L2      10.1.12.0/24 [115/30] via 10.1.45.5, 00:06:00, Serial1/2
i L2      10.1.25.0/24 [115/20] via 10.1.45.5, 00:06:00, Serial1/2
i L2      10.1.52.0/24 [115/20] via 10.1.45.5, 00:06:00, Serial1/2
```

Task 2

Configure ISIS to support routing connectivity between ALL routers for ALL IPv6 prefixes.

On R1

```
R1(config)#ipv6 unicast-routing

R1(config)#int lo0
R1(config-if)#ipv6 router isis

R1(config-if)#int s1/2
R1(config-if)#ipv6 router isis
R1(config-if)#end
```

On R2

```
R2(config)#ipv6 uni

R2(config)#int lo0
R2(config-if)#ipv6 router isis

R2(config-if)#int s1/0
R2(config-if)#ipv6 router isis

R2(config-if)#int s1/1
R2(config-if)#ipv6 router isis

R2(config-if)#int s1/2
R2(config-subif)#ipv6 router isis
R2(config-subif)#end
```

On R5

```
R5(config)#ipv6 uni
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R5(config-if)#int lo0
R5(config-if)#ipv6 router isis

R5(config-if)#int s1/0
R5(config-if)#ipv6 router isis

R5(config-if)#int s1/1
R5(config-if)#ipv6 router isis

R5(config-if)#int s1/2
R5(config-subif)#ipv6 router isis

R5(config-subif)#end
```

On R4

```
R4(config)#ipv6 uni

R4(config)#int lo0
R4(config-if)#ipv6 router isis

R4(config-if)#int s1/2
R4(config-if)#ipv6 router isis
R4(config-if)#end
```

Verifying Connectivity:

On R1

```
R1#Sh ipv6 route isis | s I2

       B - BGP, R - RIP, I1 - ISIS L1, I2 - ISIS L2
I2  2:2:2::2/128 [115/20]
     via FE80::2, Serial1/2
I2  4:4:4::4/128 [115/40]
     via FE80::2, Serial1/2
I2  5:5:5::/64 [115/30]
     via FE80::2, Serial1/2
I2  10:1:25::/64 [115/20]
     via FE80::2, Serial1/2
I2  10:1:45::/64 [115/30]
     via FE80::2, Serial1/2
I2  10:1:52::/64 [115/20]
     via FE80::2, Serial1/2
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2

```
R2#Sh ipv6 route | s I2

       B - BGP, R - RIP, I1 - ISIS L1, I2 - ISIS L2
I2  1:1:1::1/128 [115/20]
I2  4:4:4::4/128 [115/30]
     via FE80::5, Serial1/0
     via FE80::5, Serial1/1
I2  5:5:5::/64 [115/20]
     via FE80::5, Serial1/0
     via FE80::5, Serial1/1
I2  10:1:45::/64 [115/20]
     via FE80::5, Serial1/0
     via FE80::5, Serial1/1
```

On R5

```
R5#Sh ipv6 route isis | s I2

       B - BGP, R - RIP, I1 - ISIS L1, I2 - ISIS L2
I2  1:1:1::1/128 [115/30]
     via FE80::2, Serial1/0
     via FE80::2, Serial1/1
I2  2:2:2::2/128 [115/20]
     via FE80::2, Serial1/0
     via FE80::2, Serial1/1
I2  4:4:4::4/128 [115/20]
     via FE80::4, Serial1/2
I2  10:1:12::/64 [115/20]
     via FE80::2, Serial1/0
     via FE80::2, Serial1/1
```

On R4

```
R4#Sh ipv6 route | s I2

       B - BGP, R - RIP, I1 - ISIS L1, I2 - ISIS L2
I2  1:1:1::1/128 [115/40]
I2  2:2:2::2/128 [115/30]
     via FE80::5, Serial1/2
I2  5:5:5::/64 [115/20]
     via FE80::5, Serial1/2
I2  10:1:12::/64 [115/30]
     via FE80::5, Serial1/2
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
I2  10:1:25::/64 [115/20]
     via FE80::57, Serial1/2
I2  10:1:52::/64 [115/20]
     via FE80::5, Serial1/2
```

A similar pattern of routing has occurred. Both routes to the R1 and R4 loopback interfaces are load balaced across the S0/1 and S1/1 interfaces that interconnect R2 and R5.

Task 3

Modify the paths of the IPv4 and IPv6 paths to allow IPv4 treffic to use the s1/0 path between R2 and R5. IPv6 traffic should take the S1/1 interface path. The traffic should be Symetric in nature. Test failover without causing ANY physical failures.

On R2

```
R2#Trace 4.4.4.4

Type escape sequence to abort.
Tracing the route to 4.4.4.4

  1 10.1.25.5 24 msec
    10.1.52.5 4 msec
    10.1.25.5 12 msec
  2 10.1.45.4 8 msec *  16 msec
```

On R5

```
R5#Trace 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 10.1.25.2 8 msec
    10.1.52.2 4 msec
    10.1.25.2 12 msec
  2 10.1.12.1 8 msec *  32 msec
```

On R2

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config)#int s1/0
R2(config-if)# isis metric 5

R2(config-if)#int s1/1
R2(config-if)#isis ipv6 metric 5

R2(config)#router isis
R2(config-router)#address-family ipv6
R2(config-router-af)#multi-topology
```

The Isis Multi-Topoly command must be configured under the ISIS Instance on ALL of the routers that are participating with the IPv4 and IPv6 protocols. When this command is configured, it will allow the user to have independent control over the ISIS metrics from each of the protocols.

On R5

```
R5(config)#int s1/0
R5(config-if)# isis metric 5

R5(config)#int s1/1
R5(config-if)#isis ipv6 metric 5

R5(config)#router isis
R5(config-router)#address-family ipv6
R5(config-router-af)#multi-topology
```

On R1

```
R1(config)#router isis
R1(config-router)#address-family ipv6
R1(config-router-af)#multi-topology
```

On R4

```
R4(config)#router isis
R4(config-router)#address-family ipv6
R4(config-router-af)#multi-topology
```

To Test the Policy:

On R2

```
R2#Sh ip route | i L2

       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
i L2     1.1.1.1 [115/20] via 10.1.12.1, 00:18:42, Serial1/2
i L2     4.4.4.4 [115/25] via 10.1.25.5, 00:02:37, Serial1/0
i L2     5.5.5.5 [115/15] via 10.1.25.5, 00:02:37, Serial1/0
i L2     10.1.45.0/24 [115/15] via 10.1.25.5, 00:02:37, Serial1/0

R2#Trace 4.4.4.4

Type escape sequence to abort.
Tracing the route to 4.4.4.4

  1 10.1.25.5 24 msec 4 msec 8 msec
  2 10.1.45.4 8 msec *  32 msec
```

On R5

```
R5#Sh ipv6 route isis

IPv6 Routing Table - default - 13 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, I1 - ISIS L1, I2 - ISIS L2
       IA - ISIS interarea, IS - ISIS summary, D - EIGRP, EX - EIGRP external
       ND - Neighbor Discovery
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
I2  1:1:1::1/128 [115/25]
     via FE80::2, Serial1/1
I2  2:2:2::2/128 [115/15]
     via FE80::2, Serial1/1
I2  4:4:4::4/128 [115/20]
     via FE80::4, Serial1/2
I2  10:1:12::/64 [115/15]
     via FE80::2, Serial1/1

R5#Trace 1:1:1::1

Type escape sequence to abort.
Tracing the route to 1:1:1::1

  1 10:1:52::2 12 msec 4 msec 8 msec
  2 10:1:12::1 12 msec 20 msec 8 msec
```

On R4

```
R4#Trace 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
  1 10.1.45.5 52 msec 64 msec 24 msec
  2 10.1.25.2 12 msec 8 msec 8 msec
  3 10.1.12.1 76 msec *  24 msec

R4#Trace 1:1:1::1

Type escape sequence to abort.
Tracing the route to 1:1:1::1

  1 10:1:45::5 16 msec 4 msec 4 msec
  2 10:1:52::2 12 msec 4 msec 12 msec
  3 10:1:12::1 8 msec 16 msec 8 msec
```

On R1

```
R1#Trace 4.4.4.4

Type escape sequence to abort.
Tracing the route to 4.4.4.4

  1 10.1.12.2 12 msec 4 msec 12 msec
  2 10.1.25.5 8 msec 12 msec 8 msec
  3 10.1.45.4 24 msec *  12 msec

R1#Trace 4:4:4::4

Type escape sequence to abort.
Tracing the route to 4:4:4::4

  1 10:1:12::2 4 msec 8 msec 4 msec
  2 10:1:52::5 8 msec 16 msec 4 msec
 3 10:1:45::4 24 msec 8 msec 16 msec
```

The traceroute does seem to work from each of the endpoints. CEF may show different results from the endpoints but will eventually use both paths.

To Test the Failover:

On R2

```
R2(config)#int s1/0
R2(config-if)#isis protocol shutdown
```

On R1

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1#Trace 4.4.4.4

Type escape sequence to abort.
Tracing the route to 4.4.4.4

  1 10.1.12.2 20 msec 4 msec 8 msec
  2 10.1.52.5 20 msec 8 msec 12 msec
  3 10.1.45.4 12 msec *  32 msec

R1#Trace 4:4:4::4

Type escape sequence to abort.
Tracing the route to 4:4:4::4

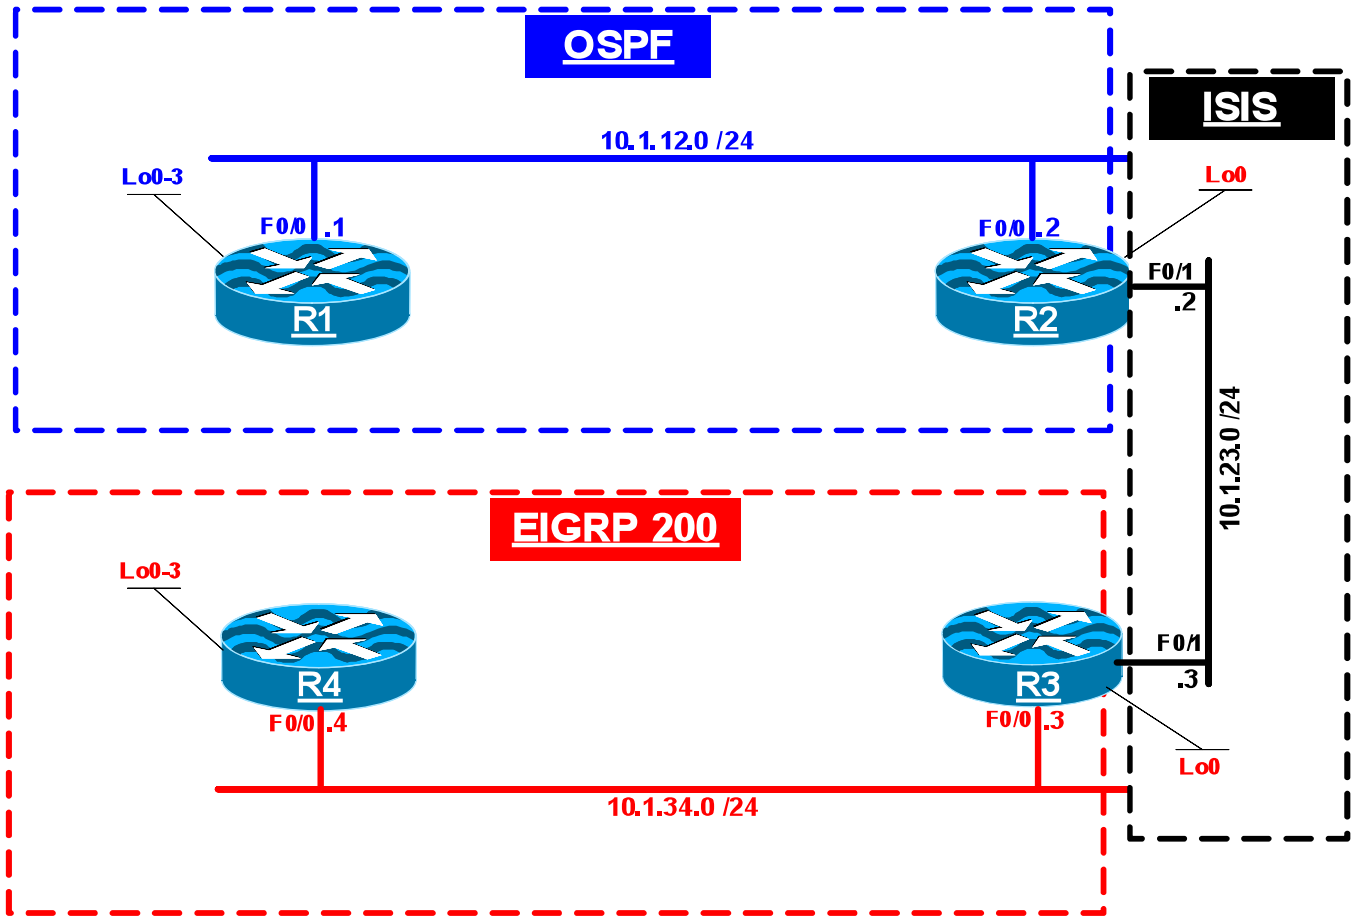
  1 10:1:12::2 12 msec 8 msec 4 msec
  2 10:1:52::5 8 msec 16 msec 4 msec
  3 10:1:45::4 8 msec 4 msec 12 msec
```

Task 4

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 4
Redistribution and Summarization



Lab Setup:

- Configure the routers based on the above topology.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP Addressing:

Router	Interface	IP address
R1	Lo0 Lo1 Lo2 Lo3 F0/0	1.1.0.1 /24 1.1.1.1 /24 1.1.2.1 /24 1.1.3.1 /24 10.1.12.1 /24
R2	Lo0 F0/0 F0/1	2.2.2.2 /8 10.1.12.2 /24 10.1.23.2 /24
R3	Lo0 F0/0 F0/1	3.3.3.3 /8 10.1.34.3 /24 10.1.23.3 /24
R4	Lo0 Lo1 Lo2 Lo3 F0/0	4.4.0.4 /24 4.4.1.4 /24 4.4.2.4 /24 4.4.3.4 /24 10.1.34.4 /24

Task 1

Configure OSPF process 1 on R1 and run all directly connected interfaces in area 0; R2 should be configured to run OSPF area 0 on its connection to R1.
R4 should be configured in EIGRP AS200, this router should advertise its directly connected interfaces in this AS; R3 should be configured to establish an Eigrp adjacency with R4.

On R1

```
R1(config)#router ospf 1
R1(config-router)#netw 1.1.0.0 0.0.255.255 area 0
R1(config-router)#netw 10.1.12.1 0.0.0.0 area 0
```

On R2

```
R2(config)#router ospf 1
R2(config-router)#netw 10.1.12.2 0.0.0.0 area 0
```

On R3

```
R3(config)#router eigrp 200
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-router)#netw 10.1.34.3 0.0.0.0
R3(config-router)#NO auto
```

On R4

```
R4(config)#router eigrp 200
R4(config-router)#netw 10.1.34.4 0.0.0.0
R4(config-router)#netw 4.0.0.0
R4(config-router)#NO auto
```

To verify the configuration:

On R3

```
R3#Sh ip route eigrp | in D

D       4.4.0.0 [90/409600] via 10.1.34.4, 00:00:28, FastEthernet0/0
D       4.4.1.0 [90/409600] via 10.1.34.4, 00:00:28, FastEthernet0/0
D       4.4.2.0 [90/409600] via 10.1.34.4, 00:00:28, FastEthernet0/0
D       4.4.3.0 [90/409600] via 10.1.34.4, 00:00:28, FastEthernet0/0
```

On R2

```
R2#Sh ip route ospf | i O

O       1.1.1.1 [110/11] via 10.1.12.1, 00:00:50, FastEthernet0/0
O       1.1.0.1 [110/11] via 10.1.12.1, 00:00:50, FastEthernet0/0
O       1.1.3.1 [110/11] via 10.1.12.1, 00:00:50, FastEthernet0/0
O       1.1.2.1 [110/11] via 10.1.12.1, 00:00:50, FastEthernet0/0
```

Task 2

Configure ISIS on the interfaces between R2 and R3 and both loopback interfaces. The NET addresses should be 49.0123 for the area. R3 should use a System ID of 3333.0000.3333 and R2 should use 2222.0000.2222. These routers should NOT use the default metric-style to accomplish this task.

On R2

```
R2(config)#router isis
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config-router)#net 49.0123.2222.0000.2222.00
R2(config-router)#metric-style wide
R2(config-router)#is-type level-2-only

R2(config-router)#int f0/1
R2(config-if)#ip router isis

R2(config)#int lo0
R2(config-if)#ip router isis
```

On R3

```
R3(config)#router isis
R3(config-router)#net 49.0123.3333.0000.3333.00
R3(config-router)#is-type level-2-only
R3(config-router)#metric-style wide

R3(config)#int f0/1
R3(config-if)#ip router isis

R3(config)#int lo0
R3(config-if)#ip router isis
```

The routers are performing as backbone routers that are in the same area.

To verify the configuration:

On R2

```
R2#Show isis neighbors

System Id      Type Interface   IP Address      State Holdtime Circuit Id
R3             L2   Fa0/1       10.1.23.3       UP    9        R3.01

R2#Show isis topology

IS-IS paths to level-2 routers
System Id            Metric     Next-Hop        Interface   SNPA
R2                   --
R3                   10         R3              Fa0/1       0011.92a0.3a11

R2#Show isis rib

IPv4 local RIB for IS-IS process
3.0.0.0/8
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
  [115/L2/20] via 10.1.23.3(FastEthernet0/1), from 3.3.3.3, tag 0, LSP[1/9]
10.1.23.0/24
  [115/L2/20] via 10.1.23.3(FastEthernet0/1), from 3.3.3.3, tag 0, LSP[1/9]
```

The isis rib stores the best paths to a destination that are to be compared. The best path is injected into the routing table.

The [115/L2/20]] represents the Administrative distance (115), The type of route (Level-2) and the metric (20).

The tag will be manipulated in another portion of this lab and defaults to "0".

The LSP [1/9] represents action that is performed on the internal process only. The "1" represents the instance assigned for specific group of LSA updates. If there were more neighbors with multiple paths to the destination, this number would change to track the different sources.

The "9" represents the version of the update within that instance. Whenever the LSPDU changes or is refreshed, the number will increment by the number of changes. If a change is by 1 peer, it will increment once per change. If by two, the number should increment twice per change if the path is equal.

Task 3

Configure basic redistribution on R2 and R3 to ensure that R1 and R4 can ping all of the loopback addresses.

On R2

```
R2(config)#router isis
R2(config-router)#redistribute ospf 1
```

The generic redistribution is performed without any filter. All of the Loopback addresses should be redistributed with the default metric of "0"

```
R2(config)#router ospf 1
R2(config-router)#redistribute isis subnets
```

On R3

```
R3(config)#router isis
R3(config-router)#redistribute eigrp 200
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config)#router eigrp 200
R3(config-router)#redistribute isis metric 1 1 1 1 1
```

To verify the configuration:

On R2

```
R2#Show ip route isis

i L2 3.0.0.0/8 [115/20] via 10.1.23.3, FastEthernet0/1
     4.0.0.0/24 is subnetted, 4 subnets
i L2    4.4.0.0 [115/10] via 10.1.23.3, FastEthernet0/1
i L2    4.4.1.0 [115/10] via 10.1.23.3, FastEthernet0/1
i L2    4.4.2.0 [115/10] via 10.1.23.3, FastEthernet0/1
i L2    4.4.3.0 [115/10] via 10.1.23.3, FastEthernet0/1
     10.0.0.0/24 is subnetted, 3 subnets
i L2    10.1.34.0 [115/10] via 10.1.23.3, FastEthernet0/1
```

The output of the above show command reveals the internal Level-2 route and a cost of "10" due to the crossing of the f0/1 interface.

On R3

```
R3#Show isis database detail | Inc 1.1.[0-3]

  Metric: 0          IP 1.1.0.1/32
  Metric: 0          IP 1.1.1.1/32
  Metric: 0          IP 1.1.2.1/32
  Metric: 0          IP 1.1.3.1/32
```

Here we see that the 1.1.0.0 prefixes have a metric of "0" after redistribution has been accomplished.

```
R3#Show ip route isis

     1.0.0.0/32 is subnetted, 4 subnets
i L2    1.1.1.1 [115/10] via 10.1.23.2, FastEthernet0/1
i L2    1.1.0.1 [115/10] via 10.1.23.2, FastEthernet0/1
i L2    1.1.3.1 [115/10] via 10.1.23.2, FastEthernet0/1
i L2    1.1.2.1 [115/10] via 10.1.23.2, FastEthernet0/1
i L2 2.0.0.0/8 [115/20] via 10.1.23.2, FastEthernet0/1
     10.0.0.0/24 is subnetted, 3 subnets
i L2    10.1.12.0 [115/10] via 10.1.23.2, FastEthernet0/1
```

In the RIB of R3 the cost of the routes has incremented to "10" for the same reason as R2.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1

```
R1#Ping 4.4.3.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.3.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

On R4

```
R4#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Task 4

Modify the cost of the redistributed routes from the default value to ensure the cost is 34 for the routes sourced by R4 and a cost of 16777216 for the loopback of R2 in ISIS without configuring a value explicitly on the loopback interface.

On R3

```
R3(config)#access-list 4 permit 4.4.0.0 0.0.255.255

R3(config)#route-map TST4 permit 10
R3(config-route-map)#match ip address 4
R3(config-route-map)#set metric 24

R3(config)#router isis
R3(config-router)#redistribute eigrp 200 route-map TST4

R3(config)#int f0/1
R3(config-if)#isis metric 1
```

The default value of a redistribute route into ISIS is "0". Therfore the value of "24" should be used in the route-map.

On R2

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config)#int lo0
R2(config-if)#isis metric maximum
```

ISIS uses a 24 metric when the metric-style wide command is used. The "maximum" keyword forces the vlaue of the metric to be 2 to the 24th power (16777215). The cost of the link on R3's interface is "10" by default, which = (16777225). Since the value CANNOT be explicity configured to change the value, we must modify the cost on the R3's fa0/1 interface to a value of "1" to change the value to be 16777216 in the R3 routing table.

To verify the configuration:

On R2

```
R2#Show clns int lo0 | S Level-2 Metric:

    Level-2 Metric: 16777215, Priority: 64, Circuit ID: R2.02

R2#Show isis database detail | Inc Metric: 16

  Metric: 16777215   IP 2.0.0.0/8
```

On R3

```
R3#Show ip route isis | Inc 2.0.0.0

i L2 2.0.0.0/8 [115/16777216] via 10.1.23.2, FastEthernet0/1
```

The cost is incremented due to the cost change on the R3 f0/1 link.

```
R3#Show isis database detail | Inc Metric: 24

  Metric: 24         IP 4.4.0.0/24
  Metric: 24         IP 4.4.1.0/24
  Metric: 24         IP 4.4.2.0/24
  Metric: 24         IP 4.4.3.0/24
```

The redistributed metric that is sourced by R3 shows a cost of 24 due to the route-map.

On R2

```
R2#Show ip route isis | Inc 4.4.

i L2    4.4.0.0 [115/34] via 10.1.23.3, FastEthernet0/1
i L2    4.4.1.0 [115/34] via 10.1.23.3, FastEthernet0/1
i L2    4.4.2.0 [115/34] via 10.1.23.3, FastEthernet0/1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
i L2    4.4.3.0 [115/34] via 10.1.23.3, FastEthernet0/1
```

After the routes cross the link with a cost of "10", they show up as "34".

Task 5

Limit the routing table entries in both R3 and R2 by summarizing the loopback interfaces of R1 and R4 respectively. The summarized route for the 4.0.0.0 networks should have a feature that allows matching on other than the prefix itself.

On R3

```
R3(config)#router isis
R3(config-router)#summary-address 4.4.0.0 255.255.252.0 tag 4444
```

On R2

```
R2(config)#router isis
R2(config-router)#summary-address 1.1.0.0 255.255.252.0 level-2
```

To verify the configuration:

On R2

```
R2#Show ip route isis

     1.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
i su    1.1.0.0/22 [115/0] via 0.0.0.0, Null0
i L2 3.0.0.0/8 [115/20] via 10.1.23.3, FastEthernet0/1
     4.0.0.0/22 is subnetted, 1 subnets
i L2    4.4.0.0 [115/34] via 10.1.23.3, FastEthernet0/1
```

The summary address is routed to a null interface as it should not be used as the best path while a more specific prefix to the destination exist. In this case, the route learned through OSPF.

```
R2#Show isis rib | S 4.4.0.0

4.4.0.0/22
  [115/L2/34] via 10.1.23.3(FastEthernet0/1), from 3.3.3.3, tag 4444, LSP[1/20]
```

The tag that was added to the summarized route on R3 cannot be viewed in the default routing table. It must be viewed in the isis rib. This tag can be used to match instead of a prefix.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3

```
R3#Show ip route isis

     1.0.0.0/22 is subnetted, 1 subnets
i L2    1.1.0.0 [115/1] via 10.1.23.2, FastEthernet0/1
i L2 2.0.0.0/8 [115/16777216] via 10.1.23.2, FastEthernet0/1
     4.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
i su    4.4.0.0/22 [115/24] via 0.0.0.0, Null0
     10.0.0.0/24 is subnetted, 3 subnets
i L2    10.1.12.0 [115/1] via 10.1.23.2, FastEthernet0/1
```

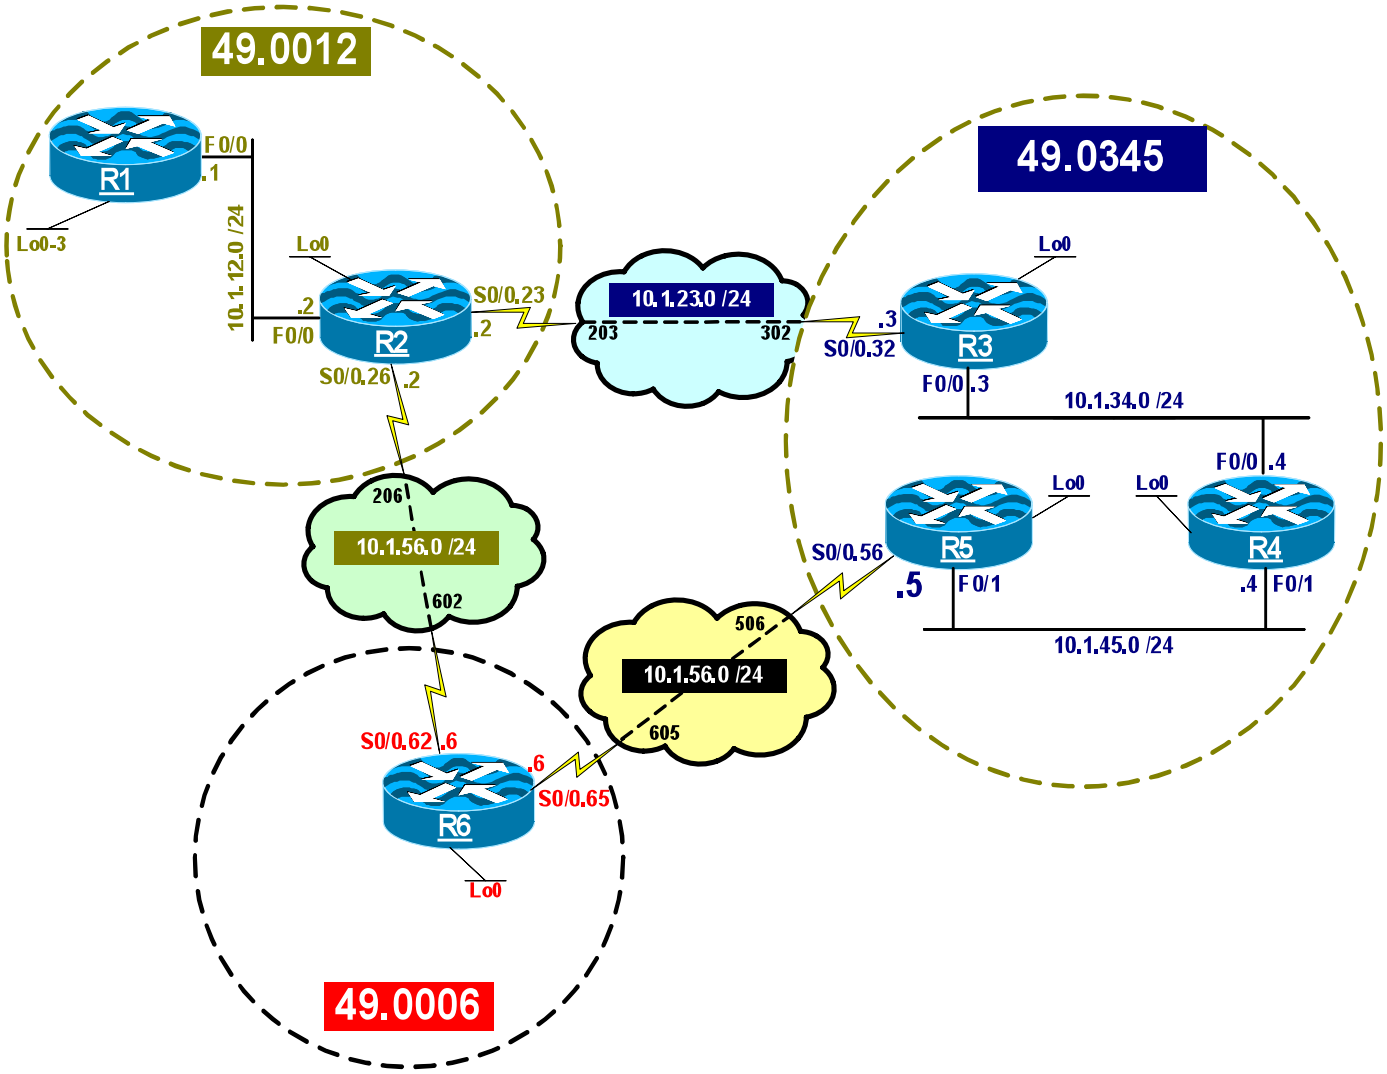
The "su" keyword in the rib shows that this is a summarized prefix.

Task 6

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 5
ISIS Efficiency Mechanisms



Lab Setup:

- Connection between R1 F0/0 and R2 F0/0 are to be configured in VLAN 12
- Connection between R3 F0/0 and R4 F0/0 are to be configured in VLAN 34
- Connection between R4 F0/1 and R5 F0/0 are to be configured in VLAN 45
- All serial connections should be configured in a frame-relay point-to-point mannor.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP Addressing:

Router	Interface	IP address
R1	Lo0 Lo1 Lo2 Lo3 F0/0	1.1.0.1 /24 1.1.1.1 /24 1.1.2.1 /24 1.1.3.1 /24 10.1.12.1/24
R2	Lo0 S0/0.23 S0/0.26 F0/0	2.2.2.2 /8 10.1.23.2 /24 10.1.26.2/24 10.1.12.2 /24
R3	Lo0 S0/0.32 F0/0	3.3.3.3 /8 10.1.23.3 /24 10.1.34.3 /24
R4	Lo0 F0/0 F0/1	4.4.4.4/8 10.1.34.4 /24 10.1.45.4/24
R5	Lo0 S0/0.56 F0/0	5.5.5.5/8 10.1.56.5/24 10.1.45.5/24
R6	Lo0 S0/0.62 S0/0.65	6.6.6.6/8 10.1.26.6/24 10.1.56.6/24

Task 1

Configure ISIS on all routers. R1 and R2 are in area 49.0012, R3, R4 and R5 are in area 49.0345, and R6 is in area 49.0006. The system-ids should all be in the following format: (<routernumber.0000.routernumber), for R5 it would be 5555.0000.5555. ALL routers should be configured with optimal ISIS level adjacencies. These routers should advertise all their directly connected interface/s in their identified Domain/Area.

On R1

```
R1(config)#router isis
R1(config-router)#net 49.0012.1111.0000.1111.00
R1(config-router)#metric-style wide
R1(config-router)#is-type level-1

R1(config-router)#int f0/0
R1(config-if)#ip router isis
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config-if)#int lo0
R1(config-if)#ip router isis

R1(config-if)#int lo1
R1(config-if)#ip router isis

R1(config-if)#int lo2
R1(config-if)#ip router isis

R1(config-if)#int lo3
R1(config-if)#ip router isis
```

On R2

```
R2(config)#router isis
R2(config-router)#net 49.0012.2222.0000.2222.00
R2(config-router)#metric-style wide

R2(config)#int fa0/0
R2(config-if)#ip router isis
R2(config-if)#isis circuit-type level-1

R2(config)#int s0/0.23
R2(config-subif#ip router isis
R2(config-subif)#isis circuit-type level-2-only

R2(config)#int s0/0.26
R2(config-subif#ip router isis
R2(config-subif)#isis circuit-type level-2-only

R1(config-subif)#int lo0
R1(config-if)#ip router isis
```

On R3

```
R3(config)#router isis
R3(config-router)#net 49.0345.3333.0000.3333.00
R3(config-router)#metric-style wide

R3(config)#int f0/0
R3(config-if)#ip router isis
R3(config-if)#isis circuit-type level-1

R3(config)#int lo0
R3(config-if)#ip router isis
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-if)#int s0/0.32
R3(config-subif)#ip router isis
R3(config-subif)#isis circuit-type level-2-only
```

On R4

```
R4(config)#router isis
R4(config-router)#net 49.0345.4444.0000.4444.00
R4(config-router)#metric-style wide
R4(config-router)#is-type level-1

R4(config)#int f0/0
R4(config-if)#ip router isis

R4(config)#int f0/1
R4(config-if)#ip router isis

R4(config)#int lo0
R4(config-if)#ip router isis
```

On R5

```
R5(config)#router isis
R5(config-router)#net 49.0345.5555.0000.5555.00
R5(config-router)#metric-style wide

R5(config)#int f0/0
R5(config-if)#ip router isis
R5(config-if)#isis circuit-type level-1

R5(config)#int lo0
R5(config-if)#ip router isis

R5(config)#int s0/0.56
R5(config-subif)#ip router isis
R5(config-subif)#isis circuit-type level-2-only
```

On R6

```
R6(config)#router isis
R6(config-router)#net 49.0006.6666.0000.6666.00
R6(config-router)#metric-style wide
R6(config-router)#is-type level-2-only

R6(config)#int s0/0.65
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R6(config-subif#ip router isis

R6(config)#int s0/0.62
R6(config-subif#ip router isis

R6(config)#int lo0
R6(config-if)#ip router isis
```

To verify the configuration:

On R1

```
R1#Show ip route isis

i L1 2.0.0.0/8 [115/20] via 10.1.12.2, FastEthernet0/0
i*L1 0.0.0.0/0 [115/10] via 10.1.12.2, FastEthernet0/0
```

On R2

```
R2#Show ip route isis

     1.0.0.0/24 is subnetted, 4 subnets
i L1    1.1.0.0 [115/20] via 10.1.12.1, FastEthernet0/0
i L1    1.1.1.0 [115/20] via 10.1.12.1, FastEthernet0/0
i L1    1.1.2.0 [115/20] via 10.1.12.1, FastEthernet0/0
i L1    1.1.3.0 [115/20] via 10.1.12.1, FastEthernet0/0
i L2 3.0.0.0/8 [115/20] via 10.1.23.3, Serial0/0.23
i L2 4.0.0.0/8 [115/30] via 10.1.23.3, Serial0/0.23
i L2 5.0.0.0/8 [115/30] via 10.1.26.6, Serial0/0.26
i L2 6.0.0.0/8 [115/20] via 10.1.26.6, Serial0/0.26
     10.0.0.0/24 is subnetted, 6 subnets
i L2    10.1.45.0 [115/30] via 10.1.26.6, Serial0/0.26
                  [115/30] via 10.1.23.3, Serial0/0.23
i L2    10.1.34.0 [115/20] via 10.1.23.3, Serial0/0.23
i L2    10.1.56.0 [115/20] via 10.1.26.6, Serial0/0.26
```

On R3

```
R3#Show ip route isis

     1.0.0.0/24 is subnetted, 4 subnets
i L2    1.1.0.0 [115/30] via 10.1.23.2, Serial0/0.32
i L2    1.1.1.0 [115/30] via 10.1.23.2, Serial0/0.32
i L2    1.1.2.0 [115/30] via 10.1.23.2, Serial0/0.32
i L2    1.1.3.0 [115/30] via 10.1.23.2, Serial0/0.32
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
i L2 2.0.0.0/8 [115/20] via 10.1.23.2, Serial0/0.32
i L1 4.0.0.0/8 [115/20] via 10.1.34.4, FastEthernet0/0
i L1 5.0.0.0/8 [115/30] via 10.1.34.4, FastEthernet0/0
i L2 6.0.0.0/8 [115/30] via 10.1.23.2, Serial0/0.32
     10.0.0.0/24 is subnetted, 6 subnets
i L2    10.1.12.0 [115/20] via 10.1.23.2, Serial0/0.32
i L2    10.1.26.0 [115/20] via 10.1.23.2, Serial0/0.32
i L1    10.1.45.0 [115/20] via 10.1.34.4, FastEthernet0/0
i L2    10.1.56.0 [115/30] via 10.1.23.2, Serial0/0.32
```

On R4

```
R4#Show ip route isis

i L1 3.0.0.0/8 [115/20] via 10.1.34.3, FastEthernet0/0
i L1 5.0.0.0/8 [115/20] via 10.1.45.5, FastEthernet0/1
i*L1 0.0.0.0/0 [115/10] via 10.1.45.5, FastEthernet0/1
               [115/10] via 10.1.34.3, FastEthernet0/0
```

On R5

```
R5#Show ip route isis

     1.0.0.0/24 is subnetted, 4 subnets
i L2    1.1.0.0 [115/40] via 10.1.56.6, Serial0/0.56
i L2    1.1.1.0 [115/40] via 10.1.56.6, Serial0/0.56
i L2    1.1.2.0 [115/40] via 10.1.56.6, Serial0/0.56
i L2    1.1.3.0 [115/40] via 10.1.56.6, Serial0/0.56
i L2 2.0.0.0/8 [115/30] via 10.1.56.6, Serial0/0.56
i L1 3.0.0.0/8 [115/30] via 10.1.45.4, FastEthernet0/1
i L1 4.0.0.0/8 [115/20] via 10.1.45.4, FastEthernet0/1
i L2 6.0.0.0/8 [115/20] via 10.1.56.6, Serial0/0.56
     10.0.0.0/24 is subnetted, 6 subnets
i L2    10.1.12.0 [115/30] via 10.1.56.6, Serial0/0.56
i L2    10.1.26.0 [115/20] via 10.1.56.6, Serial0/0.56
i L2    10.1.23.0 [115/30] via 10.1.56.6, Serial0/0.56
i L1    10.1.34.0 [115/20] via 10.1.45.4, FastEthernet0/1
```

On R6

```
R6#Show ip route isis

     1.0.0.0/24 is subnetted, 4 subnets
i L2    1.1.0.0 [115/30] via 10.1.26.2, Serial0/0.62
i L2    1.1.1.0 [115/30] via 10.1.26.2, Serial0/0.62
i L2    1.1.2.0 [115/30] via 10.1.26.2, Serial0/0.62
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
i L2    1.1.3.0 [115/30] via 10.1.26.2, Serial0/0.62
i L2 2.0.0.0/8 [115/20] via 10.1.26.2, Serial0/0.62
i L2 3.0.0.0/8 [115/30] via 10.1.26.2, Serial0/0.62
i L2 4.0.0.0/8 [115/30] via 10.1.56.5, Serial0/0.65
i L2 5.0.0.0/8 [115/20] via 10.1.56.5, Serial0/0.65
     10.0.0.0/24 is subnetted, 6 subnets
i L2    10.1.12.0 [115/20] via 10.1.26.2, Serial0/0.62
i L2    10.1.23.0 [115/20] via 10.1.26.2, Serial0/0.62
i L2    10.1.45.0 [115/20] via 10.1.56.5, Serial0/0.65
i L2    10.1.34.0 [115/30] via 10.1.56.5, Serial0/0.65
                  [115/30] via 10.1.26.2, Serial0/0.62
```

Task 2

The default Designated Intermediate System (DIS) router should be configured to be the non-default peer between R1 and R2. Do not modify the router-id or R1 to accomplish this task.

Verify and change the current DIS:

On R2

```
R2#Show isis database level-1

IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00              0x00000004   0x48EC        1178              0/0/0
R2.00-00            * 0x00000005   0x4933        999               1/0/0
R2.01-00            * 0x00000002   0xD619        1044              0/0/0
```

R2 has been selected as the DIS router due to the higher router-id.

```
R2(config)#int f0/0
R2(config-if)#isis priority 1

R2#Show isis database

IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00              0x00000005   0x4632        1148              0/0/0
R1.01-00              0x00000001   0x92A4        1148              0/0/0
R2.00-00            * 0x00000006   0x3689        1150              1/0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2.01-00          * 0x00000002   0x5E13        0 (1150)          0/0/0

IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R2.00-00          * 0x00000007   0x5D31        826               0/0/0
R3.00-00            0x00000005   0x6DE9        618               0/0/0
R5.00-00            0x00000006   0x0190        616               0/0/0
R6.00-00            0x00000007   0xEE91        1158              0/0/0
```

Notice how ISIS shows the database differently than OSPF in that the timers decrement from the default value of 1200. Another difference is the way the database displays a change when it occurs. In this case, the DIS router has been changed due to the priority of R2 being lowered on F0/0. R1 is now the DIS but the R2.01-00 LSA shows that the LSA has decremented to "0" automatically to show the LSA is not being used currently. However, the real time that the LSA will be purged is in parenthesis (1193 seconds).

Task 3

Ensure that the default route for R4 to leave the area is through R5 only.

Verify and change the default Level-2 router for R4

On R4

```
R4#Show ip route isis

i L1 3.0.0.0/8 [115/20] via 10.1.34.3, FastEthernet0/0
i L1 5.0.0.0/8 [115/20] via 10.1.45.5, FastEthernet0/1
i*L1 0.0.0.0/0 [115/10] via 10.1.45.5, FastEthernet0/1
               [115/10] via 10.1.34.3, FastEthernet0/0

R4(config)#int f0/1
R4(config-if)#isis metric 5

R4#Show ip route isis

i L1 3.0.0.0/8 [115/20] via 10.1.34.3, FastEthernet0/0
i L1 5.0.0.0/8 [115/15] via 10.1.45.5, FastEthernet0/1
i*L1 0.0.0.0/0 [115/5] via 10.1.45.5, FastEthernet0/1
```

On R1

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 80/80/80 ms
```

Shutdown the s0/0.62 interface on R6 and proceed to the next statement!!

On R6

```
R6(config)#Int S0/0.62
R6(config-subif)#Shutdown
```

On R1

```
R1#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

The ping fails due to a sub optimal path that is now present in ISIS. The route from R1 should choose the following path to R4: R1 > R2 > R3 > R4. However the path from R4 to R1 follows the following path: R4 > R5 > R6 > R2 > R1. We verified this by shutting the s0/0.62 interface and repeating the test.

Task 4

Fix the problem of sub-optimal paths from Task 3 for the 1.1.1.1 loopback interface only. You are not permitted to use statics, modify any cost on any links or inject a default route with any cost change. Ensure that flooding of LSA's takes place before convergence finishes on R6. Once the convergence takes place R6 should only run SPF on changes within 10 seconds.

On R3

```
R3(config)#access-list 1 permit 1.1.1.0 0.0.0.255

R3(config)#route-map TST1111 permit 10
R3(config-route-map)#match ip address 1
R3(config-route-map)#router isis
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-router)#redistribute isis ip level-2 into level-1 route-map TST1111
```

The configuration used above is called "Route Leaking and is used to overcome sub-optimal paths due to weak ISIS metric selection. The "clear isis *" command MAY be necessary on R4 or R3 or both.

Route Leaking tends to not work when routers are in the same area. Also, the option to perform route-leaking on R2 into Level-2 would only fail due to the default route given by R3.

To verify the configuration:

On R4

```
R4#Sh ip route

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is 10.1.45.5 to network 0.0.0.0

     1.0.0.0/24 is subnetted, 1 subnets
i ia    1.1.1.0 [115/40] via 10.1.34.3, FastEthernet0/0
i L1 3.0.0.0/8 [115/20] via 10.1.34.3, FastEthernet0/0
C    4.0.0.0/8 is directly connected, Loopback0
i L1 5.0.0.0/8 [115/15] via 10.1.45.5, FastEthernet0/1
     10.1.0.0/24 is subnetted, 2 subnets
C       10.1.34.0 is directly connected, FastEthernet0/0
C       10.1.45.0 is directly connected, FastEthernet0/1
i*L1 0.0.0.0/0 [115/5] via 10.1.45.5, FastEthernet0/1
```

Notice the presence of the new ISIS inter area (ia) route.

On R1

```
R1#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
Packet sent with a source address of 1.1.0.1
….
Success rate is 0 percent (0/5)
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1#Ping 4.4.4.4 source 1.1.1.1

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 12/20/48 ms
```

The route is now symmetrical and flows optimally due to the more accurate routing information R4 now has !!

On R6

```
R6(config)#router isis
R6(config-router)#fast flood
R6(config-router)#ispf level-2 10
```

Task 5

Remove the filter created in the previous Task. Restore the connection from R6 to R2. Verify the suboptimal path from R4 to the 1.1.1.1 prefix.

On R3

```
R3(config)#router isis
R3(config-router)#NO redistribute isis ip level-2 into level-1 route-map
TST1111
R3(config-router)#exit
R3(config)#NO route-map TST1111
```

On R4

```
R4#Trace 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 10.1.45.5 4 msec 16 msec 24 msec
  2 10.1.56.6 16 msec 28 msec 40 msec
  3 10.1.26.2 24 msec 36 msec 16 msec
  4 10.1.12.1 20 msec *  40 msec
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 6

Allow fast convergence to the 1.1.1.1 prefix from R4 due to a direct backbone link failure.

On R4

```
R4#Sh ip route isis

i L1 3.0.0.0/8 [115/20] via 10.1.34.3, FastEthernet0/0
i L1 5.0.0.0/8 [115/15] via 10.1.45.5, FastEthernet0/1
i*L1 0.0.0.0/0 [115/5] via 10.1.45.5, FastEthernet0/1
```

The default route currently follows the path through R5.

On R5

```
R5#Sh isis database level-1

IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R3.00-00              0x00000017   0x0645        1188              1/0/0
R4.00-00              0x00000016   0x69DE        705               0/0/0
R4.01-00              0x0000000D   0xF15B        736               0/0/0
R5.00-00            * 0x00000014   0xC7A2        1103              1/0/0
R5.01-00            * 0x0000000F   0x89F2        995               0/0/0
```

The attachment bit is set on level-1 routes by R3 and R5. This bit is set to allow the area to converge in the event of a failure. If the R5 backbone links fails, the attachment bit will NOT be set by R5 and will not allow R4 to use this path. R3 will then be used to leave the area.

```
R5(config)#int s0/0.56
R5(config-subif)#shut
```

On R4

```
R4#Sh isis database

IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R3.00-00              0x00000017   0x0645        737               1/0/0
R4.00-00            * 0x00000017   0x67DF        942               0/0/0
R4.01-00            * 0x0000000E   0xEF5C        964               0/0/0
R5.00-00              0x00000015   0xBDB3        1108              0/0/0
R5.01-00              0x0000000F   0x89F2        540               0/0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

As you can see, R4 no longer receives the route from R5 with the attachment bit set.

On R5

```
R5(config)#int s0/0.56
R5(config-subif)#NO shut
```

The bit will be set again now that the link is restored. This action is a result of the attached bit being tied to the backbone links ONLY.

On R6

```
R6(config)#int s0/0.62
R6(config-subif)#shut
```

On R5

```
R5#Sh isis database level-1

IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R3.00-00              0x00000018   0x0446        777               1/0/0
R4.00-00              0x00000018   0x65E0        853               0/0/0
R4.01-00              0x0000000F   0xED5D        1013              0/0/0
R5.00-00            * 0x00000016   0xC3A4        744               1/0/0
R5.01-00            * 0x00000010   0x87F3        575               0/0/0
```

The bit is still set due to the indirect failure on a different backbone router.

On R4

```
R4#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
U.U.U
Success rate is 0 percent (0/5)
```

On R6

```
R6(config)#int s0/0.62
R6(config-subif)#NO shut
```

On R5

```
R5#Sh clns route
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Codes: C - connected, S - static, d - DecnetIV
       I - ISO-IGRP,  i - IS-IS,  e - ES-IS
       B - BGP,       b - eBGP-neighbor

C  49.0345.5555.0000.5555.00 [1/0], Local IS-IS NET
C  49.0345 [2/0], Local IS-IS Area

R5(config)#int s0/0.56
R5(config-subif)#clns router isis
```

On R6

```
R6(config)#int s0/0.65
R6(config-subif)#clns router isis

R6(config)#int s0/0.62
R6(config-subif)#clns router isis
```

On R2

```
R2(config)#int s0/0.26
R2(config-subif)#clns router isis
```

On R5

```
R5#Sh clns route | i i

Codes: C - connected, S - static, d - DecnetIV
       I - ISO-IGRP,  i - IS-IS,  e - ES-IS
       B - BGP,       b - eBGP-neighbor
i  49.0012 [110/20]
      via R6, Serial0/0.56
i  49.0006 [110/10]
      via R6, Serial0/0.56

R5(config)#clns filter-set TST permit 49.0012
R5(config-router)#route-map CLNS permit 10
R5(config-route-map)#match clns address TST

R5(config-route-map)#router isis
R5(config-router)#set-attached-bit route-map CLNS
```

On R6

```
R6(config)#int s0/0.62
R6(config-subif)#shut
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R5

```
R5#Sh isis database level-1 | i R5

R5.00-00              * 0x00000019   0x9E68        1199              0/0/0
R5.01-00              * 0x00000011   0x85F4        838               0/0/0
```

On R4

```
R4#Sh ip route | i 0.0.0.0

Gateway of last resort is 10.1.34.3 to network 0.0.0.0
     10.0.0.0/24 is subnetted, 2 subnets
i*L1 0.0.0.0/0 [115/10] via 10.1.34.3, FastEthernet0/0

R4#Ping 1.1.1.1

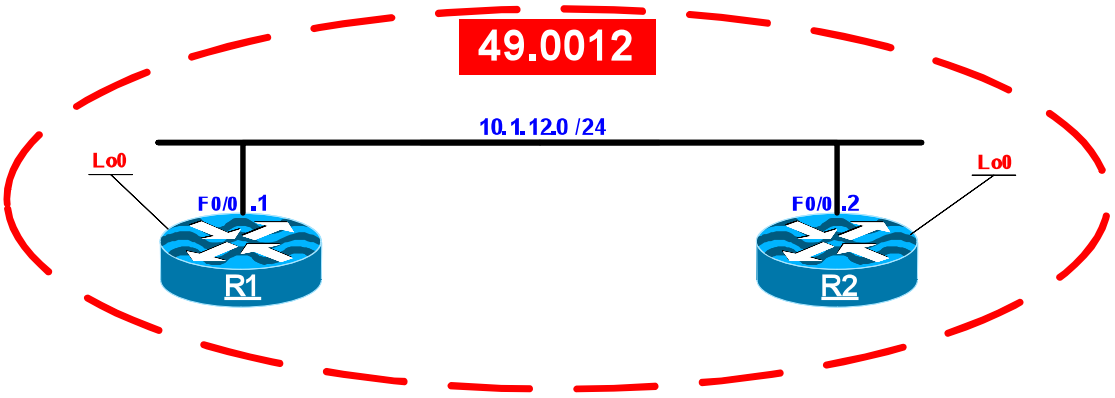
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/60/116 ms
```

Task 8

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 6
Authentication



Lab Setup:

Router	Interface	IP address
R1	Lo0 F0/0	1.1.1.1 /8 10.1.12.1/24
R2	Lo0 F0/0	2.2.2.2 /8 10.1.12.2 /24

Task 1

Configure ISIS on R1 and R2 with an area of 49.0012. The area should use Extended TLV type 135 and allow the fewest adjacencies necessary. Use system ID of 1111.0000.1111 for R1 and 2222.0000.2222 for R2. Both Loopbacks and F0/0 ports should participate in updates

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1

```
R1(config)#router isis
R1(config-router)#net 49.0012.1111.0000.1111.00
R1(config-router)#metric-style wide
R1(config-router)#is-type level-1

R1(config-router)#int fa0/0
R1(config-if)#ip router isis

R1(config-if)#int lo0
R1(config-if)#ip router isis

R1(config-if)#router isis
R1(config-router)#passive lo0
```

On R2

```
R2(config)#router isis
R2(config-router)#net 49.0012.2222.0000.2222.00
R2(config-router)#metric-style wide
R2(config-router)#is-type level-1

R2(config)#int f0/0
R2(config-if)#ip router isis

R2(config-if)#int lo0
R2(config-if)#ip router isis

R2(config)#router isis
R2(config-router)#passive lo0
```

On R1

```
R1#Ping 2.2.2.2 source 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 12/20/48 ms
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 2

Allow password authentication for all sequence numbered packets in the area with a password of "cisco". Use the stronger encryption type.

On R1

```
R1(config)#key chain TST
R1(config-keychain)#key 1
R1(config-keychain-key)#key-string cisco

R1(config-keychain-key)#int f0/0
R1(config-if)#isis authentication mode md5
R1(config-if)#isis authentication key-chain TST level-1
```

These commands can be applied as an area based configuration if used directly under the isis process. Very similar to OSPF, except the area authentication is not required for ISIS to work.

To Verify that the password is not vaidated on R2:

```
R2#Sh isis neighbor
```

System Id	Type	Interface	IP Address	State	Holdtime	Circuit Id
R1	L1	Fa0/0	10.1.12.1	INIT	28	R2.02

On R2

```
R2(config)#key chain TST
R2(config-keychain)#key 1
R2(config-keychain-key)#key-string cisco
R2(config-keychain-key)#int f0/0
R2(config-if)#isis authentication mode md5
R2(config-if)#isis authentication key-chain TST level-1

R2#Sh isis n
```

System Id	Type	Interface	IP Address	State	Holdtime	Circuit Id
R1	L1	Fa0/0	10.1.12.1	UP	27	R2.02

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task3

Allow password authentication for all sequence numbered packets in the area with a password of "cisco". You may not use key-strings to accomplish the task.

On Both Routers:

```
RX(config)#int f0/0
RX(config-if)#NO isis authentication mode MD5
RX(config-if)#NO isis authentication key-chain TST level-1
RX(config-if)#exit

RX(config)#NO key chain TEST
```

To Verify that the password is validated on R2:

On R1

```
R1#Sh isis n
```

System Id	Type	Interface	IP Address	State	Holdtime	Circuit Id
R2	L1	Fa0/0	10.1.12.2	UP	8	R2.01

On R2 *(example)*

```
R2#Show key chain TST

Key-chain TST:
   key 1 -- text "cisco"
      accept lifetime (always valid) - (always valid) [valid now]
      send lifetime (always valid) - (always valid) [valid now]
```

The benefit to moving to key chains was having the ability to age out the key gracefully and replace it with another key without the session dropping.

Reconfigure password within the area On Both Routers:

```
(config)#router isis
(config-router)#area-password CCIE
```

This configuration would represent performing encryption with the fewest key strokes.

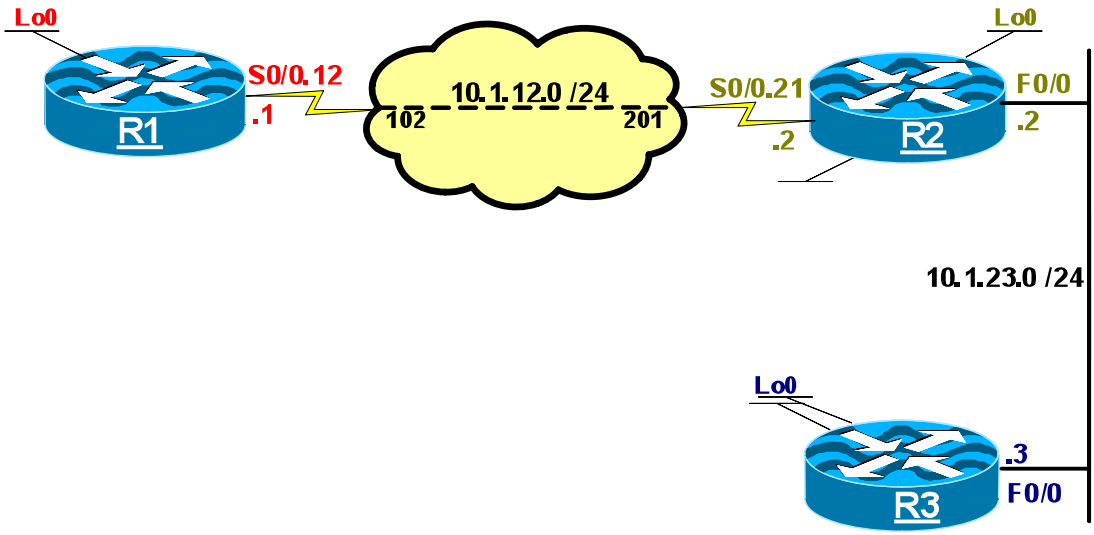
©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task4

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 7
Advanced Configuration I



IP Addressing:

Router	Interface	IP address
R1	Lo0 S0/0.12	1.1.1.1 /8 10.1.12.1/24
R2	Lo0 S0/0.21 F0/0	2.2.2.2 /8 10.1.12.1 /24 10.1.23.2/ 24
R3	Lo0 F0/0	3.3.3.3 /8 10.1.23.3/ 24

Lab Setup:

- The frame Relay circuit should be configured as a multipoint sub-interface.
- R1 and R2 are configured in VLAN 23.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 1

Configure R1 in area 49.0011, R2 and R3 in area 49.0023. The System ID's must be unique. System ID's may reflect that of the previous lab for simplicity. All interfaces must be configured to run ISIS.

On R1

```
R1(config)#router isis
R1(config-router)#net 49.0011.1111.0000.1111.00
R1(config-router)#is-type level-2-only

R1(config-router)#int s0/0.12 multipoint
R1(config-subif)#ip router isis

R1(config-subif)#int lo0
R1(config-if)#ip router isis

R1(config-if)#router isis
R1(config-router)#passive lo0
```

On R2

```
R2(config)#router isis
R2(config-router)#net 49.0023.2222.0000.2222.00
R2(config-router)#metric-style wide

R2(config-router)#int s0/0.21 multipoint
R2(config-subif)#ip router isis
R2(config-subif)#isis circuit-type level-2-only

R2(config-subif)#int f0/0
R2(config-if)#ip router isis
R2(config-if)#isis circuit-type level-1

R2(config-if)#router isis
R2(config-router)#passive lo0
```

On R3

```
R3(config)#router isis
R3(config-router)#net 49.0023.3333.0000.3333.00

R3(config-router)#is-type level-1
R3(config-router)#metric-style wide
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-router)#int f0/0
R3(config-if)#ip router isis

R3(config-if)#int lo0
R3(config-if)#ip router isis

R3(config-if)#router isis
R3(config-router)#passive lo0
```

To Verify the Configuration:

On R2

```
R2#Sh isis n
```

System Id	Type	Interface	IP Address	State	Holdtime	Circuit Id
R3	L1	Fa0/0	10.1.23.3	UP	9	R2.02

The peering session with R1 is not established even though the mapping of Frame-Relay has been completed. The CLNS messages must also be passed to share topology data with each of the ISIS routers. The broadcast key word is required to have scenario work just as it did in the normal frame map statement.

Task 2

Configure the routers to have complete ISIS reachability.

To Verify the current Connection:

On R2

```
R2#Sh isis n
```

System Id	Type	Interface	IP Address	State	Holdtime	Circuit Id
R3	L1	Fa0/0	10.1.23.3	UP	9	R3.01

R2 does not see R1 as a peer.

```
R2#Sh frame pvc | s PVC
```

PVC Statistics for interface Serial0/0 (Frame Relay DTE)

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

DLCI = 201, DLCI USAGE = LOCAL, PVC STATUS = **ACTIVE**, INTERFACE = Serial0/0.21

```
R2#Sh frame map

Serial0/0.21 (up): ip 10.1.12.1 dlci 201(0xC9,0x3090), static,
              broadcast,
              CISCO, status defined, active
```

Frame Relay shows that Layer 2 is operating correctly and mapping has taken place.

On R1

```
R1#Sh isis n
```

System Id Type Interface IP Address State Holdtime Circuit Id

```
R1#Show frame pvc | s DLCI = 102

PVC Statistics for interface Serial0/0 (Frame Relay DTE)
DLCI = 102, DLCI USAGE = LOCAL, PVC STATUS = ACTIVE, INTERFACE = Serial0/0.12
```

R1 also does not show peering but Layer 2 looks fine.

Configuring ISIS frame mapping:

On R1

```
R1(config)#int s0/0.12 multipoint
R1(config-subif)#frame map clns 102 broadcast
```

On R2

```
R2(config)#int s0/0.21 multipoint
R2(config-subif)#frame map clns 201 broadcast
```

The CLNS data must also be mapped for frame relay to work properly for ISIS.

```
R2#Sh isis n
```

System Id	Type	Interface	IP Address	State	Holdtime	Circuit Id
R1	L2	Se0/0.21	10.1.12.1	UP	27	R2.01
R3	L1	Fa0/0	10.1.23.3	UP	7	R3.01

The peer is now operational.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2#Sh ip route | in L

       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
i L2     1.0.0.0 [115/10] via 10.1.12.1, Serial0/0.21
C        2.2.2.2 is directly connected, Loopback0
i L1    3.0.0.0/8 [115/10] via 10.1.23.3, FastEthernet0/0
```

On R1

```
R1#Sh ip route | i L2

       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
i L2    2.0.0.0 [115/10] via 10.1.12.2, Serial0/0.12
i L2    3.0.0.0/8 [115/20] via 10.1.12.2, Serial0/0.12
i L2    10.1.23.0 [115/20] via 10.1.12.2, Serial0/0.12
```

On R3

```
R3#Ping 1.1.1.1 source 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 3.3.3.3
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 20/22/28 ms
```

Task 3

Modify the MTU on R1 to be 2000. Ensure correct ISIS communication without changing the physical MTU on R1 or R2. R1 should not pay attention to ISIS errors.

On R1

```
R1(config)#int s0/0
R1(config)#shut
R1(config-if)#mtu 2000
R1(config)#NO shut

R1#Sh isis n

System Id      Type   Interface   IP Address      State  Holdtime  Circuit Id
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2          L2      Se0/0.12    10.1.12.2       INIT   26        R2.02

R1#Debug isis adj-packets
IS-IS Adjacency related packets debugging is on

*Mar  1 00:04:49.319: ISIS-Adj: Sending L2 LAN IIH on Serial0/0.12, length 2000

*Mar  1 00:04:52.287: ISIS-Adj: Rec L2 IIH from DLCI 102 (Serial0/0.12), cir
type L2, cir id 2222.0000.2222.02, length 1500

R1#No debug all

*Mar  1 00:04:57.891: ISIS-Adj: Sending L2 LAN IIH on Serial0/0.12,
length 2000
```

When an ISIS routers sends hello packets, it pads the packets up to the MTU size by default. R1 will see the hello packets from R2 as 1500 and will process them. In this case, R1 has a higher MTU and will not see a problem. R2 will see an MTU of 2000 bytes and will have issues with this as it exceeds the MTU on the interface.

On R1

```
R1(config)#int s0/0.12 multipoint
R1(config-subif)#clns mtu 1500

*Mar  1 00:09:34.375: ISIS-Adj: Sending L2 LAN IIH on Serial0/0.12, length 1500

*Mar  1 00:09:36.515: ISIS-Adj: Rec L2 IIH from DLCI 102 (Serial0/0.12), cir
type L2, cir id 2222.0000.2222.02, length 1500
```

The change is made by modifying the clns mtu on R1 to match the other side. Now, when R1 sends padded hello packets, R2 will not see the problem anymore. However, R1 will continue to pad the hello's to the newly configured value of 1500 indefinately.

```
R1(config)#int s0/0.12 multipoint
R1(config-subif)#NO isis hello padding

R1#Debug isis adj-packets

*Mar  1 00:14:40.443: ISIS-Adj: Rec L2 IIH from DLCI 102 (Serial0/0.12), cir
type L2, cir id 2222.0000.2222.02, length 1500   This was the fourth hello!!

R1#No debug all

*Mar  1 00:14:43.251: ISIS-Adj: Rec L2 IIH from DLCI 102 (Serial0/0.12), cir
type L2, cir id 2222.0000.2222.02, length 1500
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1#NO debug all

*Mar  1 00:14:47.555: ISIS-Adj: Sending L2 LAN IIH on Serial0/0.12, length 55
This was the sixth hello!!
```

After disabling the hello pading on R1, the first 5 hello's are padded to make sure any MTU problem will be found. The sixth MTU moving forward will not be padded to ensure efficient bandwidth usage.

ISIS error handling:

On R1

```
R1(config)#router isis
R1(config-router)#ignore-lsp-errors
```

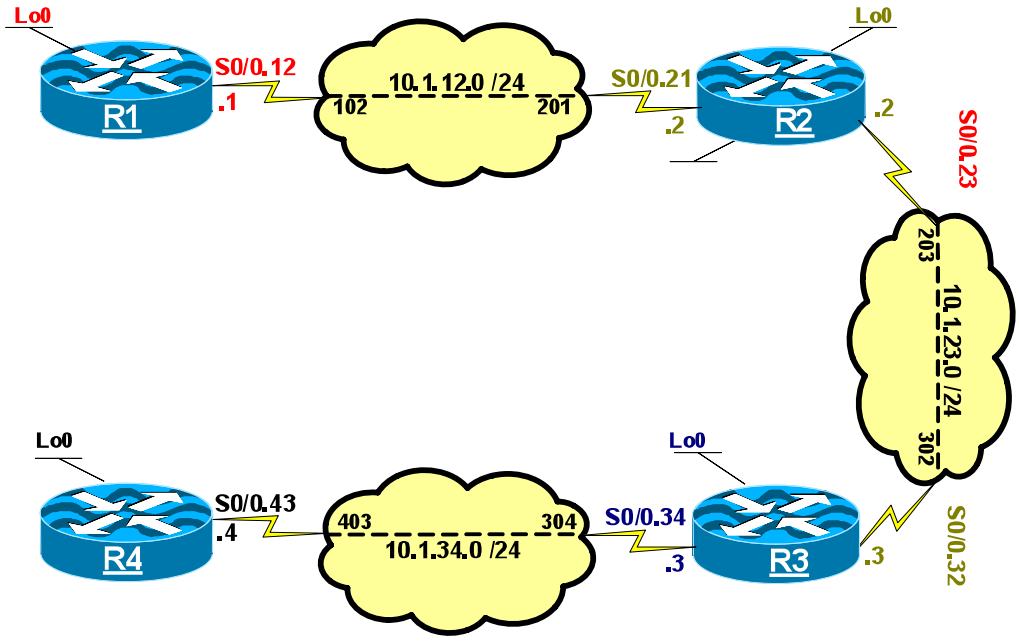
This command is used to prevent an isis router from purging an LSP in the dadatbase due to a bad checksum. The action of purging an LSA will only result in the peering router to resend the lsa again and again. By ignoring the LSPDU, we prevent peering router from creating the LSP again until the default time to do so.

Task4

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 8
Advanced Configuration II



IP Addressing:

Router	Interface	IP address
R1	Lo0 S0/0.12	1.1.1.1 /8 10.1.12.1/24
R2	Lo0 S0/0.21 S0/0.23	2.2.2.2 /8 10.1.12.1 /24 10.1.23.2/ 24
R3	Lo0 S0/0.32 S0/0.34	3.3.3.3 /8 10.1.23.3/ 24 10.1.34.3/ 24
R4	Lo0 S0/0.43	4.4.4.4 /8 10.1.34.4/ 24

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab Setup:

- The lab must be configured according to the table above
- No sub-interfaces are permitted on R4. All other Frame-Relay interfaces should be configured in a point-to-point manor.
- R2 and R4 should be in VLAN 24

Task 1

Configure R1, R2, R3 and R4 in area 49.1234. Level 2 adjacencies between R1 and R2 as well R3 to R4. The session between R2 and R3 should be level-1. The System ID's must be unique. System ID's may reflect that of the previous lab for simplicity. All interfaces must be configured to run ISIS .

On R1

```
R1(config)#router isis
R1(config-router)#net 49.1234.1111.0000.1111.00
R1(config-router)#metric-style wide
R1(config-router)#is-type level-2-only

R1(config-router)#int s0/0.12
R1(config-subif)#ip router isis

R1(config-subif)#int lo0
R1(config-if)#ip router isis
```

On R2

```
R2(config)#router isis
R2(config-router)#net 49.1234.2222.0000.2222.00
R2(config-router)#metric-style wide

R2(config-router)#int s0/0.21
R2(config-subif)#ip router isis
R2(config-subif)#isis circuit type level-2-only

R2(config-subif)#int s0/0.23
R2(config-subif)#ip router isis
R2(config-subif)#isis circuit type level-1

R2(config-subif)#int lo0
R2(config-if)#ip router isis
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3

```
R3(config)#router isis
R3(config-router)#net 49.1234.3333.0000.3333.00
R3(config-router)#metric-style wide

R3(config-router)#int s0/0.32
R3(config-subif)#ip router isis
R3(config-subif)#isis circuit type level-1

R3(config-subif)#int s0/0.34
R3(config-subif)#ip router isis
R3(config-subif)#isis circuit type level-2-only

R3(config-subif)#int lo0
R3(config-if)#ip router isis
```

On R4

```
R4(config)#router isis
R4(config-router)#net 49.1234.4444.0000.4444.00
R4(config-router)#metric-style wide
R4(config-router)#is-type level-2-only

R4(config-router)#int s0/0
R4(config-subif)#ip router isis

R1(config-subif)#int lo0
R1(config-if)#ip router isis
```

Verify the Configuration:

On R4

```
R4#Sh isis n

System Id      Type Interface IP Address      State Holdtime Circuit Id

R4#Sh ip route | b Gate
Gateway of last resort is not set

C    4.0.0.0/8 is directly connected, Loopback0
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.34.0 is directly connected, Serial0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4#Debug isis adj-packets s0/0
IS-IS Adjacency related packets debugging is on

*Mar  1 00:36:15.975: ISIS-Adj: Encapsulation failed for L2 LAN IIH on Serial0/0

*Mar  1 00:36:19.391: ISIS-Adj: Rec serial IIH from DLCI 403 (Serial0/0), cir
type L2, cir id 01, length 1499

*Mar  1 00:36:19.395: ISIS-Adj: Point-to-point IIH received on multi-point
interface: ignored IIH

R4#U all
```

On R4

```
R4(config)#int s0/0
R4(config-if)#NO ip address

R4(config-if)#frame interface-dlci 403 ppp virtual-Template 43
R4#(config-fr-dlci)#exit

R4(config-if)#int virtual-template 43
R4(config-if)#ip address 10.1.34.4 255.255.255.0
R4(config-if)#encap ppp
R4(config-if)#ip router isis
```

On R3

```
R3(config)#int s0/0.34
R3(config-subif)#NO ip address
R3(config-subif)#frame interface-dlci 304 ppp virtual-Template 34

R3(config-fr-dlci)#int virtual-template 34
R3(config-if)#ip address 10.1.34.3 255.255.255.0
R3(config-if)#encap ppp
R3(config-if)#ip router isis

R3#Sh isis n | b Sys

System Id      Type Interface IP Address      State Holdtime Circuit Id
R2             L1   Se0/0.32  10.1.23.2       UP    26       01
R4             L1L2 Vi1       10.1.34.4       UP    28       01
```

On R4

```
R4#Sh isis n | b Sys
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
System Id      Type Interface IP Address      State Holdtime Circuit Id
R3               L2 Vi2       10.1.34.3       UP    22       03
```

Task 2

Configure R1, R2, R3 and R4 in area 49.1234. Level 2 adjacencies between R1 and R2 as well R3 to R4. The session between R2 and R3 should be level-1. The System ID's

On R4

```
R4#Sh ip route isis

i L1 2.0.0.0/8 [115/30] via 10.1.34.3, Virtual-Access1
i L1 3.0.0.0/8 [115/20] via 10.1.34.3, Virtual-Access1
     10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
i L1    10.1.23.0/24 [115/20] via 10.1.34.3, Virtual-Access2
```

R4 does not see R1's loopback address; This establishes the point that R1 originates the route as a Level-2 prefix. When the route is passed to R2 across the L2 session the only place to cross is from one L1 router to the other, Routes cannot be redistributed between 2 routers that are in the same area unless the leaking is performed from L2 to L1 or L1 to L2. L1 does not leak into L1.(as of 12.4 Mainline or Train release)

On R3

```
R3#Sh ip route isis

i L1 2.0.0.0/8 [115/20] via 10.1.23.2, Serial0/0.32
i L2 4.0.0.0/8 [115/20] via 10.1.34.4, Virtual-Access1
```

On R2

```
R2#Sh ip route isis

i L2 1.0.0.0/8 [115/20] via 10.1.12.1, Serial0/0.21
i L1 3.0.0.0/8 [115/20] via 10.1.23.3, Serial0/0.23
     10.0.0.0/24 is subnetted, 3 subnets
i L1    10.1.34.0 [115/20] via 10.1.23.3, Serial0/0.23
```

R2 and R3 have similar problems. This is also where the problem exists.

On R1

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1#Sh ip route isis

i L2 2.0.0.0/8 [115/20] via 10.1.12.2, Serial0/0.12
i L2 3.0.0.0/8 [115/30] via 10.1.12.2, Serial0/0.12
     10.0.0.0/24 is subnetted, 3 subnets
i L2    10.1.23.0 [115/20] via 10.1.12.2, Serial0/0.12
i L2    10.1.34.0 [115/30] via 10.1.12.2, Serial0/0.12
```

R1 has no visibility to the 4.0.0.0 prefix as an ISIS route.

Task 3

Allow R1 to ping R4 loopback successfully with a normal or source ping from the loopback interface. The use of static routes are not permitted.

To Verify the Configuration:

On R3

```
R3#Sh ip route 4.4.4.4

Routing entry for 4.0.0.0/8
  Known via "isis", distance 115, metric 10, type level-2
  Redistributing via isis
  Last update from 10.1.34.4 on Serial0/0.34, 00:37:15 ago
  Routing Descriptor Blocks:
  * 10.1.34.4, from 4.4.4.4, via Serial0/0.34
      Route metric is 10, traffic share count is 1
```

On R2

```
R2#Sh ip route 4.4.4.4
% Network not in table
```

R2 does not have the 4.4.4.4 route is not in the table so it is safe to assume that R1 will not either. We can fix this through redistribution the same way we learned route-leaking earlier.

On R3

```
R3(config)#access-list 100 permit ip any any
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config)#router isis
R3(config-router)#redistribute isis ip level-2 into level-1 distribute-list 100
```

We use a distribute-list instead of the route-map for diversity in covering redistribution again with different options.

On R2

```
R2#Show ip route isis

i L2 1.0.0.0/8 [115/20] via 10.1.12.1, Serial0/0.21
i L1 3.0.0.0/8 [115/20] via 10.1.23.3, Serial0/0.23
i ia 4.0.0.0/8 [115/30] via 10.1.23.3, Serial0/0.23
     10.0.0.0/24 is subnetted, 3 subnets
i L1    10.1.34.0 [115/20] via 10.1.23.3, Serial0/0.23
```

R2 can now see the route due to redistribution. R1 as we can see below, will not see the route. R2 will not share the new ISIS Inter area route to a Level-2 session in the same area. Redistribution would not work either since R1 and R2 are in the same area.

```
R2(config)#access-list 100 permit ip any any

R2(config)#router isis
R2(config-router)#redistribute isis ip level-2 into level-1 distribute-list 100
```

On R3

```
R3#Sh ip route isis

i ia 1.0.0.0/8 [115/30] via 10.1.23.2, Serial0/0.32
i L1 2.0.0.0/8 [115/20] via 10.1.23.2, Serial0/0.32
i L2 4.0.0.0/8 [115/20] via 10.1.34.4, Virtual-Access1
     10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
i ia    10.1.12.0/24 [115/20] via 10.1.23.2, Serial0/0.32
```

R2 must perform the same action as the same problem exists in the other direction.

On R1

```
R1#Sh ip route 4.4.4.4
% Network not in table
```

The Solution:

On R2

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config)#router isis
R2(config-router)#default-information originate
```

On R3

```
R3(config)#router isis
R3(config-router)#default-information originate
```

The "default-information originate" command will inject a default into the level-2 session only since default routes will not be allowed into a L1 session from the L1 router database. We must perform this action on R3 as the source of the ping will not reachable due to the same hardships that were experienced by the 4.4.4.4 prefix in the other direction.

Verifying the working Configuration:

On R1

```
R1#Sh ip route isis | i 0.0.0.0

     10.0.0.0/24 is subnetted, 3 subnets
i*L2 0.0.0.0/0 [115/10] via 10.1.12.2, Serial0/0.12

R1#Ping 4.4.4.4  source 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 20/32/40 ms
```

On R4

```
R4>Sh ip route 0.0.0.0

Routing entry for 0.0.0.0/0, supernet
  Known via "isis", distance 115, metric 10, candidate default path, type
level-2
  Redistributing via isis
  Last update from 10.1.34.3 on Virtual-Access1, 00:00:14 ago
  Routing Descriptor Blocks:
    * 10.1.34.3, from 3.3.3.3, via Virtual-Access1
      Route metric is 10, traffic share count is 1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 4

Configure an ip address of 10.1.24.2/24 on R3 and 10.1.24.4/24 on R4 under the fast Ethernet interfaces. Layer 3 connectivity should be established. The cost of the links under ISIS should be raised to "highest value"under the new interfaces. The link should not send DIS packets on either side.

On R2

```
R2(config)#int f0/0
R2(config-if)#ip add 10.1.24.2 255.255.255.0

R2(config-if)#ip router isis
R2(config-if)#isis metric 200
R2(config-if)#isis circuit-type level-2-only
R2(config-if)#isis network point-to-point
R2(config-if)#NO shut
```

Point to Point mode prevents the DIS process from taking place.

On R4

```
R4(config)#int f0/0
R4(config-if)#ip add 10.1.24.4 255.255.255.0

R4(config-if)#ip router isis
R4(config-if)#isis metric 200
R4(config-if)#isis network point-to-point
R4(config-if)#NO shut

R4#sh isis n

System Id      Type Interface IP Address      State Holdtime Circuit Id
R2             L2   Fa0/0     10.1.24.2       28    02
R3             L2   Vi2       10.1.34.3       UP    27         03
```

On R4

```
R4#Sh ip route isis

i L2 1.0.0.0/8 [115/220] via 10.1.24.2, FastEthernet0/0
i L2 2.0.0.0/8 [115/30] via 10.1.34.3, Virtual-Access1
i L2 3.0.0.0/8 [115/20] via 10.1.34.3, Virtual-Access1
     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
i L2    10.1.12.0/24 [115/210] via 10.1.24.2, FastEthernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
i L2    10.1.23.0/24 [115/20] via 10.1.34.3, Virtual-Access1
i*L2 0.0.0.0/0 [115/10] via 10.1.34.3, Virtual-Access1
```

Task 5

The routes should continue to be advertised. However, R4 should not prefer ANY routes from R1 through R3. No static/policy routing or changes in Admin Distance are permitted. The default route has been perceived as being advertised to R2 thru R4 along the new path. This default route must be removed from R3 to complete the change

On R4

```
R4#Show isis database R3.00-00 | b IS-IS

IS-IS Level-2 LSP R3.00-00
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R3.00-00              0x00000018   0x9DDD        1066              0/0/0
```

On R3

```
R3(config-router)#set-overload-bit
```

On R4

```
R4#Sh isis database R3.00-00 | b IS-IS

IS-IS Level-2 LSP R3.00-00
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R3.00-00              0x0000000A   0xBDC7        1149              0/0/1
```

On R2

```
R2#Sh isis database level-2 R3.00-00 | b IS-IS

IS-IS Level-2 LSP R3.00-00
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R3.00-00              0x0000000A   0xBDC7        1042              0/0/1
```

On R4

```
R4#Sh ip route isis
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
i L2 1.0.0.0/8 [115/60] via 10.1.24.2, FastEthernet0/0
i L2 2.0.0.0/8 [115/30] via 10.1.34.3, Virtual-Access1
i L2 3.0.0.0/8 [115/20] via 10.1.34.3, Virtual-Access1
     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
i L2    10.1.12.0/24 [115/50] via 10.1.24.2, FastEthernet0/0
i L2    10.1.23.0/24 [115/20] via 10.1.34.3, Virtual-Access1
i*L2 0.0.0.0/0 [115/10] via 10.1.34.3, Virtual-Access1
```

On R2

```
R2#Sh ip route isis

i L2 1.0.0.0/8 [115/20] via 10.1.12.1, Serial0/0.21
i L1 3.0.0.0/8 [115/20] via 10.1.23.3, Serial0/0.23
i L2 4.0.0.0/8 [115/50] via 10.1.24.4, FastEthernet0/0
     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
i L1    10.1.34.0/24 [115/20] via 10.1.23.3, Serial0/0.23
i L2    10.1.34.4/32 [115/60] via 10.1.24.4, FastEthernet0/0
i*L2 0.0.0.0/0 [115/50] via 10.1.24.4, FastEthernet0/0
```

On R3

```
R3(config-router)#NO default-information originate
```

On R4

```
R4#Sh ip route isis | I L2 0.0.0.0

i*L2 0.0.0.0/0 [115/40] via 10.1.24.2, FastEthernet0/0
```

On R2

```
R2#Sh ip route 0.0.0.0
% Network not in table
```

Task 6- BONUS

DO not allow R1 and R4 to see the specific route. To each others loopback. Connectivity should still be maintained between both loopbacks. No filters or static routes.

On R2

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config)#router isis
R2(config-router)#set-overload-bit suppress interlevel
```

On R2

```
R2#Sh ip route isis

i L2 1.0.0.0/8 [115/20] via 10.1.12.1, Serial0/0.21
i L1 3.0.0.0/8 [115/20] via 10.1.23.3, Serial0/0.23
i L2 4.0.0.0/8 [115/50] via 10.1.24.4, FastEthernet0/0
     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
i L1    10.1.34.0/24 [115/20] via 10.1.23.3, Serial0/0.23
i L2    10.1.34.4/32 [115/60] via 10.1.24.4, FastEthernet0/0
```

On R1

```
R1#Sh ip route isis

i L2 2.0.0.0/8 [115/20] via 10.1.12.2, Serial0/0.12
     10.0.0.0/24 is subnetted, 3 subnets
i L2    10.1.24.0 [115/50] via 10.1.12.2, Serial0/0.12
i L2    10.1.23.0 [115/20] via 10.1.12.2, Serial0/0.12
i*L2 0.0.0.0/0 [115/10] via 10.1.12.2, Serial0/0.12
```

On R4

```
R4#Sh ip route isis

i L2 2.0.0.0/8 [115/30] via 10.1.34.3, Virtual-Access1
i L2 3.0.0.0/8 [115/20] via 10.1.34.3, Virtual-Access1
     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
i L2    10.1.12.0/24 [115/50] via 10.1.24.2, FastEthernet0/0
i L2    10.1.23.0/24 [115/20] via 10.1.34.3, Virtual-Access1
i*L2 0.0.0.0/0 [115/40] via 10.1.24.2, FastEthernet0/0
```

Task 6

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Advanced
CCIE SERVICE PROVIDER
v3.0

www.MicronicsTraining.com

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

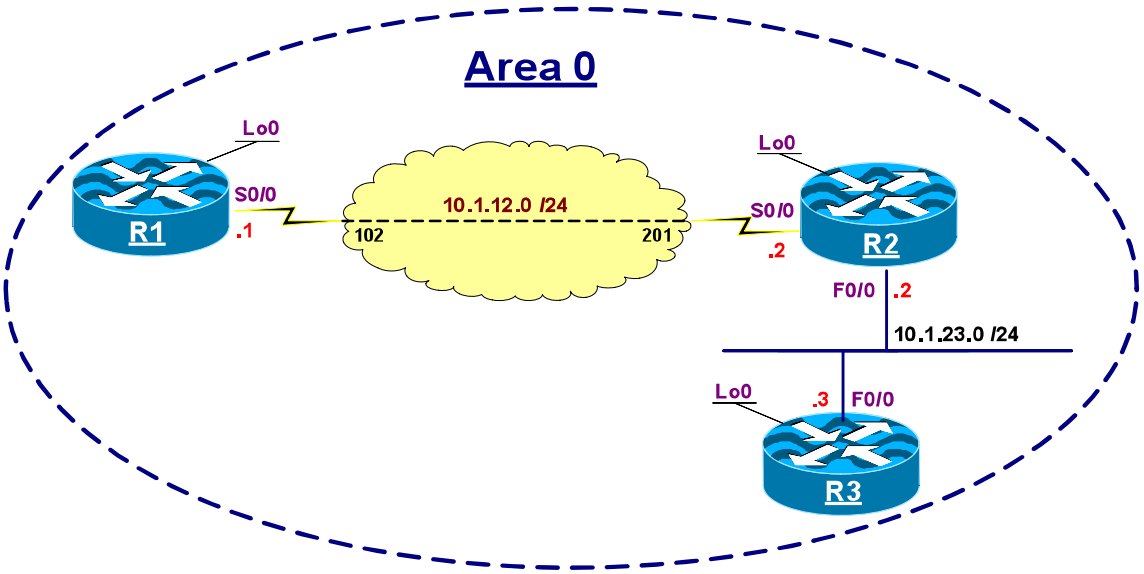
Paul Negron
CCIE #14856
SP

OSPF

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Lab 1
Advertising Networks



Lab Setup:

- Configure the frame-relay connection using the S0/0 interface
- Configure the F0/0 interface of R2 and R3 in VLAN 23
- Use the following IP addressing chart for IP address assignment

IP addressing:

Router	Interface / IP addressing
R1	S0/0 = 10.1.12.1 /24 Lo0 = 1.1.1.1 /8
R2	**S0/0 = 10.1.12.2 /24** **F0/0 = 10.1.23.2 /24** **10.2.2.2 /24 Secondary** **Lo0 = 2.2.2.2 /8**
R3	F0/0 = 10.1.23.3 /24 10.3.3.3 /24 Secondary Lo0 = 3.3.3.3 /8

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 1

Configure OSPF on these routers and run every interface (this includes the secondary interfaces) of these routers in area 0. Do not use the "network" command to accomplish this task. The loopback interface/s should be advertised with their correct mask. There should **not** be a DR election on the frame-relay network. **Do not** use point-to-multipoint network type. If this configuration is performed successfully, you should **only** see intra area routes.

On R1:

```
R1(config)#int s0/0
R1(config-if)#ip ospf network point-to-point
R1(config-if)#ip ospf 1 area 0

R1(config)#int lo0
R1(config-if)#ip ospf network point-to-point
R1(config-if)#ip ospf 1 area 0
```

On R2:

```
R2(config)#int lo0
R2(config-if)#ip ospf network point-to-point
R2(config-if)#ip ospf 1 area 0

R2(config)#int f0/0
R2(config-if)#ip ospf 1 area 0

R2(config)#int s0/0
R2(config-if)#ip ospf network point-to-point
R2(config-if)#ip ospf 1 area 0
```

On R3:

```
R3(config)#int lo0
R3(config-if)#ip ospf 1 area 0
R3(config-if)#ip ospf network point-to-point

R3(config)#int f0/0
R3(config-if)#ip ospf 1 area 0
```

To verify the configuration:

On R1:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1#Show ip route ospf | inc O

O    2.0.0.0/8 [110/65] via 10.1.12.2, 00:01:28, Serial0/0
O    3.0.0.0/8 [110/66] via 10.1.12.2, 00:00:15, Serial0/0
O       10.3.3.0 [110/66] via 10.1.12.2, 00:00:15, Serial0/0
O       10.2.2.0 [110/65] via 10.1.12.2, 00:01:28, Serial0/0
O       10.1.23.0 [110/65] via 10.1.12.2, 00:00:25, Serial0/0
```

Note: The secondary IP addresses are also advertised.

Task 2

Configure R2 and R3 such that the secondary IP addresses are **not** advertised; **Do not** use access-list or prefix-lists. You should use minimum number of commands to accomplish this task.

On R2: and R3

The "secondaries none" keyword will not advertise the secondary IP address/es.

```
Rx(config)#int f0/0
Rx(config-if)#ip ospf 1 area 0 secondaries none
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf | inc O

O    2.0.0.0/8 [110/65] via 10.1.12.2, 00:06:00, Serial0/0
O    3.0.0.0/8 [110/66] via 10.1.12.2, 00:04:47, Serial0/0
O       10.1.23.0 [110/65] via 10.1.12.2, 00:04:57, Serial0/0
```

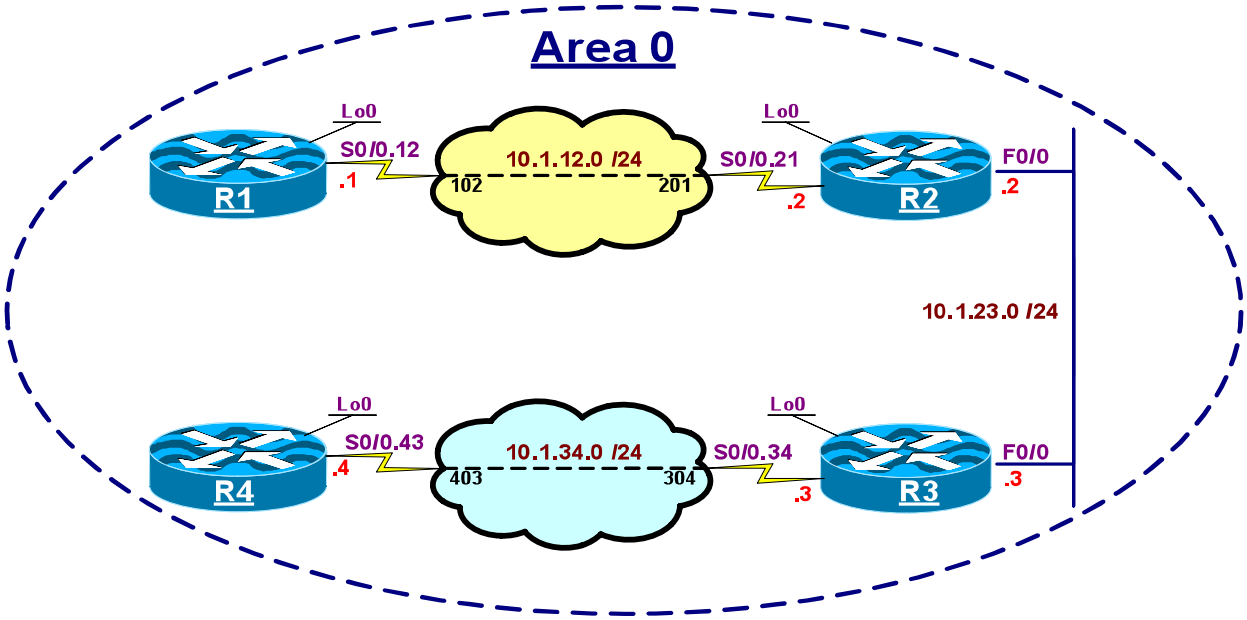
Note: The secondary prefixes are no longer advertised.

Task 3

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 2
Optimization of OSPF and Adjusting Timers



Lab Setup:

- R2 and R3's F0/0 interface should be configured in VLAN 23.
- R1 and R2, R3 and R4 should be configured in a frame-relay point-to-point manner.

IP Addressing

Router	Interface	IP address
R1	Lo0 F/R interface	1.1.1.1 /8 10.1.12.1 /24
R2	**Lo0** **F/R interface** **F0/0**	**2.2.2.2 /8** **10.1.12.2 /24** **10.1.23.2 /24**
R3	Lo0 F/R interface F0/0	3.3.3.3 /8 10.1.34.3 /24 10.1.23.3 /24
R4	**Lo0** **F/R interface**	**4.4.4.4 /8** **10.1.34.4 /24**

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 1

Configure OSPF on all routers and advertise their directly connected networks in area 0. These routers should be configured such that their existing and future interfaces are in area 0.

The following runs OSPF area 0 on all existing and future directly connected interface/s:

On R4:

```
R4(config)#router ospf 1
R4(config-router)#netw 0.0.0.0 0.0.0.0 are 0
```

On R3:

```
R3(config-if)#router ospf 1
R3(config-router)#network 0.0.0.0 0.0.0.0 are 0
```

On R2:

```
R2(config-if)#router ospf 1
R2(config-router)#netw 0.0.0.0 0.0.0.0 are 0
```

On R1:

```
R1(config-if)#router ospf 1
R1(config-router)#netw 0.0.0.0 0.0.0.0 are 0
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf

     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/65] via 10.1.12.2, 00:00:59, Serial0/0.12
     3.0.0.0/32 is subnetted, 1 subnets
O       3.3.3.3 [110/66] via 10.1.12.2, 00:00:09, Serial0/0.12
     4.0.0.0/32 is subnetted, 1 subnets
O       4.4.4.4 [110/130] via 10.1.12.2, 00:00:09, Serial0/0.12
     10.0.0.0/24 is subnetted, 3 subnets
O       10.1.23.0 [110/65] via 10.1.12.2, 00:00:19, Serial0/0.12
O       10.1.34.0 [110/129] via 10.1.12.2, 00:00:09, Serial0/0.12
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 2

R4 is getting flooded with LSA type 6 packets; ensure that R4 does not generate a syslog message for these packets.

Cisco routers do not support LSA type 6 packets and each time an OSPF router receives an MOSPF (LSA type 6) packet it sends a syslog message. If the routers receive many LSA type 6 packets they will generate a large number of syslog messages. This feature should be disabled to prevent this from occurring.

On R4:

```
R4(config)#router ospf 1
R4(config-router)#ignore lsa mospf
```

Task 3

To ensure fast detection of a neighbor being down, configure R2 and R3 to send their hellos every 250ms with a hold time of one second for their ethernet link.

On R2:

```
R2(config)#int f0/0
R2(config-if)#ip ospf dead-interval minimal hello-multiplier 4
```

On R3:

```
R3(config-if)#int f0/0
R3(config-if)#ip ospf dead-interval minimal hello-multiplier 4
```

The dead interval is advertised in OSPF hello packets. The values of this parameter must be the same for two routers in order for them to form a neighbor adjacency. By specifying the "minimal" and "hello-multiplier" keywords with a multiplier value, you are enabling OSPF fast hello packets. The "minimal" keyword sets the dead interval to 1 second and the "hello-multiplier" value sets the number of hello packets sent during that 1 second. Since one second is 1000ms, if 1000 is divided by 250 (the requested hello interval) the result should be the hello multiplier.

To verify the configuration:

On R3:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3#Show ip ospf interface f0/0 | inc Timer

 Timer intervals configured, Hello 250 msec, Dead 1, Wait 1, Retransmit 5
```

Task 4

Ensure that these routers lookup DNS names for use in all OSPF show commands, test this task to ensure proper operation. Since there are no DNS servers in this lab, you should use the routers for that purpose.

To test the OSPF "show" commands before implementing this feature:

```
R1#Show ip ospf database router | inc Advert

   Advertising Router: 1.1.1.1
   Advertising Router: 2.2.2.2
   Advertising Router: 3.3.3.3
   Advertising Router: 4.4.4.4
```

To configure the task:

On All Routers

```
Rx(config)#ip ospf name-lookup

Rx(config)#ip host R1 1.1.1.1
Rx(config)#ip host R2 2.2.2.2
Rx(config)#ip host R3 3.3.3.3
Rx(config)#ip host R4 4.4.4.4
```

To verify the configuration:

On R1:

```
R1#Show ip ospf database router | inc Advert

   Advertising Router: R1
   Advertising Router: R2
   Advertising Router: R3
   Advertising Router: R4
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 5

Configure R2 such that if it does not receive an acknowledgment from R3 for a given LSA, it waits for 10 seconds before it resends that given LSA.

On R2:

```
R2(config)#int f0/0
R2(config-if)#ip ospf retransmit-interval 10
```

When an OSPF enabled router sends an LSA to its neighbor, it keeps the LSA until it receives an ACK from that given neighbor. If the retransmission timer expires and the router receives no ACKs, the router will resend that LSA. The default timer is set to 5 seconds, and the range is 1 – 65535.

To verify the configuration:

On R2:

```
R2#Sh ip ospf inter f0/0 | inc Timer
  Timer intervals configured, Hello 250 msec, Dead 1, Wait 1, Retransmit 10
```

Task 6

Configure R2 such that it limits the number of non-self generated LSAs that an OSPF routing process can keep in the OSPF LSDB to 900. Additional LSAs should generate a console message.

On R2:

```
R2(config-if)#router ospf 1
R2(config-router)#max-lsa 900 warning-only
```

By default, the number of non-self-generated LSAs that an OSPF routing process can keep in the database is not limited. To limit this number, we can use the "max-lsa" command in the router configuration mode. This command contains the following fields:

```
Max-lsa maximum-number [threshold-percentage] [warning-only] [ignore-time
minutes] [ignore-count #] [reset-time minutes]
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

- **Maximum-number** – The maximum number of non-self-generated LSAs that an OSPF router can keep in the OSPF database.
- **Threshold-percentage** – The percentage of the maximum LSA number, as specified by the maximum-number, at which a warning message is logged. The default is 75.
- **Warning-only** – This specifies that only a warning message be sent when the maximum limit for LSAs is exceeded.
- **Ignore-time *minutes*** – This value specifies the time, in minutes, to ignore all neighbors after the maximum limit of LSAs has been exceeded. The default is 5 minutes.
- **Ignore-count count-number** – Specifies the number of times the OSPF process can consecutively be placed into the ignore state. The default is 5 times. The router cannot exceed this number.
- **Reset-time *minutes*** - This value specifies the time, in minutes, after which the ignore count is reset to zero. The default is 10 minutes.

To verify the configuration:

On R2:

```
R2#Show ip ospf 1 | s maximum number

Maximum number of non-self-generated LSA allowed 900 (warning-only)
    Current number of non-self-generated LSA 4
    Threshold for warning message 75%
```

Task 7

R3 and R4 should exchange hellos every 15 seconds with a dead interval of 60 seconds. Do **not** use the "**ip ospf dead-interval**" command to accomplish this task.

On R3:

```
R3(config-router)#int s0/0.34
R3(config-subif)#ip ospf hello-interval 15
```

On R4:

```
R4(config-router)#int s0/0.43
R4(config-subif)#ip ospf hello-interval 15
```

Once the hello-interval is set, OSPF process will set the dead-interval to be four times the hello-interval. The default value for the hello timer is as follows:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

- **On broadcast network types, it is set to 10 seconds.**
- **On non-broadcast networks or point to multipoint network types its set to 30 seconds.**

To verify the configuration:

On R4:

```
R4#Show ip ospf int s0/0.43 | inc Timer

  Timer intervals configured, Hello 15, Dead 60, Wait 60, Retransmit 5
```

Task 8

Configure R1 and R2 such that when there is a topology change in area 0 for LSAs type ones and twos, the entire SPF is NOT recomputed. This should **only** occur for the affected part/s of the tree.

When any change occurs for LSAs type-1s and Type-2s, the entire SPF is recomputed. The incremental SPF allows the IOS to only recompute the affected part/s of the tree. This can result in faster OSPF convergence, and save resources.

On R1: & R2

```
Rx(config)#router ospf 1
Rx(config-router)#ispf
```

To verify the configuration:

On R1:

```
R1#Show ip ospf | inc Incremental

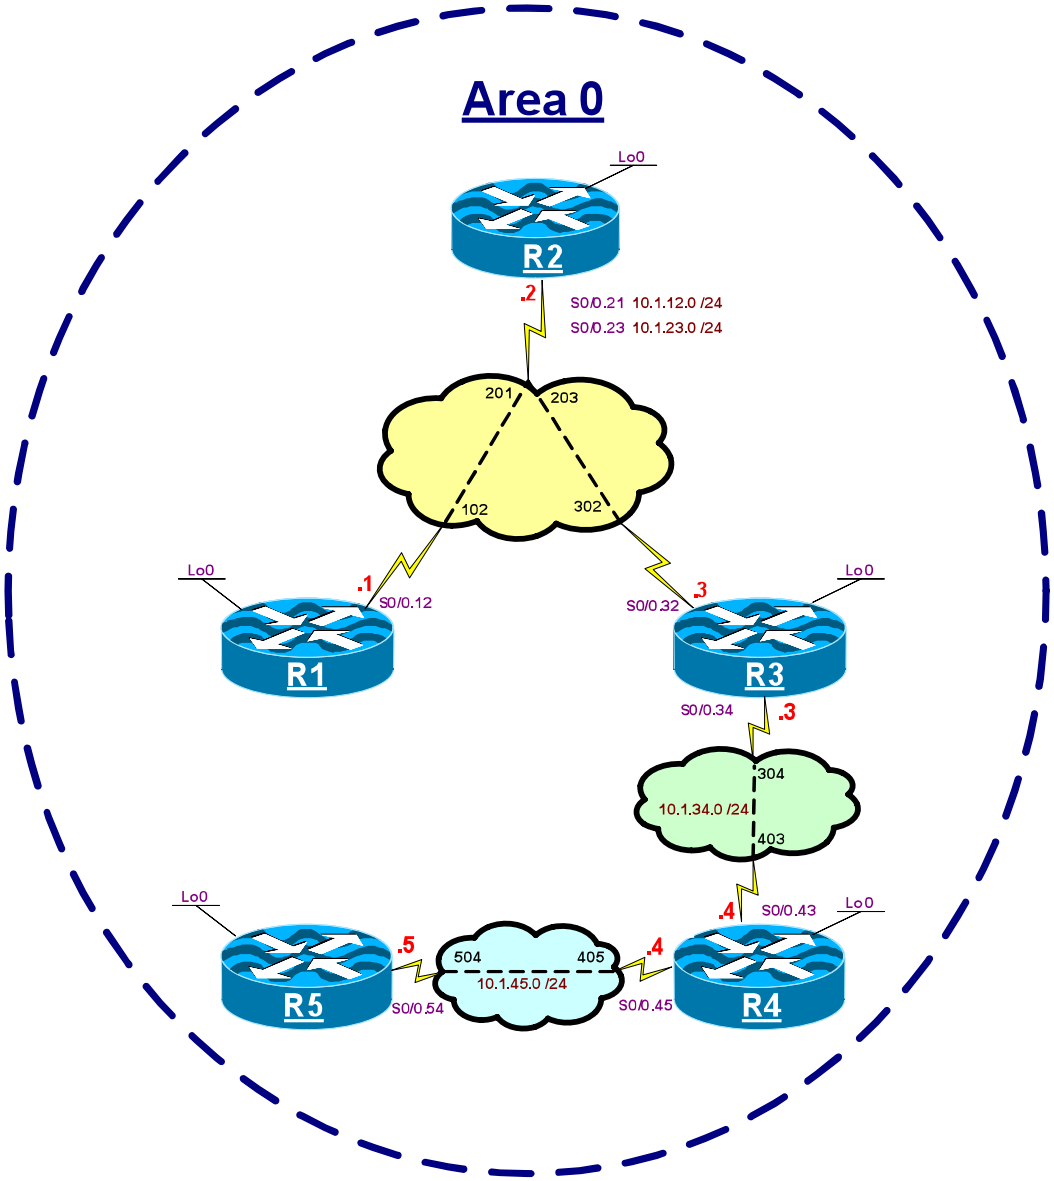
 Incremental-SPF enabled
```

Task 9

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 3 – OSPF Authentication



Lab Setup:

- Configure all frame-relay connections in a point-to-point manner.
- Use the IP addressing scheme below for IP addressing assignment.

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

IP Addressing scheme:

Routers	Interface / IP Address	Connecting to:
R1	S0/0.12 – 10.1.12.1 /24 Loopback 0 – 1.1.1.1 /24	R2
R2	S0/0.21 – 10.1.12.2 /24 S0/0.23 – 10.1.23.2 /24 Loopback 0 – 2.2.2.2 /24	R1 R3
R3	S0/0.32 – 10.1.23.3 /24 S0/0.34 – 10.1.34.3 /24 Loopback 0 – 3.3.3.3 /24	R2 R4
R4	S0/0.43 – 10.1.34.4 /24 S0/0.45 – 10.1.45.4 /24 Loopback 0 – 4.4.4.4 /24	R3 R5
R5	S0/0.54 – 10.1.45.5 /24 Loopback 0 – 5.5.5.5 /24	R4

Task 1

Configure the frame-relay interface/s and the loopback interface/s of all routers in area 0, and ensure that the loopback interfaces are advertised with their correct mask. The router-id of the routers in this area should be based on their loopback 0 interfaces' IP address.

On R1:

```
R1(config)#int lo0
R1(config-if)#ip ospf net point-to-point

R1(config)#router ospf 1
R1(config-router)#router-id 1.1.1.1

R1(config-router)#netw 1.1.1.1 0.0.0.0 are 0
R1(config-router)#netw 10.1.12.1 0.0.0.0 are 0
```

On R2:

```
R2(config)#int lo0
R2(config-if)#ip ospf network point-to-point

R2(config-if)#router ospf 1
R2(config-router)#router-id 2.2.2.2

R2(config-router)#netw 2.2.2.2 0.0.0.0 area 0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config-router)#netw 10.1.12.2 0.0.0.0 area 0
R2(config-router)#netw 10.1.23.2 0.0.0.0 area 0
```

On R3:

```
R3(config)#int lo0
R3(config-if)#ip ospf network point-to-point

R3(config-if)#router ospf 1
R3(config-router)#router-id 3.3.3.3

R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.23.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.34.3 0.0.0.0 area 0
```

On R4:

```
R4(config)#int lo0
R4(config-if)#ip ospf network point-to-point

R4(config-if)#router ospf 1
R4(config-router)#router-id 4.4.4.4

R4(config-router)#netw 4.4.4.4 0.0.0.0 area 0
R4(config-router)#netw 10.1.45.4 0.0.0.0 area 0
R4(config-router)#netw 10.1.34.4 0.0.0.0 area 0
```

On R5:

```
R5(config)#int lo0
R5(config-if)#ip ospf network point-to-point

R5(config-if)#router ospf 1
R5(config-router)#router-id 5.5.5.5

R5(config-router)#netw 10.1.45.5 0.0.0.0 area 0
R5(config-router)#netw 5.5.5.5 0.0.0.0 area 0
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf | inc O

O       2.2.2.0 [110/65] via 10.1.12.2, 00:04:08, Serial0/0.12
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O       3.3.3.0 [110/129] via 10.1.12.2, 00:03:16, Serial0/0.12
O       4.4.4.0 [110/193] via 10.1.12.2, 00:02:28, Serial0/0.12
O       5.5.5.0 [110/257] via 10.1.12.2, 00:00:42, Serial0/0.12
O       10.1.23.0 [110/128] via 10.1.12.2, 00:03:58, Serial0/0.12
O       10.1.45.0 [110/256] via 10.1.12.2, 00:00:52, Serial0/0.12
O       10.1.34.0 [110/192] via 10.1.12.2, 00:03:16, Serial0/0.12
```

Task 2

Configure plain text authentication on all the frame-relay links in this area. You should use a router configuration command as part of the solution to this task. Use "Cisco" as the password for this authentication.

OSPF supports two types of authentication, plain text (64 bit password) and MD5 (which consists of a key ID and 128 bit password). In OSPF, authentication must be enabled and then applied.
In OSPF, enabling authentication can be configured in two different ways; one way to enable OSPF authentication is to configure it in the router configuration mode. In which case authentication is enabled globally on all OSPF enabled interfaces in the specified area. The second way to enable authentication is to configure it directly on the interface for which authentication is required.

On R1:

```
R1(config)#router ospf 1
R1(config-router)#area 0 authentication

R1(config-router)#int s0/0.12
R1(config-subif)#ip ospf authentication-key Cisco
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#area 0 authentication

R2(config-router)#int s0/0.21
R2(config-subif)#ip ospf authentication-key Cisco

R2(config-subif)#int s0/0.23
R2(config-subif)#ip ospf authentication-key Cisco
```

To verify the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2:

```
R2#Show ip ospf interface s0/0.21 | inc authentication

  Simple password authentication enabled
```

Note: The output of the above "show" command verifies that a simple password authentication is enabled and applied to this interface.

```
R2#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
1.1.1.1           0   FULL/  -        00:00:38    10.1.12.1       Serial0/0.21

R2#show ip route ospf | inc O

O       1.1.1.0 [110/65] via 10.1.12.1, 00:05:00, Serial0/0.21
```

Let's configure R3 and R4:

On R3:

```
R3(config)#router ospf 1
R3(config-router)#area 0 authentication

R3(config)#int s0/0.32
R3(config-subif)#ip ospf authentication-key Cisco

R3(config)#int s0/0.34
R3(config-subif)#ip ospf authentication-key Cisco
```

To verify the configuration:

On R3:

```
R3#Show ip route ospf | inc O

O       1.1.1.0 [110/129] via 10.1.23.2, 00:01:36, Serial0/0.32
O       2.2.2.0 [110/65] via 10.1.23.2, 00:01:36, Serial0/0.32
O       10.1.12.0 [110/128] via 10.1.23.2, 00:01:36, Serial0/0.32
```

On R4:

```
R4(config-subif)#router ospf 1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config-router)#area 0 authentication

R4(config)#int s0/0.43
R4(config-subif)#ip ospf authentication-key Cisco

R4(config-subif)#int s0/0.45
R4(config-subif)#ip ospf authentication-key Cisco
```

To verify the configuration:

On R4:

```
R4#Show ip route ospf | inc O

O       1.1.1.0 [110/193] via 10.1.34.3, 00:00:21, Serial0/0.43
O       2.2.2.0 [110/129] via 10.1.34.3, 00:00:21, Serial0/0.43
O       3.3.3.0 [110/65] via 10.1.34.3, 00:00:21, Serial0/0.43
O       10.1.12.0 [110/192] via 10.1.34.3, 00:00:21, Serial0/0.43
O       10.1.23.0 [110/128] via 10.1.34.3, 00:00:21, Serial0/0.43
```

Let's configure R5:

On R5:

```
R5(config)#router ospf 1
R5(config-router)#area 0 authentication

R5(config-router)#int s0/0.54
R5(config-subif)#ip ospf authentication-key Cisco
```

To verify the configuration:

On R5:

```
R5#Show ip route ospf | inc O

O       1.1.1.0 [110/257] via 10.1.45.4, 00:00:44, Serial0/0.54
O       2.2.2.0 [110/193] via 10.1.45.4, 00:00:44, Serial0/0.54
O       3.3.3.0 [110/129] via 10.1.45.4, 00:00:44, Serial0/0.54
O       4.4.4.0 [110/65] via 10.1.45.4, 00:00:44, Serial0/0.54
O       10.1.12.0 [110/256] via 10.1.45.4, 00:00:44, Serial0/0.54
O       10.1.23.0 [110/192] via 10.1.45.4, 00:00:44, Serial0/0.54
O       10.1.34.0 [110/128] via 10.1.45.4, 00:00:44, Serial0/0.54
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 3

Remove the authentication configuration from the previous task and ensure that every router sees every route advertised in area 0.

On All Routers

```
Rx(config)#router ospf 1
Rx(config-router)#No area 0 authentication
```

On R1:

```
R1(config)#int s0/0.12
R1(config-subif)#No ip ospf authentication-key Cisco
```

On R2:

```
R2(config-subif)#int s0/0.21
R2(config-subif)#No ip ospf authentication-key Cisco

R2(config-subif)#int s0/0.23
R2(config-subif)#No ip ospf authentication-key Cisco
```

On R3:

```
R3(config-router)#int s0/0.32
R3(config-subif)#No ip ospf authentication-key Cisco

R3(config-subif)#int s0/0.34
R3(config-subif)#No ip ospf authentication-key Cisco
```

On R4:

```
R4(config)#int s0/0.43
R4(config-subif)#NO ip ospf authentication-key Cisco

R4(config)#int s0/0.45
R4(config-subif)#No ip ospf authentication-key Cisco
```

On R5:

```
R5(config)#int s0/0.54
R5(config-subif)#No ip ospf authentication-key Cisco
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the configuration:

On R1:

```
R1#Show ip route ospf | inc O

O        2.2.2.0 [110/65] via 10.1.12.2, 00:00:10, Serial0/0.12
O        3.3.3.0 [110/129] via 10.1.12.2, 00:00:10, Serial0/0.12
O        4.4.4.0 [110/193] via 10.1.12.2, 00:00:10, Serial0/0.12
O        5.5.5.0 [110/257] via 10.1.12.2, 00:00:10, Serial0/0.12
O        10.1.23.0 [110/128] via 10.1.12.2, 00:00:10, Serial0/0.12
O        10.1.45.0 [110/256] via 10.1.12.2, 00:00:10, Serial0/0.12
O        10.1.34.0 [110/192] via 10.1.12.2, 00:00:10, Serial0/0.12
```

Task 4

Configure **MD5** authentication on all the frame-relay links in this area. You should use a router configuration command as part of the solution to this task. Use "Cisco" as the password for this authentication.

The following command enables MD5 authentication on the routers using the router configuration mode:

On All Routers

```
Rx(config)#router ospf 1
Rx(config-router)#area 0 authentication message-digest
```

On R1:

```
R1(config)#int s0/0.12
R1(config-subif)#ip ospf message-digest-key 1 MD5 Cisco
```

On R2:

```
R2(config)#int s0/0.21
R2(config-subif)#ip ospf message-digest-key 1 MD5 Cisco

R2(config-subif)#int s0/0.23
R2(config-subif)#ip ospf message-digest-key 1 MD5 Cisco
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the configuration:

On R2:

```
R2#Show ip ospf interface S0/0.21 | b Message

  Message digest authentication enabled
    Youngest key id is 1
```

Note: the output of the above "show" command reveals that MD5 authentication is enabled and applied and the key 1 is in use.

```
R2#Show ip route ospf

     1.0.0.0/24 is subnetted, 1 subnets
O       1.1.1.0 [110/65] via 10.1.12.1, 00:01:50, Serial0/0.21
```

On R3:

```
R3(config)#int s0/0.32
R3(config-subif)#ip ospf message-digest-key 1 md5 Cisco

R3(config)#int s0/0.34
R3(config-subif)#ip ospf message-digest-key 1 md5 Cisco
```

To verify the configuration:

On R3:

```
R3#Show ip route ospf | inc O

O       1.1.1.0 [110/129] via 10.1.23.2, 00:00:11, Serial0/0.32
O       2.2.2.0 [110/65] via 10.1.23.2, 00:00:11, Serial0/0.32
O       10.1.12.0 [110/128] via 10.1.23.2, 00:00:11, Serial0/0.32
```

On R4:

```
R4(config)#int s0/0.45
R4(config-subif)#ip ospf message-digest-key 1 MD5 Cisco

R4(config)#int s0/0.43
R4(config-subif)#ip ospf message-digest-key 1 MD5 Cisco
```

To verify the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4:

```
R4#Show ip route ospf | inc O

O       1.1.1.0 [110/193] via 10.1.34.3, 00:00:21, Serial0/0.43
O       2.2.2.0 [110/129] via 10.1.34.3, 00:00:21, Serial0/0.43
O       3.3.3.0 [110/65] via 10.1.34.3, 00:00:21, Serial0/0.43
O       10.1.12.0 [110/192] via 10.1.34.3, 00:00:21, Serial0/0.43
O       10.1.23.0 [110/128] via 10.1.34.3, 00:00:21, Serial0/0.43
```

On R5:

```
R5(config)#int s0/0.54
R5(config-subif)#ip ospf message-digest-key 1 md5 Cisco
```

To verify the configuration:

On R5:

```
R5#Show ip route ospf | inc O

O       1.1.1.0 [110/257] via 10.1.45.4, 00:00:42, Serial0/0.54
O       2.2.2.0 [110/193] via 10.1.45.4, 00:00:42, Serial0/0.54
O       3.3.3.0 [110/129] via 10.1.45.4, 00:00:42, Serial0/0.54
O       4.4.4.0 [110/65] via 10.1.45.4, 00:00:42, Serial0/0.54
O       10.1.12.0 [110/256] via 10.1.45.4, 00:00:42, Serial0/0.54
O       10.1.23.0 [110/192] via 10.1.45.4, 00:00:42, Serial0/0.54
O       10.1.34.0 [110/128] via 10.1.45.4, 00:00:42, Serial0/0.54
```

Task 5

Remove the authentication configuration from the previous task and ensure that every router sees every route advertised in area 0.

On All Routers:

```
Rx(config)#router ospf 1
Rx(config-router)#No area 0 authentication message-digest
```

On R1:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config)#int s0/0.12
R1(config-if)#No ip ospf message-digest-key 1 md5 Cisco
```

On R2:

```
R2(config)#int s0/0.21
R2(config-subif)#No ip ospf message-digest-key 1 md5 Cisco

R2(config-subif)#int s0/0.23
R2(config-subif)#No ip ospf message-digest-key 1 md5 Cisco
```

On R3:

```
R3(config)#int s0/0.32
R3(config-subif)#No ip ospf message-digest-key 1 md5 Cisco

R3(config)#int s0/0.34
R3(config-subif)#No ip ospf message-digest-key 1 md5 Cisco
```

On R4:

```
R4(config)#int s0/0.43
R4(config-subif)#NO ip ospf message-digest-key 1 md5 Cisco

R4(config)#int s0/0.45
R4(config-subif)#No ip ospf message-digest-key 1 md5 Cisco
```

On R5:

```
R5(config)#int s0/0.54
R5(config-subif)#No ip ospf message-digest-key 1 md5 Cisco
```

To verify the configuration:

On R5:

```
R5#Show ip route ospf | inc O

O       1.1.1.0 [110/257] via 10.1.45.4, 00:09:13, Serial0/0.54
O       2.2.2.0 [110/193] via 10.1.45.4, 00:09:13, Serial0/0.54
O       3.3.3.0 [110/129] via 10.1.45.4, 00:09:13, Serial0/0.54
O       4.4.4.0 [110/65] via 10.1.45.4, 00:09:13, Serial0/0.54
O       10.1.12.0 [110/256] via 10.1.45.4, 00:09:13, Serial0/0.54
O       10.1.23.0 [110/192] via 10.1.45.4, 00:09:13, Serial0/0.54
O       10.1.34.0 [110/128] via 10.1.45.4, 00:09:13, Serial0/0.54
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 6

Configure MD5 authentication on the frame-relay link connecting R1 to R2, you should use a router configuration command as part of the solution to this task. The password should be "ccie".

On Both Routers:

```
Rx(config)#router ospf 1
Rx(config-router)#area 0 authentication message-digest
```

On R1:

```
R1(config)#int s0/0.12
R1(config-subif)#ip ospf message-digest-key 1 md5 ccie
```

On R2:

```
R2(config)#int s0/0.21
R2(config-subif)#ip ospf message-digest-key 1 md5 ccie
```

To verify the configuration:

On R2:

```
R2#Show ip route ospf | inc O

O        1.1.1.0 [110/65] via 10.1.12.1, 00:00:43, Serial0/0.21
```

Note: Because authentication is enabled in the router configuration mode, it is applied to every interface in area 0. Therefore, every router in area 0 must have authentication enabled. Since R3 does not have authentication enabled, R2 will drop its adjacency with R3 when the dead interval expires, therefore, they will not exchange updates. You should see the following console message on R2:

```
%OSPF-5-ADJCHG: Process 1, Nbr 3.3.3.3 on Serial0/0.23 from FULL to DOWN,
Neighbor Down: Dead timer expired
```

Task 7

Configure these routers such that every router has every prefix advertised in this topology in their routing table and link state database. **Do not** remove the "area 0 authentication message-digest" router configuration commands from the previous task to accomplish this task.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3, R4 and R5:

```
Rx(config)#router ospf 1
Rx(config-router)#area 0 authentication message-digest
```

To verify the configuration:

On R2:

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 3.3.3.3 on Serial0/0.23 from LOADING to
FULL, Loading Done

R2#Show ip route ospf | inc O

O       1.1.1.0 [110/65] via 10.1.12.1, 00:01:01, Serial0/0.21
O       3.3.3.0 [110/65] via 10.1.23.3, 00:01:01, Serial0/0.23
O       4.4.4.0 [110/129] via 10.1.23.3, 00:01:01, Serial0/0.23
O       5.5.5.0 [110/193] via 10.1.23.3, 00:01:01, Serial0/0.23
O       10.1.45.0 [110/192] via 10.1.23.3, 00:01:01, Serial0/0.23
O       10.1.34.0 [110/128] via 10.1.23.3, 00:01:01, Serial0/0.23
```

Note: Once the authentication is enabled on the other routers, they will form adjacency and exchange routes.

Task 8

Remove the solution from the previous task, and reconfigure R2 such that every router has every prefix advertised in this topology in their routing table and link state database. **Do not** remove the authentication that is applied to the link between R1 and R2.

On R3, R4 and R5:

```
Rx(config)#router ospf 1
Rx(config-router)#No area 0 authentication message-digest
```

To verify the configuration:

On R2:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2#Show ip route ospf | inc O

O       1.1.1.0 [110/65] via 10.1.12.1, 00:00:06, Serial0/0.21
```

To configure this task, we must disable authentication on R2 on its interface facing R3 using the "IP OSPF authentication null" interface configuration command. Meaning that there is no need to have authentication downstream S0/0.23 interface of R2. Therefore, R3, R4 and/or R5 does not need to have authentication enabled.

On R2:

```
R2(config)#int s0/0.23
R2(config-subif)#ip ospf authentication null
```

You should see the following console message on R2:

```
%OSPF-5-ADJCHG: Process 1, Nbr 3.3.3.3 on Serial0/0.23 from LOADING to
FULL, Loading Done
```

To verify the configuration:

On R2:

```
R2#Show ip route ospf | inc O

O       1.1.1.0 [110/65] via 10.1.12.1, 00:03:02, Serial0/0.21
O       3.3.3.0 [110/65] via 10.1.23.3, 00:03:02, Serial0/0.23
O       4.4.4.0 [110/129] via 10.1.23.3, 00:03:02, Serial0/0.23
O       5.5.5.0 [110/193] via 10.1.23.3, 00:03:02, Serial0/0.23
O       10.1.45.0 [110/192] via 10.1.23.3, 00:03:02, Serial0/0.23
O       10.1.34.0 [110/128] via 10.1.23.3, 00:03:02, Serial0/0.23
```

Task 9

Re-configure the authentication password on R1 and R2 to be "CCIERS" without interrupting the links operation.

To see the current configuration:

On R1:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1#Show run int s0/0.12 | inc ip ospf

ip ospf message-digest-key 1 md5 ccie
```

On R2:

```
R2#Show run int s0/0.21 | inc ip ospf

ip ospf message-digest-key 1 md5 ccie

R2#Show ip route ospf | inc O

O       1.1.1.0 [110/65] via 10.1.12.1, 00:06:26, Serial0/0.21
O       3.3.3.0 [110/65] via 10.1.23.3, 00:06:26, Serial0/0.23
O       4.4.4.0 [110/129] via 10.1.23.3, 00:06:26, Serial0/0.23
O       5.5.5.0 [110/193] via 10.1.23.3, 00:06:26, Serial0/0.23
O       10.1.45.0 [110/192] via 10.1.23.3, 00:06:26, Serial0/0.23
O       10.1.34.0 [110/128] via 10.1.23.3, 00:06:26, Serial0/0.23
```

In order to change the password without any interruption to the link, the second key is entered with the required password.

On R1:

```
R1(config)#int s0/0.12
R1(config-subif)#ip ospf message-digest-key 2 md5 CCIERS
```

To verify the configuration:

On R1:

```
R1#Show run int s0/0.12 | inc ip ospf

ip ospf message-digest-key 1 md5 ccie
ip ospf message-digest-key 2 md5 CCIERS

R1#Show ip ospf inter s0/0.12 | b Message

  Message digest authentication enabled
    Youngest key id is 2
    Rollover in progress, 1 neighbor(s) using the old key(s):
      key id 1
```

Note: Even though the second key (key 2) is only configured on R1, R1 and R2 are still authenticating based on the first key (key 1); this is revealed in the second line.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

However, the router knows that the second key is configured (the second line in the above display) and it knows that the rollover is in progress (the third line), but the other end (R2) has not been configured yet.

On R2:

```
R2(config-subif)#int s0/0.21
R2(config-if)#ip ospf message-digest-key 2 md5 CCIERS
```

To verify the configuration:

On R2:

```
R2#Sh ip ospf inter s0/0.21 | b Message

  Message digest authentication enabled
    Youngest key id is 2
```

Note: once R2 is configured, both routers (R1 and R2) will switchover and use the second key for their authentication.

On R1:

```
R1#Show ip ospf interface s0/0.12 | b Message

  Message digest authentication enabled
    Youngest key id is 2
```

Once R1 and R2's key rollover is completed and both routers display the same youngest key without the "rollover in progress" message, we can safely remove the prior key, in this case key id 1. Remember that the newest key is not determined based on the numerically higher value.

On R1:

```
R1#Show run int s0/0.12 | inc ip ospf

 ip ospf message-digest-key 1 md5 ccie
 ip ospf message-digest-key 2 md5 CCIERS

R1(config)#int s0/0.12
R1(config-subif)#no ip ospf message-digest-key 1 md5 ccie
```

On R2:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2#Show run int s0/0.21 | inc ip ospf

 ip ospf message-digest-key 1 md5 ccie
 ip ospf message-digest-key 2 md5 CCIERS

R2(config)#int s0/0.21
R2(config-subif)#no ip ospf message-digest-key 1 md5 ccie
```

Task 10

Re-configure the MD5 authentication on R1 and R2 using "cisco" as the password. You should **noy** use a router configuration mode to accomplish this task.

On R1: and R2

```
Rx(config)#router ospf 1
Rx(config-router)#No area 0 authentication message-digest
```

On R1:

```
R1(config)#int s0/0.12
R1(config-if)#No ip ospf message-digest-key 2 md5 CCIERS
R1(config-subif)#No ip ospf message-digest-key 1 md5 ccie
```

On R2:

```
R2(config)#int s0/0.23
R2(config-subif)#No ip ospf authentication null

R2(config-subif)#int s0/0.21
R2(config-subif)#No ip ospf message-digest-key 2 md5 CCIERS
R2(config-subif)#NO ip ospf message-digest-key 1 md5 ccie
```

Note: The following command enables authentication directly under S0/0.21 interface and not in router configuration mode, this is called per interface authentication.

```
R2(config)#int s0/0.21
R2(config-subif)#ip ospf authentication message-digest
```

On R1:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config)#int s0/0.12
R1(config-subif)#ip ospf authentication message-digest
```

To verify the configuration:

On R1:

```
R1#Show run inter s0/0.12 | inc ip ospf

 ip ospf authentication message-digest
```

On R2:

```
R2#Show run int s0/0.21 | inc ip ospf

 ip ospf authentication message-digest
```

Note: When authentication is enabled directly under a given interface, it enables authentication on that given interface only, therefore, only the neighbor/s through that interface should have authentication enabled. This is called per-interface authentication.

Let's apply the authentication to the link between R1 and R2:

On R1:

```
R1(config)#int s0/0.12
R1(config-subif)#ip ospf message-digest-key 1 md5 cisco
```

On R2:

```
R1(config)#int s0/0.21
R1(config-subif)#ip ospf message-digest-key 1 md5 cisco
```

To test the configuration:

On R1:

```
R1#Show ip route ospf | inc O

O       2.2.2.0 [110/65] via 10.1.12.2, 00:14:36, Serial0/0.12
O       3.3.3.0 [110/129] via 10.1.12.2, 00:14:36, Serial0/0.12
O       4.4.4.0 [110/193] via 10.1.12.2, 00:14:36, Serial0/0.12
O       5.5.5.0 [110/257] via 10.1.12.2, 00:14:36, Serial0/0.12
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O        10.1.23.0 [110/128] via 10.1.12.2, 00:14:36, Serial0/0.12
O        10.1.45.0 [110/256] via 10.1.12.2, 00:14:36, Serial0/0.12
O        10.1.34.0 [110/192] via 10.1.12.2, 00:14:36, Serial0/0.12
```

On R2:

```
R2#Show ip route ospf | inc O

O        1.1.1.0 [110/65] via 10.1.12.1, 00:17:32, Serial0/0.21
O        3.3.3.0 [110/65] via 10.1.23.3, 00:17:32, Serial0/0.23
O        4.4.4.0 [110/129] via 10.1.23.3, 00:17:32, Serial0/0.23
O        5.5.5.0 [110/193] via 10.1.23.3, 00:17:32, Serial0/0.23
O        10.1.45.0 [110/192] via 10.1.23.3, 00:17:32, Serial0/0.23
O        10.1.34.0 [110/128] via 10.1.23.3, 00:17:32, Serial0/0.23
```

Task 11

Re-configure the routers using the following chart; Configure OSPF router-id based on their Lo0 interface; ensure that every router has every prefix advertised in this routing domain in their routing table and link state database:

Router	Interface	Area
R1	S0/0.12 Loopback 0	0 0
R2	S0/0.21 S0/0.23 Loopback 0	0 1 1
R3	S0/0.32 S0/0.34 Loopback 0	1 2 2
R4	S0/0.43 S0/0.45 Loopback 0	2 3 3
R5	S0/0.54 Loopback 0	3 3

On R1:

```
R1(config)#No router ospf 1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config)#int s0/0.12
R1(config-subif)#No ip ospf authentication message-digest
R1(config-subif)#No ip ospf message-digest-key 1 md5 cisco

R1(config)#router ospf 1
R1(config-router)#router-id 1.1.1.1

R1(config-router)#netw 10.1.12.1 0.0.0.0 area 0
R1(config-router)#netw 1.1.1.1 0.0.0.0 area 0
```

On R2:

```
R2(config)#No router ospf 1

R2(config)#int s0/0.21
R2(config-sub)#No ip ospf authentication message-digest
R2(config-subif)#No ip ospf message-digest-key 1 md5 cisco

R2(config)#int s0/0.23
R2(config-subif)#No ip ospf authentication-key Cisco

R2(config)#router ospf 1
R2(config-router)#router-id 2.2.2.2

R2(config-router)#netw 10.1.12.2 0.0.0.0 area 0
R2(config-router)#netw 10.1.23.2 0.0.0.0 area 1
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 1
R2(config-router)#area 1 virtual-link 3.3.3.3
```

On R3:

```
R3(config)#No router ospf 1

R3(config)#int s0/0.32
R3(config-subif)#No ip ospf authentication-key Cisco

R3(config)#int s0/0.34
R3(config-subif)#No ip ospf authentication-key Cisco

R3(config)#router ospf 1
R3(config-router)#router-id 3.3.3.3

R3(config-router)#netw 10.1.23.3 0.0.0.0 area 1
R3(config-router)#netw 10.1.34.3 0.0.0.0 area 2
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 2
R3(config-router)#area 1 virtual-link 2.2.2.2
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-router)#area 2 virtual-link 4.4.4.4
```

On R4:

```
R4(config)#No router ospf 1

R4(config)#int s0/0.43
R4(config-subif)#No ip ospf authentication-key Cisco

R4(config)#int s0/0.45
R4(config-subif)#No ip ospf authentication-key Cisco

R4(config)#router ospf 1
R4(config-router)#router-id 4.4.4.4

R4(config-router)#netw 10.1.34.4 0.0.0.0 area 2
R4(config-router)#netw 4.4.4.4 0.0.0.0 area 3
R4(config-router)#netw 10.1.45.4 0.0.0.0 area 3
R4(config-router)#area 2 virtual-link 3.3.3.3
```

On R5:

```
R5(config)#No router ospf 1

R5(config)#int s0/0.54
R5(config-subif)#No ip ospf authentication-key Cisco

R5(config)#router ospf 1
R5(config-router)#router-id 5.5.5.5

R5(config-router)#netw 10.1.45.5 0.0.0.0 area 3
R5(config-router)#netw 5.5.5.5 0.0.0.0 area 3
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf | inc O

O IA    2.2.2.0 [110/65] via 10.1.12.2, 00:00:32, Serial0/0.12
O IA    3.3.3.0 [110/129] via 10.1.12.2, 00:00:32, Serial0/0.12
O IA    4.4.4.0 [110/193] via 10.1.12.2, 00:00:32, Serial0/0.12
O IA    5.5.5.0 [110/257] via 10.1.12.2, 00:00:32, Serial0/0.12
O IA    10.1.23.0 [110/128] via 10.1.12.2, 00:00:32, Serial0/0.12
O IA    10.1.45.0 [110/256] via 10.1.12.2, 00:00:32, Serial0/0.12
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O IA    10.1.34.0 [110/192] via 10.1.12.2, 00:00:32, Serial0/0.12
```

On R5:

```
R5#Show ip route ospf | inc O

O IA    1.1.1.0 [110/257] via 10.1.45.4, 00:00:28, Serial0/0.54
O IA    2.2.2.0 [110/193] via 10.1.45.4, 00:00:28, Serial0/0.54
O IA    3.3.3.0 [110/129] via 10.1.45.4, 00:00:28, Serial0/0.54
O         4.4.4.0 [110/65] via 10.1.45.4, 00:00:28, Serial0/0.54
O IA    10.1.12.0 [110/256] via 10.1.45.4, 00:00:28, Serial0/0.54
O IA    10.1.23.0 [110/192] via 10.1.45.4, 00:00:28, Serial0/0.54
O IA    10.1.34.0 [110/128] via 10.1.45.4, 00:00:28, Serial0/0.54
```

Task 12

Configure MD5 authentication on the link between R1 and R2 in area 0, the password for this authentication should be set to "**micronics**", you should use router configuration mode to accomplish this task.

On R1: and R2

```
Rx(config)#router ospf 1
Rx(config-router)#area 0 authentication message-digest
```

On R1:

```
R1(config)#int s0/0.12
R1(config-subif)#ip ospf message-digest-key 1 md5 micronics
```

On R2:

```
R2(config)#int s0/0.21
R2(config-subif)#ip ospf message-digest-key 1 md5 micronics
```

To verify the configuration:

On R2:

```
R2#Show ip route ospf | inc O
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O       1.1.1.0 [110/65] via 10.1.12.1, 00:02:32, Serial0/0.21
```

If you see other routes in the routing table, wait till you see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 3.3.3.3 on OSPF_VL0 from FULL to DOWN,
Neighbor Down: Dead timer expired
```

Note: R2 does not have any other prefix in its routing table; this is because authentication is enabled directly under the router configuration mode of R1 and R2, which is applied to all links in area 0. Since a virtual link is configured to connect areas 2 and 3 to area 0, R3 and R4 (The ABRs) become area 0 routers. Therefore, authentication must be enabled on these routers as well.

On R3 and R4:

```
Rx(config)#router ospf 1
Rx(config-router)#area 0 authentication message-digest
```

To verify the configuration:

On R5:

```
R5#Show ip route ospf | inc O

O IA    1.1.1.0 [110/257] via 10.1.45.4, 00:00:21, Serial0/0.54
O IA    2.2.2.0 [110/193] via 10.1.45.4, 00:08:43, Serial0/0.54
O IA    3.3.3.0 [110/129] via 10.1.45.4, 00:08:43, Serial0/0.54
O       4.4.4.0 [110/65] via 10.1.45.4, 00:08:43, Serial0/0.54
O IA    10.1.12.0 [110/256] via 10.1.45.4, 00:00:21, Serial0/0.54
O IA    10.1.23.0 [110/192] via 10.1.45.4, 00:08:43, Serial0/0.54
O IA    10.1.34.0 [110/128] via 10.1.45.4, 00:08:43, Serial0/0.54
```

On R2:

```
R2#Show ip route ospf | inc O

O       1.1.1.0 [110/65] via 10.1.12.1, 00:06:18, Serial0/0.21
O IA    3.3.3.0 [110/65] via 10.1.23.3, 00:01:28, Serial0/0.23
O IA    4.4.4.0 [110/129] via 10.1.23.3, 00:01:28, Serial0/0.23
O IA    5.5.5.0 [110/193] via 10.1.23.3, 00:01:28, Serial0/0.23
O IA    10.1.45.0 [110/192] via 10.1.23.3, 00:01:28, Serial0/0.23
O IA    10.1.34.0 [110/128] via 10.1.23.3, 00:01:28, Serial0/0.23
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 13

Remove all authentications and configure MD5 authentication on the link between R1 and R2 using “micronics” as the password. Ensure that every router in this routing domain has all the prefixes advertised by all the other routers in their routing table and link state database. You should **not** configure the other routers to accomplish this task.

On R1, R2, R3 and R4

```
Rx(config)#router ospf 1
Rx(config-router)#No area 0 authentication message-digest
```

On R1:

```
R1(config-router)#int s0/0.12
R1(config-subif)#ip ospf authentication message-digest
```

On R2:

```
R2(config)#int s0/0.21
R2(config-subif)#ip ospf authentication message-digest
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf | inc O

O IA    2.2.2.0 [110/65] via 10.1.12.2, 00:01:41, Serial0/0.12
O IA    3.3.3.0 [110/129] via 10.1.12.2, 00:01:42, Serial0/0.12
O IA    4.4.4.0 [110/193] via 10.1.12.2, 00:01:42, Serial0/0.12
O IA    5.5.5.0 [110/257] via 10.1.12.2, 00:01:42, Serial0/0.12
O IA    10.1.23.0 [110/128] via 10.1.12.2, 00:01:42, Serial0/0.12
O IA    10.1.45.0 [110/256] via 10.1.12.2, 00:01:41, Serial0/0.12
O IA    10.1.34.0 [110/192] via 10.1.12.2, 00:01:41, Serial0/0.12
```

On R5:

```
R5#Show ip route ospf | inc O

O IA    1.1.1.0 [110/257] via 10.1.45.4, 00:02:36, Serial0/0.54
O IA    2.2.2.0 [110/193] via 10.1.45.4, 00:02:36, Serial0/0.54
O IA    3.3.3.0 [110/129] via 10.1.45.4, 00:02:45, Serial0/0.54
O       4.4.4.0 [110/65] via 10.1.45.4, 00:02:45, Serial0/0.54
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O IA    10.1.12.0 [110/256] via 10.1.45.4, 00:02:36, Serial0/0.54
O IA    10.1.23.0 [110/192] via 10.1.45.4, 00:02:36, Serial0/0.54
O IA    10.1.34.0 [110/128] via 10.1.45.4, 00:02:45, Serial0/0.54
```

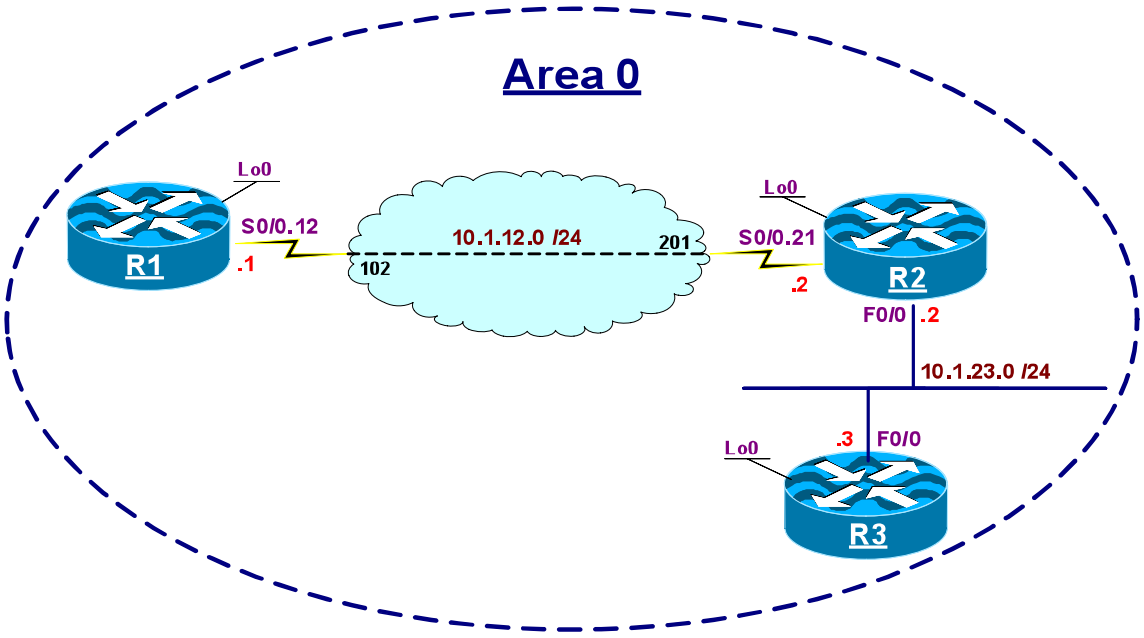
Note: When configuring per-interface authentication, the other routers on the OSPF routing domain do not need to have authentication enabled.

Task 14

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 4
OSPF Cost



Lab Setup:

- Configure R2 and R3 in VLAN 23
- Configure the frame-relay connection between R1 and R2 in a point-to-point manner.

Ip addressing

Router	Interface and IP address
R1	Lo0 = 1.1.1.1 /8 S0/0.12 = 10.1.12.1 /24
R2	**Lo0 = 2.2.2.2 /8** **S0/0.21 = 10.1.12.2 /24** **F0/0 = 10.1.23.2 /24**
R3	Lo0 = 3.3.3.3 /8 F0/0 = 10.1.23.3 /24

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 1

Configure all three routers in OSPF area 0 and advertise their directly connected networks in this area. Ensure that all routers have NLRI to every advertised network. Ensure that loopback interfaces are advertised with their correct mask.

On R1:

```
R1(config)#router ospf 1
R1(config-router)#netw 0.0.0.0 0.0.0.0 are 0

R1(config-router)#int lo0
R1(config-if)#ip ospf net point-to-point
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#netw 0.0.0.0 0.0.0.0 are 0

R2(config-router)#int lo0
R2(config-if)#ip ospf network point-to-point
```

On R3:

```
R3(config)#router ospf 1
R3(config-router)#netw 0.0.0.0 0.0.0.0 are 0

R3(config-router)#int lo0
R3(config-if)#ip ospf network point-to-point
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf

O    2.0.0.0/8 [110/65] via 10.1.12.2, 00:00:50, Serial0/0.12
O    3.0.0.0/8 [110/66] via 10.1.12.2, 00:00:00, Serial0/0.12
     10.0.0.0/24 is subnetted, 2 subnets
O       10.1.23.0 [110/65] via 10.1.12.2, 00:00:10, Serial0/0.12
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 2

Configure R1 such that it advertises a cost of 20 for its loopback 0 interface.

You should check the cost of network 1.0.0.0 /8 that is advertised to R2 by R1.

On R2:

```
R2#Show ip route 1.0.0.0 | inc metric

Known via "ospf 1", distance 110, metric 65, type intra area
      Route metric is 65, traffic share count is 1
```

Note: The cost of R1's loopback 0 interface is 65; this is the result of adding the cost of the serial interface (100,000,000 / 1,544,000 = 64, remember to drop the fractions) plus the cost of the loopback interface (100,000,000 / 8000,000,000 = 1).

Enter the following to change the cost of the lo0 interface on R1:

On R1:

```
R1(config-if)#int lo0
R1(config-if)#ip ospf cost 20
```

To verify the configuration:

On R2:

```
R2#show ip route 1.0.0.0 | inc metric

  Known via "ospf 1", distance 110, metric 84, type intra area
      Route metric is 84, traffic share count is 1
```

Note: The cost is 84, which is the sum of 64 (the total cost of R2 to the Lo0 interface of R1) plus 20.

Task 3

In the future, a gigabit interface will be installed on one of the routers in this routing domain. Ensure that the costs of the other interfaces are adjusted proportionally.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On All Routers

```
Rx(config-if)#router ospf 1
Rx(config-router)#auto-cost reference-bandwidth 1000

Rx#Clear ip ospf proc
Reset ALL OSPF processes? [no]: y
```

By default, OSPF calculates the cost of an interface by dividing the bandwidth of the interface into 100 million bps. The default value is not recommended when your network has interface/s with a bandwidth greater than 100 mbps. This is because, OSPF will not be able to differentiate between 100 mbps interface and an interface with a bandwidth that is greater than 100 mbps, resulting in sub-optimal routing. "IP OSPF COST" command enables you to change the OSPF cost for an interface, but a better way to accomplish this is to change the default reference value. This can be accomplished by modifying the "auto-cost reference-bandwidth" command. If you are planning to use this command, remember that every router in the OSPF routing domain must be configured as well.

Task 4

Remove the "auto-cost reference-bandwidth" command configured in the previous task.

On All Routers

```
Rx(config-if)#router ospf 1
Rx(config-router)#NO auto-cost reference-bandwidth 1000

Rx#Clear ip ospf proc
Reset ALL OSPF processes? [no]: y
```

Task 5

Configure the routers such that the fast ethernet interface of these routers will have a cost of 70. The other interfaces should have their cost calculated proportionally.

The equation used by OSPF is as follows:

Reference/Bandwidth = Cost

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Rearranging the formula, we get the following equation:

Reference = Cost * Bandwidth Which means 70 * 100,000,000 = 7,000,000,000

The configured reference value is in units of Mbps, and by default its set to 100, which means 100,000,000 bps, now we have to divide the result by 1000,000 to get the actual reference, which is 7000.

Therefore, the cost reference should be changed to 7000

On All Routers:

```
Rx(config-router)#router ospf 1
Rx(config-router)#auto-cost reference-bandwidth 7000
```

To verify the configuration:

On R3:

```
R3#Show ip ospf int f0/0 | inc Cost

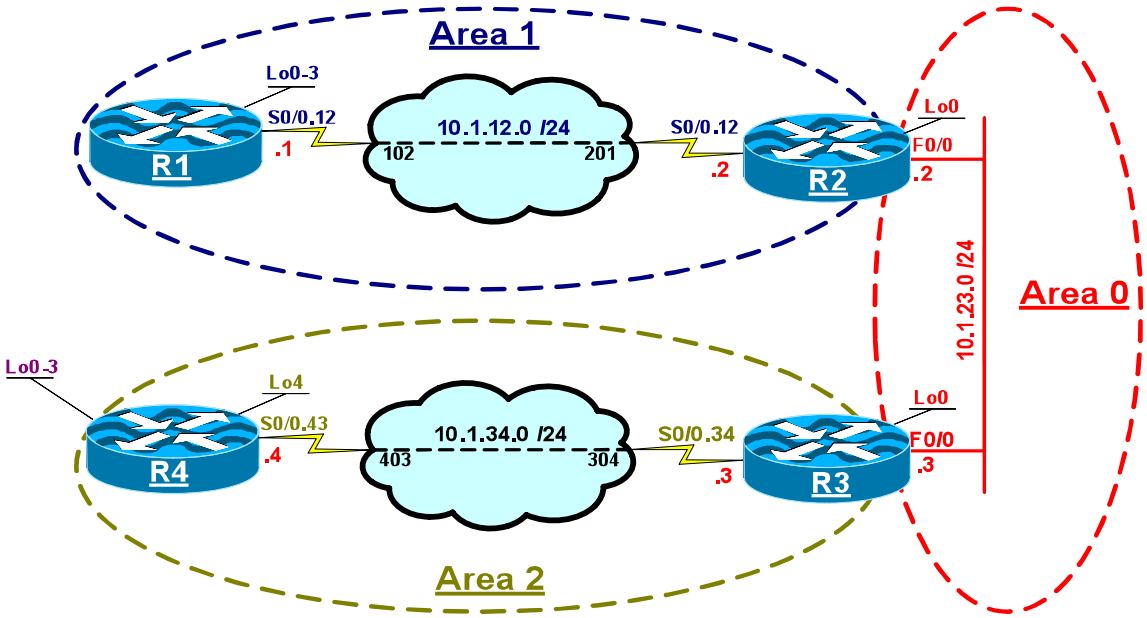
  Process ID 1, Router ID 3.3.3.3, Network Type BROADCAST, Cost: 70
```

Task 6

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 5
OSPF Summarization



Lab setup:

- Configure R2' F0/0 and R3's F0/0 in VLAN 23
- Configure the frame-relay connection between R1, R2 and R3, R4 in a point-to-point manner.

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

IP Addressing

Router	Interface	IP address	Area
R1	Lo0	1.1.0.1 /24	Area 1
	Lo1	1.1.1.1 /24	Area 1
	Lo2	1.1.2.1 /24	Area 1
	Lo3	1.1.3.1 /24	Area 1
	S0/0.12	10.1.12.1 /24	Area 1
R2	Lo0	2.2.2.2 /8	Area 0
	S0/0.21	10.1.12.2 /24	Area 1
	F0/0	10.1.23.2 /24	Area 0
R3	Lo0	3.3.3.3 /8	Area 0
	S0/0.34	10.1.34.3 /24	Area 2
	F0/0	10.1.23.3 /24	Area 0
R4	Lo0	4.4.0.4 /24	External
	Lo1	4.4.1.4 /24	External
	Lo2	4.4.2.4 /24	External
	Lo3	4.4.3.4 /24	External
	Lo4	4.4.4.4 /24	Area 2
	S0/0.43	10.1.34.4 /24	Area 2

Task 1

Configure the routers as follows:

- R4 should redistribute the four-loopback interfaces (4.4.0.4 /24 – 4.4.3.4 /24) in the OSPF routing domain.
- R4 should advertise its loopback 4 and frame-relay interface in area 2.
- R1 should advertise all of its interfaces in OSPF area 1. Use minimum number of commands to accomplish this.
- R2 should advertise its loopback0, F0/0 interface in area 0 and the frame-relay interface in area 1.
- R3 should advertise its loopback0, F0/0 interface in area 0, and its frame-relay interface in area 2.
- R1 should use 1.1.1.1, R2 should use 2.2.2.2, R3 should use 3.3.3.3 and R4 should use 4.4.4.4 as their OSPF router-id

On R1:

```
R1(config)#router ospf 1
R1(config-router)#router-id 1.1.1.1
R1(config-router)#netw 0.0.0.0 0.0.0.0 area 1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2:

```
R2(config)#router ospf 1
R2(config-router)#router-id 2.2.2.2
R2(config-router)#netw 2.2.2.2 0.0.0.0 are 0
R2(config-router)#netw 10.1.23.2 0.0.0.0 are 0
R2(config-router)#netw 10.1.12.2 0.0.0.0 are 1
```

On R3:

```
R3(config)#router ospf 1
R3(config-router)#router-id 3.3.3.3
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.23.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.34.3 0.0.0.0 area 2
```

On R4:

```
R4(config)#route-map TST permit 10
R4(config-route-map)#match int lo0 lo1 lo2 lo3

R4(config-if)#router ospf 1
R4(config-router)#router-id 4.4.4.4
R4(config-router)#netw 4.4.4.4 0.0.0.0 are 2
R4(config-router)#netw 10.1.34.4 0.0.0.0 are 2
R4(config-router)#redistribute connected subnets route-map TST
```

When redistributing routes into OSPF, the subnets keyword will redistribute all the subnetted networks into OSPF. If this keyword is omitted, then, only classful networks will be redistributed into OSPF.

To verify the configuration:

On R1:

```
R1#Show ip route ospf | inc O

O IA    2.2.2.2 [110/65] via 10.1.12.2, 00:03:55, Serial0/0.12
O IA    3.3.3.3 [110/66] via 10.1.12.2, 00:02:49, Serial0/0.12
O IA    4.4.4.4/32 [110/130] via 10.1.12.2, 00:00:49, Serial0/0.12
O E2    4.4.0.0/24 [110/20] via 10.1.12.2, 00:00:18, Serial0/0.12
O E2    4.4.1.0/24 [110/20] via 10.1.12.2, 00:00:18, Serial0/0.12
O E2    4.4.2.0/24 [110/20] via 10.1.12.2, 00:00:18, Serial0/0.12
O E2    4.4.3.0/24 [110/20] via 10.1.12.2, 00:00:18, Serial0/0.12
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O IA    10.1.23.0 [110/65] via 10.1.12.2, 00:03:55, Serial0/0.12
O IA    10.1.34.0 [110/129] via 10.1.12.2, 00:02:39, Serial0/0.12
```

Task 2

Configure the appropriate router in area 2 to summarize all external (E2) routes.

In OSPF, summarization can be configured on two types of routers: ABR/s and/or ASBR/s. The internal OSPF routes can only be summarized on ABR/s, whereas, the external (redistributed) routes can be summarized on the router that originates the external routes.

If the external routes are LSA type 7s, they must be summarized on an ASBR, whereas, if the external routes are LSA type 5s, they can be summarized on the router that originates type 5s, this could be an ABR (in an NSSA area) or an ASBR (in a normal area).

When summarizing internal routes on ABRs the "area xx range" command must be used, where xx is the area id.

Summarization of external routes can be accomplished by using the "summary-address" command. In this case, since R4 is the router that originates type-5s, summarization can only be performed On R4: and R4 ONLY.

On R4:

```
R4(config)#router ospf 1
R4(config-router)#summary-address 4.4.0.0 255.255.252.0
```

Note: Whenever summarization is performed on a given router, the IOS will auto-inject a summary route to null0 to avoid forwarding loops.

To see the Null0 route for the summary:

On R4:

```
R4#Show ip route ospf | inc O

O IA    1.1.1.1 [110/130] via 10.1.34.3, 00:04:45, Serial0/0.43
O IA    1.1.0.1 [110/130] via 10.1.34.3, 00:04:45, Serial0/0.43
O IA    1.1.3.1 [110/130] via 10.1.34.3, 00:04:45, Serial0/0.43
O IA    1.1.2.1 [110/130] via 10.1.34.3, 00:04:45, Serial0/0.43
O IA    2.2.2.2 [110/66] via 10.1.34.3, 00:17:54, Serial0/0.43
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O IA     3.3.3.3 [110/65] via 10.1.34.3, 00:17:54, Serial0/0.43
O        4.4.0.0/22 is a summary, 00:09:31, Null0
O IA     10.1.12.0 [110/129] via 10.1.34.3, 00:17:54, Serial0/0.43
O IA     10.1.23.0 [110/65] via 10.1.34.3, 00:17:54, Serial0/0.43
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf | inc O

O IA     2.2.2.2 [110/65] via 10.1.12.2, 00:11:58, Serial0/0.12
O IA     3.3.3.3 [110/66] via 10.1.12.2, 00:10:52, Serial0/0.12
O IA     4.4.4.4/32 [110/130] via 10.1.12.2, 00:08:52, Serial0/0.12
O E2     4.4.0.0/22 [110/20] via 10.1.12.2, 00:00:36, Serial0/0.12
O IA     10.1.23.0 [110/65] via 10.1.12.2, 00:11:58, Serial0/0.12
O IA     10.1.34.0 [110/129] via 10.1.12.2, 00:10:42, Serial0/0.12
```

Note: The external routes are summarized.

Task 3

Configure Area 1 such that networks (1.1.0.0 /24, 1.1.1.0 /24, 1.1.2.0 /24 and 1.1.3.0 /24) are summarized into the OSPF routing domain.

Note: These routes are originated by R1, but they can only be summarized by the ABR of the area that originated the specific routes, in this topology the ABR is R2. Since the routes that are being summarized originated in area 1, the "area range" command must specify area 1 range" followed by the summary network address (1.1.0.0) and then the subnet mask (255.255.252.0).

On R2:

```
R2(config)#router ospf 1
R2(config-router)#area 1 range 1.1.0.0 255.255.252.0
```

Note: Whenever summarization is performed on a given router, the IOS will auto-inject a summary route to null0 to avoid forwarding loops.

To see the Null0 route for the summary:

On R2:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2#Show ip route ospf | inc O

O        1.1.1.1/32 [110/65] via 10.1.12.1, 00:01:48, Serial0/0.21
O        1.1.0.0/22 is a summary, 00:01:48, Null0
O        1.1.0.1/32 [110/65] via 10.1.12.1, 00:01:48, Serial0/0.21
O        1.1.3.1/32 [110/65] via 10.1.12.1, 00:01:48, Serial0/0.21
O        1.1.2.1/32 [110/65] via 10.1.12.1, 00:01:48, Serial0/0.21
O        3.3.3.3 [110/2] via 10.1.23.3, 00:01:48, FastEthernet0/0
O IA     4.4.4.4/32 [110/66] via 10.1.23.3, 00:01:48, FastEthernet0/0
O E2     4.4.0.0/22 [110/20] via 10.1.23.3, 00:01:48, FastEthernet0/0
O IA     10.1.34.0 [110/65] via 10.1.23.3, 00:01:48, FastEthernet0/0
```

To verify the configuration:

On R3:

```
R3#Show ip route ospf | inc O

O IA     1.1.0.0 [110/66] via 10.1.23.2, 00:00:34, FastEthernet0/0
O        2.2.2.2 [110/2] via 10.1.23.2, 00:14:00, FastEthernet0/0
O        4.4.4.4/32 [110/65] via 10.1.34.4, 00:11:59, Serial0/0.34
O E2     4.4.0.0/22 [110/20] via 10.1.34.4, 00:03:43, Serial0/0.34
O IA     10.1.12.0 [110/65] via 10.1.23.2, 00:14:00, FastEthernet0/0
```

Task 4

The routers should **not** install a null 0 route in the routing table when they summarize internal or external routes.

In OSPF, the discard route is created automatically whenever a summary route is configured, there are two types of summary routes: Internal and External.

When internal summary routes are configured, OSPF will inject an internal discard route, and when an external summary route is configured, the OSPF process will create an external discard route. The discard routes are created to stop forwarding loops.

To remove the internal discard route:

On R2:

```
R2#Show ip route ospf | inc summary
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O        1.1.0.0/22 is a summary, 00:00:58, Null0
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#NO discard-route internal
```

To verify the configuration:

On R2:

```
R2#Show ip route ospf | inc summary
R2#
```

To remove the external discard route:

On R4:

```
R4#Show ip route ospf | inc summary

O        4.4.0.0/22 is a summary, 00:06:13, Null0
```

On R4:

```
R4(config)#router ospf 1
R4(config-router)#NO discard-route external
```

To Verify the configuration:

On R4:

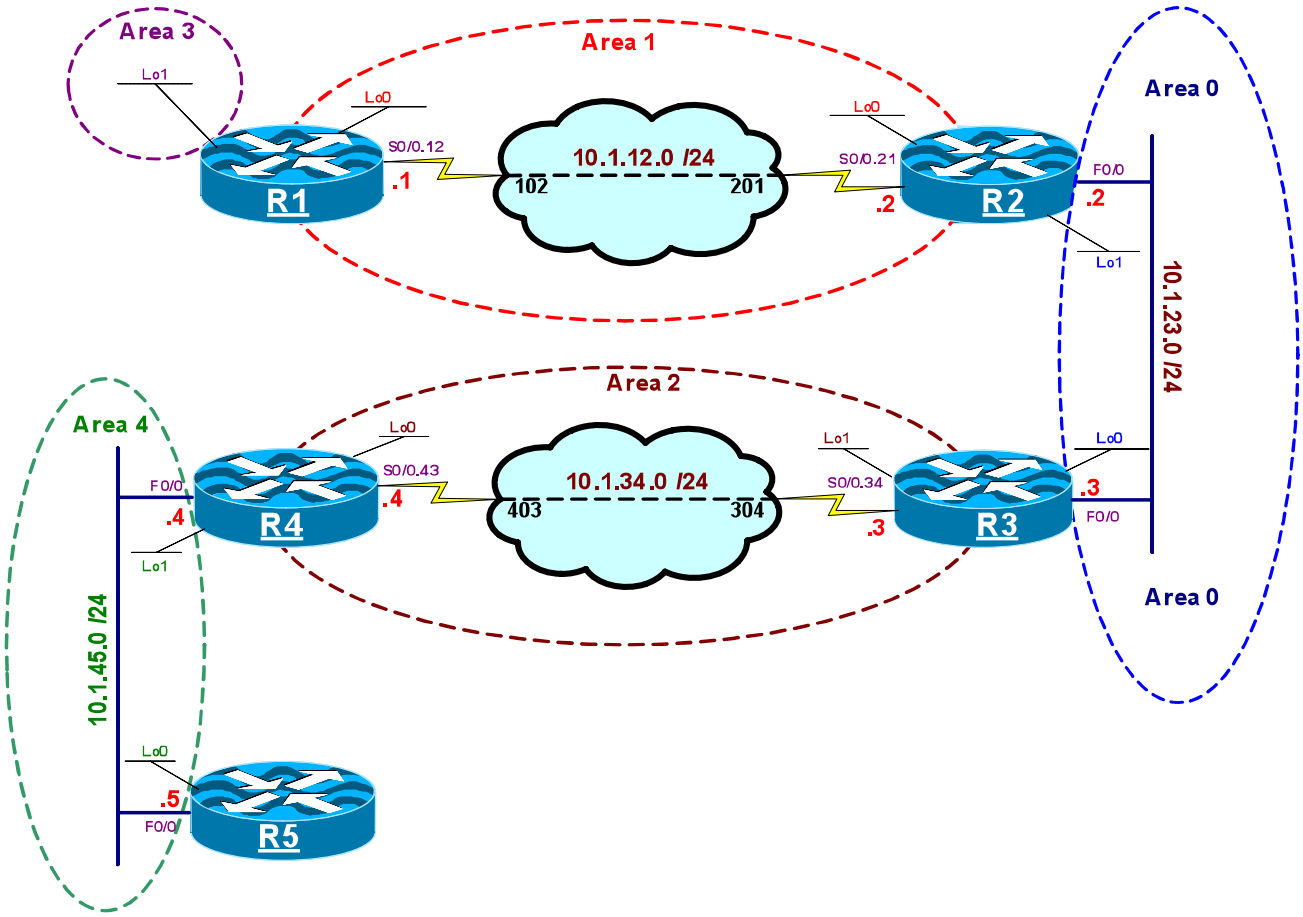
```
R4#Show ip route ospf | inc summary
R4#
```

Task 5

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 6
Virtual-Links and GRE Tunnels



Lab Setup:

- Use the IP addressing chart below to assign IP addresses to the interfaces.
- The frame-relay connection between R1 and R2 should be configured in a point-to-point manner.
- R2 and R3's F0/0 interface should be configured in VLAN 23.
- The frame-relay connection between R3 and R4 should be configured in a point-to-point manner.
- R4 and R5's F0/0 interface should be configured in VLAN 45.

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

IP Addressing:

Router	Interface	IP address
R1	Lo0 Lo1 F/R interface	1.1.1.1 /24 10.1.1.1 / 24 10.1.12.1 /24
R2	**Lo0** **Lo1** **F0/0** **F/R interface**	**2.2.2.2 /24** **20.2.2.2 /24** **10.1.23.2 /24** **10.1.12.2 /24**
R3	Lo0 Lo1 F0/0 F/R interface	3.3.3.3 /24 30.3.3.3 /24 10.1.23.3 /24 10.1.34.3 /24
R4	**Lo0** **Lo1** **F0/0** **F/R interface**	**4.4.4.4 /24** **40.4.4.4 /24** **10.1.45.4 /24** **10.1.34.4 /24**
R5	Lo0 F0/0	5.5.5.5 /24 10.1.45.5 /24

Task 1

- R1's loopback1 interface should be advertised in area 3 and its frame-relay and loopback0 interface should be advertised in area 1
- R2's loopback0 and its frame-relay interface should be configured in area 1 and its loopback1 and F0/0 interface should be configured in area 0
- R3's loopback0 and F0/0 interface should be configured in area 0 and its frame-relay and loopback1 interface should be configured in area 2.
- R4's frame-relay and loopback0 interface should be configured in area 2 and its F0/0 and loopback1 interface should be configured in area 4.
- R5's loopback and F0/0 interface should be configured in area 4.
- These routers should use their Lo0 interface as their OSPF router-id.

On R1:

```
R1(config)#router ospf 1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config-router)#router-id 1.1.1.1
R1(config-router)#netw 10.1.1.1 0.0.0.0 are 3
R1(config-router)#netw 10.1.12.1 0.0.0.0 are 1
R1(config-router)#netw 1.1.1.1 0.0.0.0 are 1
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#router-id 2.2.2.2
R2(config-router)#netw 10.1.12.2 0.0.0.0 are 1
R2(config-router)#netw 2.2.2.2 0.0.0.0 are 1
R2(config-router)#netw 10.1.23.2 0.0.0.0 are 0
R2(config-router)#netw 20.2.2.2 0.0.0.0 are 0
```

On R3:

```
R3(config-if)#router ospf 1
R3(config-router)#router-id 3.3.3.3
R3(config-router)#netw 10.1.23.3 0.0.0.0 are 0
R3(config-router)#netw 3.3.3.3 0.0.0.0 are 0
R3(config-router)#netw 10.1.34.3 0.0.0.0 are 2
R3(config-router)#netw 30.3.3.3 0.0.0.0 are 2
```

On R4:

```
R4(config)#router ospf 1
R4(config-router)#router-id 4.4.4.4
R4(config-router)#netw 10.1.34.4 0.0.0.0 area 2
R4(config-router)#netw 4.4.4.4 0.0.0.0 are 2
R4(config-router)#netw 10.1.45.4 0.0.0.0 are 4
R4(config-router)#netw 40.4.4.4 0.0.0.0 are 4
```

On R5:

```
R5(config)#router ospf 1
R5(config-router)#router-id 5.5.5.5
R5(config-router)#netw 10.1.45.5 0.0.0.0 are 4
R5(config-router)#netw 5.5.5.5 0.0.0.0 are 4
```

Task 2

Ensure that the route/s from area 3 are reachable by R2, R3, and R4. **Do not** use a GRE tunnel to accomplish this task.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Since area 3 is not connected to area 0, the other routers will not be able to see the route/s advertised by this area (10.1.1.0 / 24). A virtual-link must be created that connects R1 (The ABR of area 3) to area 0.

On R1:

```
R1(config)#router ospf 1
R1(config-router)#area 1 virtual-link 2.2.2.2
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#area 1 virtual-link 1.1.1.1
```

To Verify the configuration:

On R2:

```
R2#Show ip route ospf | inc O

O       1.1.1.1 [110/65] via 10.1.12.1, 00:08:11, Serial0/0.21
O       3.3.3.3 [110/2] via 10.1.23.3, 00:07:09, FastEthernet0/0
O IA    4.4.4.4 [110/66] via 10.1.23.3, 00:06:03, FastEthernet0/0
O IA    10.1.1.1/32 [110/65] via 10.1.12.1, 00:00:09, Serial0/0.21
O IA    10.1.34.0/24 [110/65] via 10.1.23.3, 00:06:59, FastEthernet0/0
O IA    30.3.3.3 [110/2] via 10.1.23.3, 00:06:51, FastEthernet0/0

R2#Ping 10.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/50/52 ms
```

On R3:

```
R3#Show ip rou ospf | inc O

O IA    1.1.1.1 [110/66] via 10.1.23.2, 00:01:11, FastEthernet0/0
O IA    2.2.2.2 [110/2] via 10.1.23.2, 00:08:18, FastEthernet0/0
O       4.4.4.4 [110/65] via 10.1.34.4, 00:07:15, Serial0/0.34
O       20.2.2.2 [110/2] via 10.1.23.2, 00:08:18, FastEthernet0/0
O IA    10.1.12.0/24 [110/65] via 10.1.23.2, 00:08:18, FastEthernet0/0
O IA    10.1.1.1/32 [110/66] via 10.1.23.2, 00:01:11, FastEthernet0/0

R3#Ping 10.1.1.1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/50/52 ms
```

On R4:

```
R4#Show ip route ospf | inc O

O IA     1.1.1.1 [110/130] via 10.1.34.3, 00:07:56, Serial0/0.43
O IA     2.2.2.2 [110/66] via 10.1.34.3, 00:07:56, Serial0/0.43
O IA     3.3.3.3 [110/65] via 10.1.34.3, 00:07:56, Serial0/0.43
O IA     20.2.2.2 [110/66] via 10.1.34.3, 00:07:56, Serial0/0.43
O        5.5.5.5 [110/2] via 10.1.45.5, 00:07:03, FastEthernet0/0
O IA     10.1.12.0/24 [110/129] via 10.1.34.3, 00:07:56, Serial0/0.43
O IA     10.1.1.1/32 [110/130] via 10.1.34.3, 00:01:52, Serial0/0.43
O IA     10.1.23.0/24 [110/65] via 10.1.34.3, 00:07:56, Serial0/0.43
O        30.3.3.3 [110/65] via 10.1.34.3, 00:07:56, Serial0/0.43

R4#Ping 10.1.1.1

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

Task 3

Ensure that all the advertised networks are reachable by all the routers. Use any IP addressing and **do not** use a virtual-link to accomplish this task. You should **not** remove the existing virtual-link.

The routing table of R5 reveals that only network 40.4.4.0 / 24 was propagated by R4.

On R5:

```
R5#Show ip route ospf | inc O

O        40.4.4.4 [110/2] via 10.1.45.4, 00:10:48, FastEthernet0/0
```

The reason for this behavior is as follows:
Area 4 does not have a connection (logical or physical) to area 0.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

In order to rectify this problem we must configure a virtual-link, since virtual-link is prohibited by the condition of this task, a GRE tunnel must be used.

On R4:

```
R4(config)#int tu1
R4(config-if)#ip addr 200.1.34.4 255.255.255.0
R4(config-if)#tunnel source 10.1.34.4
R4(config-if)#tunnel destination 10.1.34.3

R4(config-if)#router ospf 1
R4(config-router)#netw 200.1.34.4 0.0.0.0 are 0
```

On R3:

```
R3(config)#int tu1
R3(config-if)#ip addr 200.1.34.3 255.255.255.0
R3(config-if)#tunnel source 10.1.34.3
R3(config-if)#tunnel destination 10.1.34.4

R3(config)#router ospf 1
R3(config-router)#netw 200.1.34.3 0.0.0.0 are 0
```

To Verify the configuration:

On R5:

```
R5#Show ip route ospf | inc O

O IA    1.1.1.1 [110/11178] via 10.1.45.4, 00:00:26, FastEthernet0/0
O IA    2.2.2.2 [110/11114] via 10.1.45.4, 00:00:26, FastEthernet0/0
O IA    3.3.3.3 [110/11113] via 10.1.45.4, 00:00:26, FastEthernet0/0
O IA    4.4.4.4 [110/2] via 10.1.45.4, 00:01:10, FastEthernet0/0
O IA 200.1.34.0/24 [110/11112] via 10.1.45.4, 00:01:10, FastEthernet0/0
O IA    20.2.2.2 [110/11114] via 10.1.45.4, 00:00:26, FastEthernet0/0
O       40.4.4.4 [110/2] via 10.1.45.4, 00:16:12, FastEthernet0/0
O IA    10.1.12.0/24 [110/11177] via 10.1.45.4, 00:00:26, FastEthernet0/0
O IA    10.1.1.1/32 [110/11178] via 10.1.45.4, 00:00:26, FastEthernet0/0
O IA    10.1.23.0/24 [110/11113] via 10.1.45.4, 00:00:26, FastEthernet0/0
O IA    10.1.34.0/24 [110/65] via 10.1.45.4, 00:01:10, FastEthernet0/0
O IA    30.3.3.3 [110/66] via 10.1.45.4, 00:01:10, FastEthernet0/0
```

Note: All the routes are advertised. The IP address of the tunnel interface must be advertised in area 0 or else this solution will not work.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 4

Remove the configuration from the previous task and replace it with virtual-link.

On R4:

```
R4(config-if)#router ospf 1
R4(config-router)#No netw 200.1.34.4 0.0.0.0 are 0

R4(config)#No int tu1

R4(config)#router ospf 1
R4(config-router)#area 2 virtual-link 3.3.3.3
```

On R3:

```
R3(config)#router ospf 1
R3(config-router)#No netw 200.1.34.3 0.0.0.0 are 0

R3(config)#No int tu1

R3(config)#router ospf 1
R3(config-router)#area 2 virtual-link 4.4.4.4
```

To verify the configuration:

On R5:

```
R5#Show ip route ospf | inc O

O IA    1.1.1.1 [110/131] via 10.1.45.4, 00:00:01, FastEthernet0/0
O IA    2.2.2.2 [110/67] via 10.1.45.4, 00:00:01, FastEthernet0/0
O IA    3.3.3.3 [110/66] via 10.1.45.4, 00:00:01, FastEthernet0/0
O IA    4.4.4.4 [110/2] via 10.1.45.4, 00:00:11, FastEthernet0/0
O IA    20.2.2.2 [110/67] via 10.1.45.4, 00:00:01, FastEthernet0/0
O       40.4.4.4 [110/2] via 10.1.45.4, 00:19:33, FastEthernet0/0
O IA    10.1.12.0/24 [110/130] via 10.1.45.4, 00:00:01, FastEthernet0/0
O IA    10.1.1.1/32 [110/131] via 10.1.45.4, 00:00:01, FastEthernet0/0
O IA    10.1.23.0/24 [110/66] via 10.1.45.4, 00:00:01, FastEthernet0/0
O IA    10.1.34.0/24 [110/65] via 10.1.45.4, 00:00:11, FastEthernet0/0
O IA    30.3.3.3 [110/66] via 10.1.45.4, 00:00:11, FastEthernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 5

Configure a simple clear text authentication for the virtual-link that connects area 3 to area 0. Use "**Cisco**" as the password.

On R1:

The first "area 1 virtual-link 2.2.2.2 authentication" command enables authentication, and the second command applies the password:

```
R1(config)#router ospf 1
R1(config-router)#area 1 virtual-link 2.2.2.2 authentication
R1(config-router)#area 1 virtual-link 2.2.2.2 authentication-key Cisco
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#area 1 virtual-link 1.1.1.1 authentication
R2(config-router)#area 1 virtual-link 1.1.1.1 authentication-key Cisco
```

To verify the configuration:

On R1:

```
R1#Show ip ospf virtual-links

Virtual Link OSPF_VL0 to router 2.2.2.2 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 1, via interface Serial0/0.12, 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 00:00:03
    Adjacency State FULL (Hello suppressed)
    Index 1/2, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum is 0 msec
  Simple password authentication enabled
```

On R5:

```
R5#Show ip route ospf | inc O
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O IA     1.1.1.1 [110/131] via 10.1.45.4, 00:00:01, FastEthernet0/0
O IA     2.2.2.2 [110/67] via 10.1.45.4, 00:00:01, FastEthernet0/0
O IA     3.3.3.3 [110/66] via 10.1.45.4, 00:00:01, FastEthernet0/0
O IA     4.4.4.4 [110/2] via 10.1.45.4, 00:00:11, FastEthernet0/0
O IA     20.2.2.2 [110/67] via 10.1.45.4, 00:00:01, FastEthernet0/0
O        40.4.4.4 [110/2] via 10.1.45.4, 00:19:33, FastEthernet0/0
O IA     10.1.12.0/24 [110/130] via 10.1.45.4, 00:00:01, FastEthernet0/0
O IA     10.1.1.1/32 [110/131] via 10.1.45.4, 00:00:01, FastEthernet0/0
O IA     10.1.23.0/24 [110/66] via 10.1.45.4, 00:00:01, FastEthernet0/0
O IA     10.1.34.0/24 [110/65] via 10.1.45.4, 00:00:11, FastEthernet0/0
O IA     30.3.3.3 [110/66] via 10.1.45.4, 00:00:11, FastEthernet0/0
```

To test the configuration:

On R5:

```
R5#Ping 10.1.1.1

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

Task 6

Configure MD5 authentication for the virtual-link that connects area 4 to area 0, use "**cisco**" as the password.

On R4:

```
R4(config)#router ospf 1
R4(config-router)#area 2 virtual-link 3.3.3.3 authentication message-digest
R4(config-router)#area 2 virtual-link 3.3.3.3 message-digest-key 1 md5 cisco
```

On R3:

```
R3(config)#router ospf 1
R3(config-router)#area 2 virtual-link 4.4.4.4 authentication message-digest
R3(config-router)#area 2 virtual-link 4.4.4.4 message-digest-key 1 md5 cisco
```

To verify and test the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1:

```
R1#Show ip route ospf | inc O

O       2.2.2.2 [110/65] via 10.1.12.2, 00:27:44, Serial0/0.12
O       3.3.3.3 [110/66] via 10.1.12.2, 00:27:23, Serial0/0.12
O IA    4.4.4.4 [110/130] via 10.1.12.2, 00:12:44, Serial0/0.12
O       20.2.2.2 [110/65] via 10.1.12.2, 00:27:23, Serial0/0.12
O IA    5.5.5.5 [110/131] via 10.1.12.2, 00:12:44, Serial0/0.12
O IA    40.4.4.4 [110/130] via 10.1.12.2, 00:12:44, Serial0/0.12
O       10.1.23.0 [110/65] via 10.1.12.2, 00:27:23, Serial0/0.12
O IA    10.1.45.0 [110/130] via 10.1.12.2, 00:12:44, Serial0/0.12
O IA    10.1.34.0 [110/129] via 10.1.12.2, 00:27:23, Serial0/0.12
O IA    30.3.3.3 [110/66] via 10.1.12.2, 00:27:23, Serial0/0.12

R1#Ping 5.5.5.5

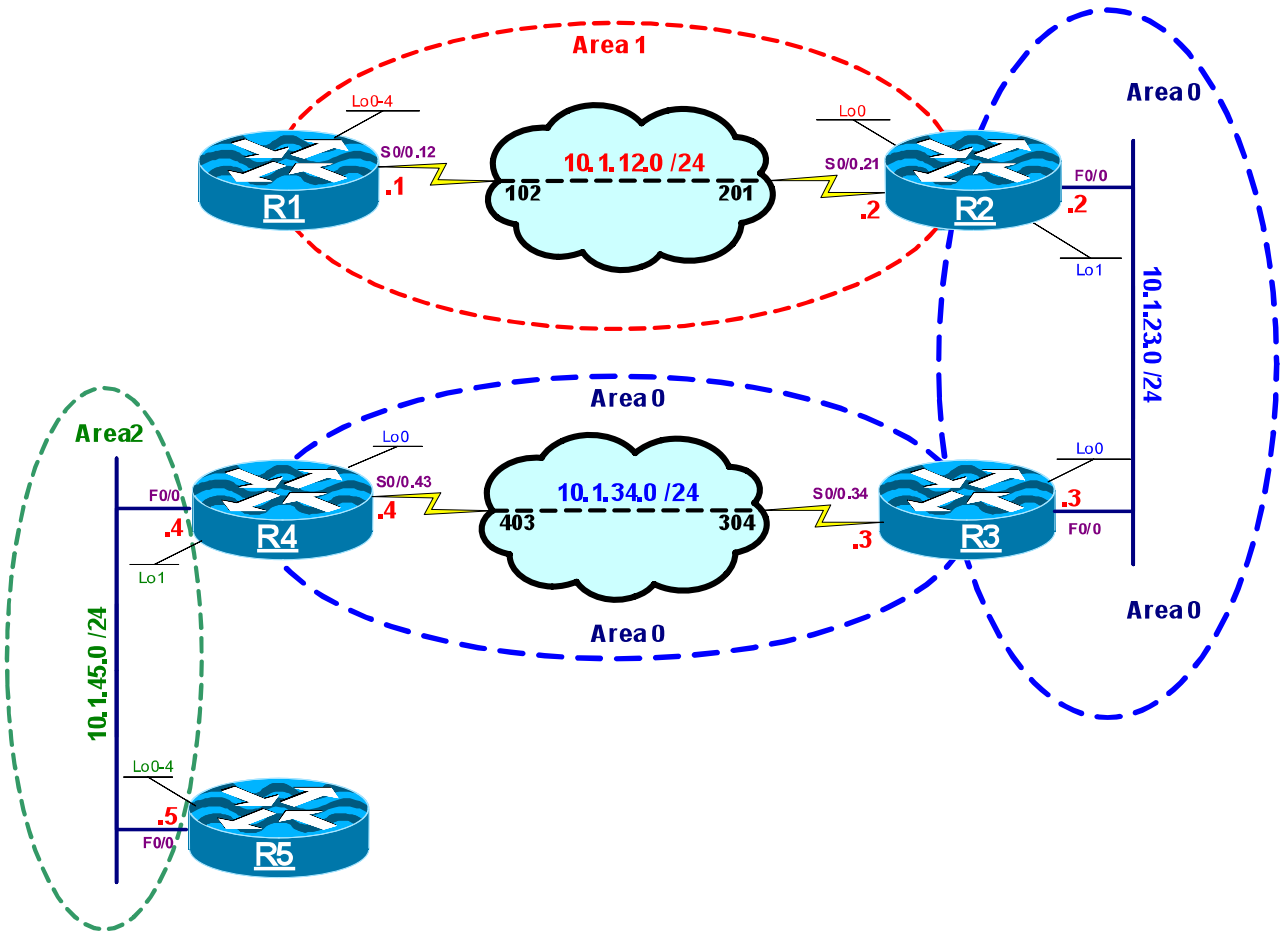
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 100/100/100 ms
```

Task 7

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 7
OSPF Stub, Totally Stubby, and NSSA Areas



Lab Setup:

- R4 and R5's F0/0 interface should be configured in VLAN 45.
- R2 and R3's F0/0 interface should be configured in VLAN 23.
- The frame-relay connection between R1, R2 and R3, R4 should be configured in a point-to-point manner.
- Use the IP addressing chart below to assign IP addresses to the routers.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP Addressing:

Router	Interface	IP address
R1	Lo0 Lo1 Lo2 Lo3 Lo4 S0/0.12	1.1.1.1 /24 11.1.0.1 /24 11.1.1.1 /24 11.1.2.1 /24 11.1.3.1 /24 10.1.12.1 /24
R2	**Lo0** **Lo1** **F0/0** **S0/0.21**	**2.2.2.2 /24** **22.2.2.2 /24** **10.1.23.2 /24** **10.1.12.2 /24**
R3	Lo0 F0/0 S0/0.34	3.3.3.3 /24 10.1.23.3 /24 10.1.34.3 /24
R4	**Lo0** **Lo1** **F0/0** **S0/0.43**	**4.4.4.4 /24** **44.4.4.4 /24** **10.1.45.4 /24** **10.1.34.4 /24**
R5	Lo0 F0/0	5.5.5.5 /24 10.1.45.5 /24

Task 1

Configure OSPF as follows:

- Configure R1's loopback0 and frame-relay interface in area 1
- Configure R2's loopback0 and frame-relay interface in area 1 and its loopback1 and F0/0 interface should be configured in area 0.
- Configure R3's loopback0, F0/0, and frame-relay interface in OSPF area 0.
- Configure R4's loopback0, and frame-relay interface in area 0, and its loopback1 and F0/0 in OSPF area 2.
- Configure R5's loopback0 and F0/0 interface in OSPF area 2.
- The loopback interfaces must be advertised with their correct mask.
- Routers should use their Lo0 as their router id.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1:

```
R1(config)#router ospf 1
R1(config-router)#router-id 1.1.1.1

R1(config-router)#netw 1.1.1.1 0.0.0.0 are 1
R1(config-router)#netw 10.1.12.1 0.0.0.0 are 1

R1(config)#interface lo0
R1(config-if)#ip ospf network point-to-point
```

On R2:

```
R2(config-if)#router ospf 1
R2(config-router)#router-id 2.2.2.2
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 1
R2(config-router)#netw 22.2.2.2 0.0.0.0 are 0
R2(config-router)#netw 10.1.12.2 0.0.0.0 are 1
R2(config-router)#netw 10.1.23.2 0.0.0.0 are 0

R2(config)#interface lo0
R2(config-if)#ip ospf network point-to-point

R2(config)#interface lo1
R2(config-if)#ip ospf network point-to-point
```

On R3:

```
R3(config-if)#router ospf 1
R3(config-router)#router-id 3.3.3.3
R3(config-router)#netw 3.3.3.3 0.0.0.0 are 0
R3(config-router)#netw 10.1.23.3 0.0.0.0 are 0
R3(config-router)#netw 10.1.34.3 0.0.0.0 are 0

R3(config)#interface lo0
R3(config-if)#ip ospf network point-to-point
```

On R4:

```
R4(config-if)#router ospf 1
R4(config-router)#router-id 4.4.4.4
R4(config-router)#netw 4.4.4.4 0.0.0.0 are 0
R4(config-router)#netw 10.1.34.4 0.0.0.0 are 0
R4(config-router)#netw 44.4.4.4 0.0.0.0 are 2
R4(config-router)#netw 10.1.45.4 0.0.0.0 are 2
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config)#interface lo0
R4(config-if)#ip ospf network point-to-point

R4(config)#interface lo1
R4(config-if)#ip ospf network point-to-point
```

On R5:

```
R5(config-if)#router ospf 1
R5(config-router)#router-id 5.5.5.5
R5(config-router)#netw 5.5.5.5 0.0.0.0 are 2
R5(config-router)#netw 10.1.45.5 0.0.0.0 are 2

R5(config)#interface lo0
R5(config-if)#ip ospf network point-to-point
```

To verify the configuration:

On R5:

```
R5#Show ip rou ospf | inc O

O IA    1.1.1.0 [110/131] via 10.1.45.4, 00:00:10, FastEthernet0/0
O IA    2.2.2.0 [110/67] via 10.1.45.4, 00:00:10, FastEthernet0/0
O IA    3.3.3.0 [110/66] via 10.1.45.4, 00:00:10, FastEthernet0/0
O IA    4.4.4.0 [110/2] via 10.1.45.4, 00:00:10, FastEthernet0/0
O IA    22.2.2.0 [110/67] via 10.1.45.4, 00:00:10, FastEthernet0/0
O IA    10.1.12.0 [110/130] via 10.1.45.4, 00:00:10, FastEthernet0/0
O IA    10.1.23.0 [110/66] via 10.1.45.4, 00:00:10, FastEthernet0/0
O IA    10.1.34.0 [110/65] via 10.1.45.4, 00:00:10, FastEthernet0/0
O       44.4.4.0 [110/2] via 10.1.45.4, 00:00:10, FastEthernet0/0
```

Task 2

Configure area 1 such that it does not receive LSA types 4 and 5, but the routers in this area should still maintain LSA type 3 in their routing table for the routes from area 1. These routers should have reachability to the existing and future external routes redistributed into this routing domain.

On R1:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config)#router ospf 1
R1(config-router)#area 1 stub
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#area 1 stub
```

Important points to understand about a STUB area:

- **A STUB area cannot be a transit area for virtual links, but a GRE tunnel can be used instead.**
- **A STUB area cannot have an ASBR.**
- **The backbone area cannot be configured as a STUB area.**
- **Every router and the ABR of that area should have "area xx stub" command.**
- **LSA type 5s (E1, or E2) are not allowed in a STUB area, but the routers in the STUB area can connect to the external routes via the default route that is injected in the area by their ABR.**
- **By default, the cost of the default route is 1; this can be verified by "show ip ospf", and "show ip route". The cost of the default route can be changed by "area xx default-cost cc", where xx is the area number, and cc is the desired cost**

To verify the configuration:

On R1:

Note the cost of the default route is 65 which is the Total of FR links cost + 1.

```
R1#Show ip route ospf | inc O

O        2.2.2.0 [110/65] via 10.1.12.2, 00:00:02, Serial0/0.12
O IA     3.3.3.0 [110/66] via 10.1.12.2, 00:00:02, Serial0/0.12
O IA     4.4.4.0 [110/130] via 10.1.12.2, 00:00:02, Serial0/0.12
O IA     5.5.5.0 [110/131] via 10.1.12.2, 00:00:02, Serial0/0.12
O IA     22.2.2.0 [110/65] via 10.1.12.2, 00:00:02, Serial0/0.12
O IA     10.1.23.0 [110/65] via 10.1.12.2, 00:00:02, Serial0/0.12
O IA     10.1.45.0 [110/130] via 10.1.12.2, 00:00:02, Serial0/0.12
O IA     10.1.34.0 [110/129] via 10.1.12.2, 00:00:02, Serial0/0.12
O IA     44.4.4.0 [110/130] via 10.1.12.2, 00:00:02, Serial0/0.12
O*IA 0.0.0.0/0 [110/65] via 10.1.12.2, 00:00:02, Serial0/0.12
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 3

Area 2 should not receive LSA types 3, 4 and 5. These routers should have reachability to the existing and future external routes redistributed into this routing domain.

On R5:

Note: You can reduce the size of the routing table further by configuring an area as totally stubby. Since all the IA and E (inter-area and external) routes can be reached via a default route, which is injected by the ABR of this area, there is no reason to maintain these routes.

```
R5(config)#router ospf 1
R5(config-router)#area 2 stub
```

The above command must be configured on all the routers within area 2.

On R4:

```
R4(config)#router ospf 1
R4(config-router)#area 2 stub no-summary
```

The above command must only be configured on the ABR of this area.

To verify the configuration:

On R5:

```
R5#Show ip route ospf | inc O

O       44.4.4.0 [110/2] via 10.1.45.4, 00:04:33, FastEthernet0/0
O*IA 0.0.0.0/0 [110/2] via 10.1.45.4, 00:04:33, FastEthernet0/0
```

Task 4

Redistribute Lo1 – 4 interfaces of R1 into this routing domain. After the redistribution, area 1 should only receive and propagate LSA types 1, 2, 3 and 7. This area should not have the ability to connect to any external routes redistributed elsewhere within this routing domain.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Note: Since a stub area can not have external routes, area 1 is converted into a NSSA.

On R1:

```
R1(config)#router ospf 1
R1(config-router)#NO area 1 stub
R1(config-router)#area 1 nssa
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#NO area 1 stub
R2(config-router)#area 1 nssa
```

On R1:

```
R1(config)#route-map TST permit 10
R1(config-route-map)#match interface Lo1 Lo2 Lo3 Lo4

R1(config)#router ospf 1
R1(config-router)#redistribute connected route-map TST subnets
```

Note: When configuring an area as an NSSA area, a default route is not injected by the ABR of that area.

To verify the configuration:

On R1:

Note: R1 does not have a default route, because the ABR did not inject one:

```
R1#Show ip route ospf | inc O

O       2.2.2.0 [110/65] via 10.1.12.2, 00:02:32, Serial0/0.12
O IA    3.3.3.0 [110/66] via 10.1.12.2, 00:02:32, Serial0/0.12
O IA    4.4.4.0 [110/130] via 10.1.12.2, 00:02:32, Serial0/0.12
O IA    5.5.5.0 [110/131] via 10.1.12.2, 00:02:32, Serial0/0.12
O IA    22.2.2.0 [110/65] via 10.1.12.2, 00:02:32, Serial0/0.12
O IA    10.1.23.0 [110/65] via 10.1.12.2, 00:02:32, Serial0/0.12
O IA    10.1.45.0 [110/130] via 10.1.12.2, 00:02:32, Serial0/0.12
O IA    10.1.34.0 [110/129] via 10.1.12.2, 00:02:32, Serial0/0.12
O IA    44.4.4.0 [110/130] via 10.1.12.2, 00:02:32, Serial0/0.12
```

On R2:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2#Show ip route ospf | inc O

O        1.1.1.0 [110/65] via 10.1.12.1, 00:03:37, Serial0/0.21
O        3.3.3.0 [110/2] via 10.1.23.3, 00:03:52, FastEthernet0/0
O        4.4.4.0 [110/66] via 10.1.23.3, 00:03:52, FastEthernet0/0
O IA     5.5.5.0 [110/67] via 10.1.23.3, 00:03:52, FastEthernet0/0
O IA     10.1.45.0 [110/66] via 10.1.23.3, 00:03:52, FastEthernet0/0
O        10.1.34.0 [110/65] via 10.1.23.3, 00:03:52, FastEthernet0/0
O N2     11.1.2.0 [110/20] via 10.1.12.1, 00:02:50, Serial0/0.21
O N2     11.1.3.0 [110/20] via 10.1.12.1, 00:02:50, Serial0/0.21
O N2     11.1.0.0 [110/20] via 10.1.12.1, 00:02:50, Serial0/0.21
O N2     11.1.1.0 [110/20] via 10.1.12.1, 00:02:50, Serial0/0.21
O IA     44.4.4.0 [110/66] via 10.1.23.3, 00:03:52, FastEthernet0/0
```

Task 5

Redistribute Lo1 – 4 interfaces On R5:. Configure R5 such that the routers in this area only maintain and propagate LSA types 1, 2, 3, 7 and a default route.

Since area 2 is a totally stub area, it can have external routes, unless it's converted into a NSSA, but the task states that the routers of this area should also receive a default route. Therefore, the ABR of this area (R4) should be configured to inject a default route using the "default-information-originate", as follows:

On R5:

```
R5(config-if)#route-map TST permit 10
R5(config-route-map)#match interface Lo1 Lo2 Lo3 Lo4

R5(config)#router ospf 1
R5(config-router)#No area 2 stub
R5(config-router)#area 2 nssa
R5(config-router)#redistribute connected subnets route-map TST
```

On R4:

```
R4(config)#router ospf 1
R4(config-router)#No area 2 stub
R4(config-router)#area 2 nssa default-information-originate
```

Note: the "default-information-originate" keyword at the end of "area 2 nssa" command will inject a default route into the area. This default route will be injected as "N2" and not "IA".

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the configuration:

On R5:

```
R5#Show ip route ospf | inc O

O IA     1.1.1.0 [110/131] via 10.1.45.4, 00:06:29, FastEthernet0/0
O IA     2.2.2.0 [110/67] via 10.1.45.4, 00:06:29, FastEthernet0/0
O IA     3.3.3.0 [110/66] via 10.1.45.4, 00:06:29, FastEthernet0/0
O IA     4.4.4.0 [110/2] via 10.1.45.4, 00:06:29, FastEthernet0/0
O IA     22.2.2.0 [110/67] via 10.1.45.4, 00:06:29, FastEthernet0/0
O IA     10.1.12.0 [110/130] via 10.1.45.4, 00:06:29, FastEthernet0/0
O IA     10.1.23.0 [110/66] via 10.1.45.4, 00:06:29, FastEthernet0/0
O IA     10.1.34.0 [110/65] via 10.1.45.4, 00:06:29, FastEthernet0/0
O        44.4.4.0 [110/2] via 10.1.45.4, 00:06:29, FastEthernet0/0
O*N2 0.0.0.0/0 [110/1] via 10.1.45.4, 00:06:29, FastEthernet0/0
```

Task 6

Area 1 should be changed such that it receives and propagates LSA types 1, 2, 7 plus a default route. This area should **not** maintain inter-area routes, but must have the ability to connect to these routes.

On R2:

```
R2(config)#router ospf 1
R2(config-router)#area 1 nssa no-summary
```

The "no-summary" keyword filters the summary LSAs, which are the LSA type 3s, and using this keyword will always inject a default route.

To verify the configuration:

On R1:

```
R1#Show ip route ospf | inc O

O        2.2.2.0 [110/65] via 10.1.12.2, 00:29:15, Serial0/0.12
O*IA 0.0.0.0/0 [110/65] via 10.1.12.2, 00:00:19, Serial0/0.12
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 7

The default route that was injected into area 1 should have a cost of 50.

To see the cost of the default route that was injected:

On R1:

```
R1#Show ip route ospf | inc 0.0.0.0/0

O*IA 0.0.0.0/0 [110/65] via 10.1.12.2, 00:02:53, Serial0/0.12
```

The cost of the default route that is injected by the ABR

On R2:

```
R2(config)#router ospf 1
R2(config-router)#area 1 default-cost 50
```

By default, the cost of the default route that is injected into a given area is 1; this can be verified by "show ip route" command; the output of the above show command reveals that the cost of the injected default route is 65. This is the cost of the link to the ABR (the cost of the frame-relay link) plus 1 (the default cost of the default route).

The default cost of the injected default route can be changed using the "area xx default-cost cc", where cc is the new cost replacing the default value.

To verify the configuration:

On R1:

```
R1#Show ip route ospf | inc 0.0.0.0/0

O*IA 0.0.0.0/0 [110/114] via 10.1.12.2, 00:01:32, Serial0/0.12
```

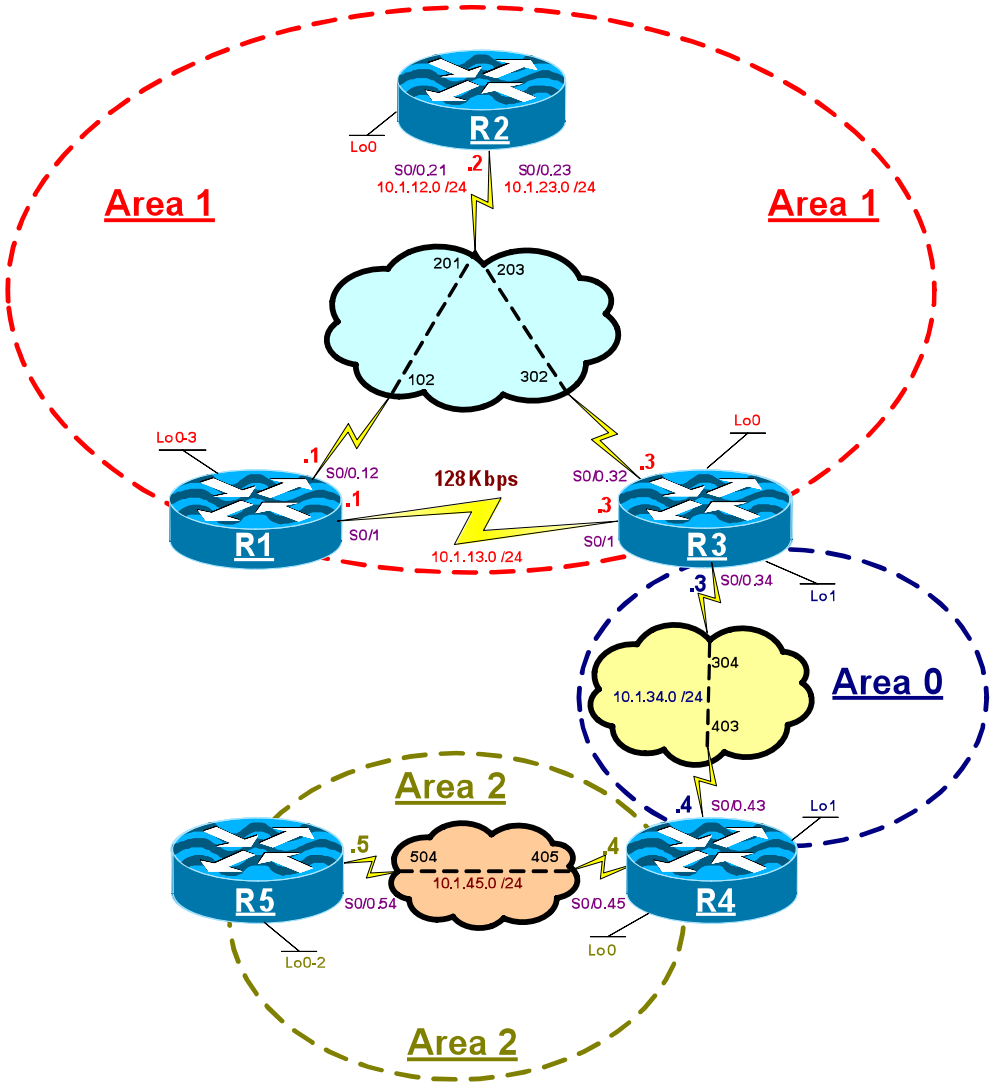
Note: The new cost of 50 + 64 (Which is the cost of the frame-relay link) = 114

Task 8

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 8 – OSPF Filtering



Lab Setup:

- Configure all frame-relay connections in a point-to-point manner.
- Configure the serial interface connecting R1 to R3 as HDLC.
- Use the IP addressing scheme below for IP addressing assignment.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP Addressing scheme:

Routers	Interface / IP Address	Connecting to:
R1	S0/0.12 – 10.1.12.1 /24 S0/1 – 10.1.13.1 /24 Lo0 – 1.1.1.1 /24 Lo1 – 11.1.1.1 /24 Lo2 – 100.1.1.1 /24 Lo3 – 111.1.1.1 /24	R2 R3
R2	S0/0.21 – 10.1.12.2 /24 S0/0.23 – 10.1.23.2 /24 Lo0 – 2.2.2.2 /24	R1 R3
R3	S0/0.32 – 10.1.23.3 /24 S0/0.34 – 10.1.34.3 /24 S0/1 – 10.1.13.3 /24 Lo0 – 3.3.3.3 /24 Lo1 – 30.3.3.3 /24	R2 R4 R1
R4	S0/0.43 – 10.1.34.4 /24 S0/0.45 – 10.1.45.4 /24 Lo0 – 4.4.4.4 /24 Lo1 – 40.4.4.4 /24	R3 R5
R5	S0/0.54 – 10.1.45.5 /24 Lo0 – 5.5.5.5 /24 Lo1 – 50.5.5.5 /24 Lo2 – 55.5.5.5 /24	R4

Task 1

Configure R1 and R2's directly connected interfaces, R3's frame-relay connection to R2 and R3's HDLC connection to R1 in Area 1. Configure the bandwidth of S0/1 interface on R1 and R3 to 128K using the "bandwidth" command. Configure 1.1.1.1, 2.2.2.2, 3.3.3.3, 4.4.4.4 and 5.5.5.5 as the router ID of R1, R2, R3, R4 and R5 respectively.

On R1:

```
R1(config)#int s0/1
R1(config-if)#bandwidth 128

R1(config)#router ospf 1
R1(config-router)#router-id 1.1.1.1
R1(config-router)#netw 0.0.0.0 0.0.0.0 area 1
```

On R2:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config)#router ospf 1
R2(config-router)#router-id 2.2.2.2
R2(config-router)#netw 0.0.0.0 0.0.0.0 area 1
```

On R3:

```
R3(config)#int s0/1
R3(config-if)#bandwidth 128

R3(config)#router ospf 1
R3(config-router)#router-id 3.3.3.3
R3(config-router)#netw 10.1.13.3 0.0.0.0 area 1
R3(config-router)#netw 10.1.23.3 0.0.0.0 area 1
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf | inc O

O       2.2.2.2 [110/65] via 10.1.12.2, 00:01:02, Serial0/0.12
O       10.1.23.0 [110/128] via 10.1.12.2, 00:01:02, Serial0/0.12
```

On R4:

```
R4(config)#router ospf 1
R4(config-router)#router-id 4.4.4.4
```

On R5:

```
R5(config)#router ospf 1
R5(config-router)#router-id 5.5.5.5
```

Task 2

Configure R3's frame-relay connection to R4 and R4's frame-relay connection to R3 in Area 0.

On R3:

```
R3(config)#router ospf 1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-router)#netw 10.1.34.3 0.0.0.0 area 0
```

On R4:

```
R4(config)#router ospf 1
R4(config-router)#netw 10.1.34.4 0.0.0.0 area 0
```

To verify the configuration:

On R4:

```
R4#Show ip route ospf | inc O

O IA    1.1.1.1 [110/193] via 10.1.34.3, 00:00:08, Serial0/0.43
O IA    2.2.2.2 [110/129] via 10.1.34.3, 00:00:08, Serial0/0.43
O IA    100.1.1.1 [110/193] via 10.1.34.3, 00:00:08, Serial0/0.43
O IA    111.1.1.1 [110/193] via 10.1.34.3, 00:00:08, Serial0/0.43
O IA    10.1.13.0 [110/845] via 10.1.34.3, 00:00:08, Serial0/0.43
O IA    10.1.12.0 [110/192] via 10.1.34.3, 00:00:08, Serial0/0.43
O IA    10.1.23.0 [110/128] via 10.1.34.3, 00:00:08, Serial0/0.43
O IA    11.1.1.1 [110/193] via 10.1.34.3, 00:00:08, Serial0/0.43
```

Task 3

Configure R4's frame-relay connection to R5 and R5's frame-relay connection to R4 in Area 2.

On R4:

```
R4(config)#router ospf 1
R4(config-router)#netw 10.1.45.4 0.0.0.0 area 2
```

On R5:

```
R5(config)#router ospf 1
R5(config-router)#netw 10.1.45.5 0.0.0.0 area 2
```

To verify the configuration:

On R5:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R5#Show ip route ospf | inc O

O IA    1.1.1.1 [110/257] via 10.1.45.4, 00:00:05, Serial0/0.54
O IA    2.2.2.2 [110/193] via 10.1.45.4, 00:00:05, Serial0/0.54
O IA    100.1.1.1 [110/257] via 10.1.45.4, 00:00:05, Serial0/0.54
O IA    111.1.1.1 [110/257] via 10.1.45.4, 00:00:05, Serial0/0.54
O IA    10.1.13.0 [110/909] via 10.1.45.4, 00:00:05, Serial0/0.54
O IA    10.1.12.0 [110/256] via 10.1.45.4, 00:00:05, Serial0/0.54
O IA    10.1.23.0 [110/192] via 10.1.45.4, 00:00:05, Serial0/0.54
O IA    10.1.34.0 [110/128] via 10.1.45.4, 00:00:05, Serial0/0.54
O IA    11.1.1.1 [110/257] via 10.1.45.4, 00:00:05, Serial0/0.54
```

Task 4

Configure R2 to filter network 1.1.1.0 /24 from its routing table. You should use R1 as part of the solution to accomplish this task. Ensure that R3, R4 and R5 have full reachability to this network. Use an IP prefix-list to accomplish this task.

On R2:

```
R2#Show ip route 1.1.1.0 255.255.255.0

% Subnet not in table
```

Note this network is not in the routing table of R2, because the loopback interfaces were NOT advertised with their correct mask. To fix this problem, we should advertise all the loopback interfaces on R1 with their correct mask. Sometimes this can be a problem where an unbelievable amount of time is spent looking for a prefix that does not exist.

On R1:

```
R1(config)#int range lo0 - 3
R1(config-if)#ip ospf network point-to-point
```

On R2:

```
R2#Show ip route 1.1.1.0 255.255.255.0

Routing entry for 1.1.1.0/24
  Known via "ospf 1", distance 110, metric 65, type intra area
  Last update from 10.1.12.1 on Serial0/0.21, 00:02:15 ago
  Routing Descriptor Blocks:
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
* 10.1.12.1, from 1.1.1.1, 00:02:15 ago, via Serial0/0.21
    Route metric is 65, traffic share count is 1
```

In this case, the following solution only affects the router that it is configured on.

On R2:

```
R2(config)#ip prefix-list TST seq 5 deny 1.1.1.0/24
R2(config)#ip prefix-list TST seq 10 permit 0.0.0.0/0 LE 32

R2(config)#router ospf 1
R2(config-router)#distribute-list prefix TST in
```

To verify the configuration:

On R2:

```
R2#Show ip route ospf | inc O

O       100.1.1.0 [110/65] via 10.1.12.1, 00:01:47, Serial0/0.21
O       111.1.1.0 [110/65] via 10.1.12.1, 00:01:47, Serial0/0.21
O       10.1.13.0 [110/845] via 10.1.23.3, 00:01:47, Serial0/0.23
O IA    10.1.45.0 [110/192] via 10.1.23.3, 00:01:47, Serial0/0.23
O IA    10.1.34.0 [110/128] via 10.1.23.3, 00:01:47, Serial0/0.23
O       11.1.1.0 [110/65] via 10.1.12.1, 00:01:47, Serial0/0.21
```

Note: The "distribute-list in" router configuration mode command can be used when filtering any type of LSA on a given router. This command only filters the prefix/es from the local router's routing table and not the database. The output of the following "show" command reveals that R3 is learning network 1.1.1.0 /24 through R2, even though this prefix is not in R2's routing table.

On R3:

```
R3#Show ip route ospf | inc O

O       1.1.1.0 [110/129] via 10.1.23.2, 00:10:35, Serial0/0.32
O       100.1.1.0 [110/129] via 10.1.23.2, 00:10:25, Serial0/0.32
O       111.1.1.0 [110/129] via 10.1.23.2, 00:10:15, Serial0/0.32
O       10.1.12.0 [110/128] via 10.1.23.2, 00:26:51, Serial0/0.32
O IA    10.1.45.0 [110/128] via 10.1.34.4, 00:20:55, Serial0/0.34
O       11.1.1.0 [110/129] via 10.1.23.2, 00:10:25, Serial0/0.32
```

Note: R3 sees network 1.1.1.0 /24 through R2; this is because the bandwidth of the S0/1 interface connecting R3 to R1 is 128Kbps, which means that the cost is much higher.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The output of the following command reveals that network 1.1.1.0/24 is in R2's database:

```
R2#Show ip ospf database router | inc _1.1.1.0

      (Link ID) Network/subnet number: 1.1.1.0
```

To test R3's reachability:

On R3:

```
R3#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
U.U.U
Success rate is 0 percent (0/5)
```

Note: R3 sees network 1.1.1.0 /24 through R2, But R2 does not have reachability to network 1.1.1.0/24. Therefore, in order to provide reachability to R3 and the other downstream routers without using a router configuration mode command, we can use PBR to accomplish this task, as follows:

On R3:

```
R3(config)#access-list 100 permit ip any 1.1.1.0 0.0.0.255

R3(config)#route-map TST permit 10
R3(config-route-map)#match ip addr 100
R3(config-route-map)#set ip next-hop 10.1.13.1

R3(config)#route-map TST permit 90

R3(config)#ip local policy route-map TST
```

To test R3's reachability:

On R3:

```
R3#Show ip route ospf | inc O

O       1.1.1.0 [110/129] via 10.1.23.2, 00:12:27, Serial0/0.32
O       2.2.2.2 [110/65] via 10.1.23.2, 00:18:38, Serial0/0.32
O       100.1.1.0 [110/129] via 10.1.23.2, 00:12:27, Serial0/0.32
O       111.1.1.0 [110/129] via 10.1.23.2, 00:12:27, Serial0/0.32
O       10.1.12.0 [110/128] via 10.1.23.2, 00:18:38, Serial0/0.32
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O IA    10.1.45.0 [110/128] via 10.1.34.4, 00:16:59, Serial0/0.34
O       11.1.1.0 [110/129] via 10.1.23.2, 00:12:27, Serial0/0.32

R3#Traceroute 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 10.1.13.1 40 msec *  36 msec
```

Note: Even though the routing table states that 10.1.23.2 is the next hop, PBR is looked at before the routing table, therefore, the next hop is R1 (10.1.13.1).

To test R4's reachability:

On R4:

```
R4#Traceroute 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 10.1.34.3 28 msec 28 msec 28 msec
  2 10.1.23.2 52 msec 56 msec 52 msec
  3 10.1.23.2 !H  *  !H
```

Note: R4 does not have reachability to network 1.1.1.0 /24. This is because when PBR was configured on R3, it was not applied to the interface facing R4. To fix this problem:

On R3:

```
R3(config)#int s0/0.34
R3(config-subif)#ip policy route-map TST
```

To verify the configuration:

On R4:

```
R4#Traceroute 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 10.1.34.3 28 msec 28 msec 28 msec
  2 10.1.13.1 64 msec *  64 msec
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R5:

```
R5#Traceroute 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 10.1.45.4 28 msec 28 msec 28 msec
  2 10.1.34.3 52 msec 56 msec 52 msec
  3 10.1.13.1 92 msec *  92 msec
```

Task 5

Configure filtering on the appropriate router/s such that the existing and future routers in area 2 filter network 11.1.1.0 /24 from their routing table or database.

The following method only works for filtering LSA type 3s and LSA type 3s ONLY. The first step is to configure a prefix-list to deny the route/s and permit everything else:

On R4:

```
R4(config)#ip prefix-list TST seq 5 deny 11.1.1.0/24
R4(config)#ip prefix-list TST seq 10 permit 0.0.0.0/0 LE 32
```

Once the prefix-list is configured, it can be applied to the area that it must be filtered from, in this case area 2. This command must be configured on an ABR. In the following configuration, the "prefix-list" filter's network 11.1.1.0 /24 from getting into area 2.

```
R4(config)#router ospf 1
R4(config-router)#area 2 filter-list prefix TST in
```

To verify the configuration:

On R5:

```
R5#Sh ip route ospf | inc O IA

O IA    1.1.1.0 [110/257] via 10.1.45.4, 00:21:09, Serial0/0.54
O IA    2.2.2.2 [110/193] via 10.1.45.4, 00:25:21, Serial0/0.54
O IA    100.1.1.0 [110/257] via 10.1.45.4, 00:21:09, Serial0/0.54
O IA    111.1.1.0 [110/257] via 10.1.45.4, 00:21:09, Serial0/0.54
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O IA    10.1.13.0 [110/909] via 10.1.45.4, 00:25:21, Serial0/0.54
O IA    10.1.12.0 [110/256] via 10.1.45.4, 00:25:21, Serial0/0.54
O IA    10.1.23.0 [110/192] via 10.1.45.4, 00:25:21, Serial0/0.54
O IA    10.1.34.0 [110/128] via 10.1.45.4, 00:25:21, Serial0/0.54
```

Note: The above "show" command reveals that R5 does not have the route in its routing table, and the following command verifies that R5 does not have the prefix in its database.

```
R5#Show ip ospf database summary 11.1.1.0

            OSPF Router with ID (5.5.5.5) (Process ID 1)
```

On R4:

```
R4#Show ip route ospf | inc O IA

O IA    1.1.1.0 [110/193] via 10.1.34.3, 00:13:32, Serial0/0.43
O IA    100.1.1.0 [110/193] via 10.1.34.3, 00:13:32, Serial0/0.43
O IA    111.1.1.0 [110/193] via 10.1.34.3, 00:13:32, Serial0/0.43
O IA    10.1.13.0 [110/845] via 10.1.34.3, 00:13:32, Serial0/0.43
O IA    10.1.12.0 [110/192] via 10.1.34.3, 00:13:32, Serial0/0.43
O IA    10.1.23.0 [110/128] via 10.1.34.3, 00:13:32, Serial0/0.43
O IA    11.1.1.0 [110/193] via 10.1.34.3, 00:13:32, Serial0/0.43
```

Note: Even though the output of the above "show" command reveals that network 11.1.1.0 /24 is in R4's routing table, the output of the following "show" command clearly shows that it is in the database of area 0 and not in the database that belongs to area 2.

```
R4#Show ip ospf database summary 11.1.1.0

     OSPF Router with ID (4.4.4.4) (Process ID 1)

                Summary Net Link States (Area 0)

  Routing Bit Set on this LSA
  LS age: 267
  Options: (No TOS-capability, DC, Upward)
  LS Type: Summary Links(Network)
  Link State ID: 11.1.1.0 (summary Network Number)
  Advertising Router: 3.3.3.3
  LS Seq Number: 80000001
  Checksum: 0x950C
  Length: 28
  Network Mask: /24
        TOS: 0  Metric: 129
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 6

Configure the appropriate router/s such that the routers in area 0 do not see network 11.1.1.0 /24 in their routing table or link state database. You should use the same solution as the one in the previous task, but it should be implemented in the **out** bound direction.

The following method is used for filtering LSA type 3s, and LSA type 3s only. A prefix-list is configured to deny network 11.1.1.0 /24 on the ABR, but in the following case, the "area filter-list" command is filtering network 11.1.1.0 /24 as it is advertised out of area 1. This means that in the future if another area is configured (Let's say area 3) that is connected to R3; the new area will also be affected.

On R3:

```
R3(config)#ip prefix-list TST seq 5 deny 11.1.1.0/24
R3(config)#ip prefix-list TST seq 10 permit 0.0.0.0/0 LE 32

R3(config)#router ospf 1
R3(config-router)#area 1 filter-list prefix TST OUT
```

To verify the configuration:

Note: The output of the following commands show that prefix 11.1.1.0 /24 is no longer in the routing table of R4 or R5.

On R4:

```
R4#Show ip route 11.1.1.0 255.255.255.0

% Network not in table
```

On R5:

```
R5#Show ip route 11.1.1.0 255.255.255.0

% Network not in table
```

Note: the prefix is still in the routing table of R3, where the filtering is performed, but the prefix is in the routing table of this router as a route from area 1 (LSA type 1) and not a prefix from area 0 (from area 0's perspective this prefix is LSA type 3).

On R3:

```
R3#Show ip route ospf | inc O
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O        1.1.1.0 [110/129] via 10.1.23.2, 00:30:02, Serial0/0.32
O        2.2.2.2 [110/65] via 10.1.23.2, 00:36:14, Serial0/0.32
O        100.1.1.0 [110/129] via 10.1.23.2, 00:30:02, Serial0/0.32
O        111.1.1.0 [110/129] via 10.1.23.2, 00:30:02, Serial0/0.32
O        10.1.12.0 [110/128] via 10.1.23.2, 00:36:14, Serial0/0.32
O IA     10.1.45.0 [110/128] via 10.1.34.4, 00:34:34, Serial0/0.34
O        11.1.1.0 [110/129] via 10.1.23.2, 00:30:02, Serial0/0.32
```



Note this is an intra-area route.

To prove this further:

On R3:

```
R3#Sh ip ospf database summary 11.1.1.0

      OSPF Router with ID (3.3.3.3) (Process ID 1)
```

Note: The output of the above "show" command reveals that network 11.1.1.0 /24 is not in area 0. If it was in area 0, it would have been in the link state database of this router as a summary LSA or LSA type 3, whereas, the following "show" command reveals that the prefix is in area 1 as a router LSA or LSA type 1.

```
R3#Show ip ospf database router | inc area 1 | 11.1.1.0

      (Link ID) Network/subnet number: 11.1.1.0
```

Task 7

Configure the appropriate router/s to filter network 111.1.1.0/24 such that the routers in area 0 or area 2 **do not** have this network in their routing table or database. Use the minimum number of commands to accomplish this task.

Note: The output of the following "show" command verifies that network 111.1.1.0 /24 is in the database of R3 that belongs to area 1, and it shows up as a router LSA or LSA type 1:

On R3:

```
R3#Show ip ospf database router | inc area 1 | 111.1.1.0

                Router Link States (Area 1)
     (Link ID) Network/subnet number: 111.1.1.0
```


©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

The output of the following "show" command reveals that prefix 111.1.1.0 /24 is also in the database of R3 as a summary LSA or LSA type 3:

```
R3#Show ip ospf database summary 111.1.1.0

            OSPF Router with ID (3.3.3.3) (Process ID 1)

                Summary Net Link States (Area 0)

  LS age: 294
  Options: (No TOS-capability, DC, Upward)
  LS Type: Summary Links(Network)
  Link State ID: 111.1.1.0 (summary Network Number)
  Advertising Router: 3.3.3.3
  LS Seq Number: 8000000A
  Checksum: 0x6CC8
  Length: 28
  Network Mask: /24
        TOS: 0  Metric: 129
```

The following command reveals that prefix 111.1.1.0 /24 is in the routing table of R3 as an intra-area route. Remember that intra-area routes take precedence over inter-area routes; therefore, this prefix shows up as an "O" route in the routing table.

```
R3#show ip route ospf | inc 111.1.1.0

O       111.1.1.0 [110/129] via 10.1.23.2, 00:47:21, Serial0/0.32
```

The following OSPF filtering mechanism works only on LSA type 1s and 2s. It filters LSA type 1s or 2s from being injected into a given area, this command only works if it is configured on an ABR and it is used for filtering LSA type 1 & 2s ONLY.

<u>On R3:</u>

```
R3(config)#router ospf 1
R3(config-router)#area 1 range 111.1.1.0 255.255.255.0 not-advertise
```

Note: The prefix is still in the routing table of R3 where the filtering is performed, but once again it shows up in the routing table as LSA type 1:

```
R3#Show ip route ospf | inc 111.1.1.0

O       111.1.1.0 [110/129] via 10.1.23.2, 00:00:34, Serial0/0.32
```

The prefix is no longer in the database of R3 as LSA type 3, which means that the routers in area 0 or any other area downstream to area 0 will not have this prefix in their routing table or link state

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

database.

```
R3#Show ip ospf database summary 111.1.1.0

            OSPF Router with ID (3.3.3.3) (Process ID 1)

R3#Show ip ospf database router | inc area 1 | 111.1.1.0

                Router Link States (Area 1)
     (Link ID) Network/subnet number: 111.1.1.0
```

On R4: or R5

```
Rx#Show ip route 111.1.1.0

% Network not in table
```

Task 8

Configure the appropriate router/s such that none of the routers except R1 can see network 100.1.1.0 /24 in their routing table; **do not** stop advertising this network to accomplish this task. You should **not** use the solution that was used in tasks 6, or 7 to accomplish this task.

On R2:

You should always display the existing access-list/s and distribute-list/s before configuring one. You do not want to include additional networks in an existing access-list/distribute-list few minutes before the end of your lab exam.

```
R2#Sh ip prefix-list

ip prefix-list TST: 2 entries
   seq 5 deny 1.1.1.0/24
   seq 10 permit 0.0.0.0/0 le 32

R2#Show run | s router

router ospf 1
 router-id 2.2.2.2
 log-adjacency-changes
 network 0.0.0.0 255.255.255.255 area 1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
distribute-list prefix TST in
```

Note: The above "show" commands verify that there is already a distribute-list configured in the router configuration mode, therefore, we should try to modify the existing prefix-list that is applied by the existing distribute-list.

```
R2(config)#ip prefix-list TST seq 7 deny 100.1.1.0/24
```

To see the prefix-list after the above addition:

On R2:

```
R2#Show run | inc ip pref

ip prefix-list TST seq 5 deny 1.1.1.0/24
ip prefix-list TST seq 7 deny 100.1.1.0/24
ip prefix-list TST seq 10 permit 0.0.0.0/0 le 32
```

To verify the configuration:

On R2:

```
R2#Show ip route 100.1.1.0

% Network not in table
```

Let's check R3:

On R3:

```
R3#Show access-list

Extended IP access list 100
    10 permit ip any 1.1.1.0 0.0.0.255 (40 matches)

R3#Show ip route | inc 100.1.1.0

O       100.1.1.0 [110/129] via 10.1.23.2, 00:16:48, Serial0/0.32
```

To configure R3:

```
R3(config)#ip prefix-list NET seq 5 deny 100.1.1.0/24
R3(config)#ip prefix-list NET seq 10 permit 0.0.0.0/0 LE 32

R3(config)#router ospf 1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-router)#distribute-list prefix NET in
```

To verify the configuration:

On R3:

```
R3#Show ip route ospf | inc 100.1.1.0
R3#
```

On R4:

On R4 the distance command can be used. Referencing network 100.1.1.0 /24 in access-list 1, and the distance command sets the AD of the network referenced in access-list 1 to 255 sourcing from R2 (3.3.3.3 0.0.0.0); AD of 255 is unreachable and won't be injected into the routing table:

```
R4(config)#access-list 1 permit 100.1.1.0 0.0.0.255

R4(config)#router ospf 1
R4(config-router)#distance 255 3.3.3.3 0.0.0.0 1
```

On R5:

```
R5#Show ip route 100.1.1.0 255.255.255.0

% Network not in table

R5#Show ip ospf da summ 100.1.1.0

        OSPF Router with ID (5.5.5.5) (Process ID 1)
```

****

Note: The above solution utilizes the behavior of OSPF when prefixes are advertised from area 0 into non-zero areas. When LSA type-3s are advertised from area 0 into other areas, the behavior is like an internal redistribution. Basically if R4 does not have the route in its routing table, it will not redistribute and as a result of this, R5 or any other router in area 2 will not see the prefixe/s in their routing table or database.

****

Task 9

Configure R5 to redistributel 0, 1 and 2 interfaces in OSPF routing domain using the default cost.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R5:

```
R5(config)#route-map TST permit 10
R5(config-route-map)#match interface lo0 lo1 lo2

R5(config)#router ospf 1
R5(config-router)#redistribute connected subnets route-map TST
```

To verify the configuration:

On R4:

```
R4#Show ip route ospf | inc E2

O E2     50.5.5.0 [110/20] via 10.1.45.5, 00:01:05, Serial0/0.45
O E2     55.5.5.0 [110/20] via 10.1.45.5, 00:01:05, Serial0/0.45
O E2     5.5.5.0  [110/20] via 10.1.45.5, 00:01:05, Serial0/0.45
```

Task 10

Configure the appropriate router such that none of the routers except R5 can see network 5.5.5.0 /24 in their routing table.

On R5:

```
R5(config)#ip prefix-list NET seq 5 deny 5.5.5.0/24
R5(config)#ip prefix-list NET seq 10 permit 0.0.0.0/0 LE 32

R5(config)#router ospf 1
R5(config-router)#distribute-list prefix NET OUT
```

To verify the configuration:

On R4:

```
R4#Show ip route ospf | inc E2

O E2     50.5.5.0 [110/20] via 10.1.45.5, 00:06:00, Serial0/0.45
O E2     55.5.5.0 [110/20] via 10.1.45.5, 00:06:00, Serial0/0.45
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1:

```
R1#Show ip route ospf | inc E2

O E2     50.5.5.0 [110/20] via 10.1.12.2, 00:07:08, Serial0/0.12
O E2     55.5.5.0 [110/20] via 10.1.12.2, 00:07:08, Serial0/0.12
```

Note: this is the only scenario where the "distribute-list OUT" command works in OSPF. This command must be configured on the ASBR or else it will not have any effect whatsoever. This command filters LSA type 5s or 7s; in this case, the specific LSA type 5 is filtered from R5's OSPF database and as a result of that. None of the other OSPF routers will see the route in their routing table or database.

```
R5#Show ip ospf da ext | inc 5

            OSPF Router with ID (5.5.5.5) (Process ID 1)
                Type-5 AS External Link States
  Link State ID: 50.5.5.0 (External Network Number )
  Advertising Router: 5.5.5.5
  Checksum: 0x51FE
  Link State ID: 55.5.5.0 (External Network Number )
  Advertising Router: 5.5.5.5
```

Task 11

Configure the appropriate router such that none of the routers except R5 can see network 50.5.5.0 /24 in their routing table or database. You should **not** use the solution that was implemented in the previous task.

The following command is used to filter LSA type 5s or 7s, this command must be configured on an ASBR or the router that originated LSA type 5, and when configured, it filters the specified prefix from the OSPF Link state database of the ASBR.

On R5:

```
R5(config)#router ospf 1
R5(config-router)#summary-address 50.5.5.0 255.255.255.0 not-advertise
```

The following show command reveals that network 50.5.5.0 /24 is not in the link state database of R5.

```
R5#Show ip ospf da ext | inc 5
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
OSPF Router with ID (5.5.5.5) (Process ID 1)
                Type-5 AS External Link States
  LS age: 534
  Link State ID: 55.5.5.0 (External Network Number )
  Advertising Router: 5.5.5.5
```

To verify the configuration:

On R1:

```
R1#Sh ip route ospf | inc E2

O E2    55.5.5.0 [110/20] via 10.1.12.2, 00:21:25, Serial0/0.12
```

On R4:

```
R4#Show ip route ospf | inc E2

O E2    55.5.5.0 [110/20] via 10.1.45.5, 00:23:17, Serial0/0.45
```

Task 12

Configure the appropriate router such that R1 does **not** have network 55.5.5.0 /24 in its routing table.

Note: R1 has the network in its routing table.

On R1:

```
R1#Show ip route ospf | inc E2

O E2    55.5.5.0 [110/20] via 10.1.12.2, 00:30:59, Serial0/0.12
```

Note: There are no access-lists or ip prefix-lists configured on this router:

```
R1(config)#ip prefix-list NET seq 5 deny 55.5.5.0/24
R1(config)#ip prefix-list NET seq 10 permit 0.0.0.0/0 LE 32

R1(config)#router ospf 1
R1(config-router)#distribute-list prefix NET in
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the configuration:

On R1:

```
R1#Show ip route ospf | inc E2
R1#
```

Note: The above "show" command verifies that prefix 55.5.5.0 /24 was filtered successfully.

Task 13

Remove any PBR configuration and all the filters applied in the previous tasks, if this configuration is performed successfully. All the routers should have every route advertised and redistributed in this lab.

On R1:

```
R1(config)#No ip prefix NET

R1(config)#router ospf 1
R1(config-router)#No distribute-list prefix NET in
```

On R2:

```
R2(config)#No ip prefix-list TST

R2(config)#router ospf 1
R2(config-router)#No distribute-list prefix TST in
```

On R3:

```
R3(config)#No route-map TST

R3(config)#No ip prefix-list NET
R3(config)#No ip prefix-list TST

R3(config)#No access-list 100

R3(config)#No ip local policy route-map TST

R3(config)#router ospf 1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-router)#No area 1 range 111.1.1.0 255.255.255.0 not-advertise
R3(config-router)#No distribute-list prefix NET in
R3(config-router)#No area 1 filter-list prefix TST out

R3(config-router)#int s0/0.34
R3(config-subif)#No ip policy route-map TST
```

On R4:

```
R4(config)#No ip prefix-list TST
R4(config)#No access-list 1

R4(config)#router ospf 1
R4(config-router)#No distance 255 3.3.3.3 0.0.0.0 1
R4(config-router)#No area 0 filter-list prefix TST out
R4(config-router)#No area 2 filter-list prefix TST in
```

On R5:

```
R5(config)#No ip prefix-list NET

R5(config)#router ospf 1
R5(config-router)#No distribute-list prefix NET out
R5(config-router)#No summary-address 50.5.5.0 255.255.255.0 not-advertise
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf | inc O

O E2    50.5.5.0 [110/20] via 10.1.12.2, 00:00:21, Serial0/0.12
O       2.2.2.2 [110/65] via 10.1.12.2, 00:13:54, Serial0/0.12
O E2    55.5.5.0 [110/20] via 10.1.12.2, 00:00:21, Serial0/0.12
O E2    5.5.5.0 [110/20] via 10.1.12.2, 00:00:21, Serial0/0.12
O       10.1.23.0 [110/128] via 10.1.12.2, 00:13:54, Serial0/0.12
O IA    10.1.45.0 [110/256] via 10.1.12.2, 00:13:54, Serial0/0.12
O IA    10.1.34.0 [110/192] via 10.1.12.2, 00:13:54, Serial0/0.12
```

On R5:

```
R5#Show ip route ospf | inc O

O IA    1.1.1.0 [110/257] via 10.1.45.4, 00:02:16, Serial0/0.54
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O IA    2.2.2.2 [110/193] via 10.1.45.4, 00:02:16, Serial0/0.54
O IA    100.1.1.0 [110/257] via 10.1.45.4, 00:02:16, Serial0/0.54
O IA    111.1.1.0 [110/257] via 10.1.45.4, 00:02:16, Serial0/0.54
O IA    10.1.13.0 [110/909] via 10.1.45.4, 00:02:16, Serial0/0.54
O IA    10.1.12.0 [110/256] via 10.1.45.4, 00:02:16, Serial0/0.54
O IA    10.1.23.0 [110/192] via 10.1.45.4, 00:02:16, Serial0/0.54
O IA    10.1.34.0 [110/128] via 10.1.45.4, 02:28:25, Serial0/0.54
O IA    11.1.1.0 [110/257] via 10.1.45.4, 00:02:16, Serial0/0.54
```

On R3:

```
R3#Show ip route ospf | inc O

O       1.1.1.0 [110/129] via 10.1.23.2, 00:09:26, Serial0/0.32
O E2    50.5.5.0 [110/20] via 10.1.34.4, 00:01:15, Serial0/0.34
O       2.2.2.2 [110/65] via 10.1.23.2, 00:09:26, Serial0/0.32
O       100.1.1.0 [110/129] via 10.1.23.2, 00:09:26, Serial0/0.32
O E2    55.5.5.0 [110/20] via 10.1.34.4, 00:01:15, Serial0/0.34
O E2    5.5.5.0 [110/20] via 10.1.34.4, 00:01:15, Serial0/0.34
O       111.1.1.0 [110/129] via 10.1.23.2, 00:09:26, Serial0/0.32
O       10.1.12.0 [110/128] via 10.1.23.2, 00:09:26, Serial0/0.32
O IA    10.1.45.0 [110/128] via 10.1.34.4, 00:09:26, Serial0/0.34
O       11.1.1.0 [110/129] via 10.1.23.2, 00:09:26, Serial0/0.32
```

Task 14

Advertise Lo0 interface of R3 in area 1, Lo1 interface of R3 and R4 in area 0, and Lo0 interface of R4 in area 2. These loopback interfaces and the Lo0 of R2 **must** be advertised with their correct mask.

On R2:

```
R2(config)#int lo0
R2(config-if)#ip ospf net point-to-point
```

On R3:

```
R3(config)#int lo0
R3(config-if)#ip ospf net point-to-point

R3(config-if)#int lo1
R3(config-if)#ip ospf net point-to-point
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-if)#router ospf 1
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 1
R3(config-router)#netw 30.3.3.3 0.0.0.0 area 0
```

On R4:

```
R4(config)#int lo0
R4(config-if)#ip ospf net point-to-point

R4(config-if)#int lo1
R4(config-if)#ip ospf net point-to-point

R4(config-if)#router ospf 1
R4(config-router)#netw 4.4.4.4 0.0.0.0 area 2
R4(config-router)#netw 40.4.4.4 0.0.0.0 area 0
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf | inc O

O E2     50.5.5.0 [110/20] via 10.1.12.2, 00:09:44, Serial0/0.12
O        2.2.2.0 [110/65] via 10.1.12.2, 00:01:26, Serial0/0.12
O        3.3.3.0 [110/129] via 10.1.12.2, 00:01:02, Serial0/0.12
O IA     4.4.4.0 [110/193] via 10.1.12.2, 00:00:31, Serial0/0.12
O E2     55.5.5.0 [110/20] via 10.1.12.2, 00:09:44, Serial0/0.12
O E2     5.5.5.0 [110/20] via 10.1.12.2, 00:09:44, Serial0/0.12
O IA     40.4.4.0 [110/193] via 10.1.12.2, 00:00:21, Serial0/0.12
O        10.1.23.0 [110/128] via 10.1.12.2, 00:23:17, Serial0/0.12
O IA     10.1.45.0 [110/256] via 10.1.12.2, 00:23:17, Serial0/0.12
O IA     10.1.34.0 [110/192] via 10.1.12.2, 00:23:17, Serial0/0.12
O        30.3.3.0 [110/129] via 10.1.12.2, 00:00:52, Serial0/0.12
```

Task 15

Configure the appropriate router/s such that the routers in area 2 **do not** see any of the networks advertised by any of the routers in this topology. Routers R1, R2 and R3 and R4 should see all the networks advertised by the existing and future router/s in area 2.

By default, all outgoing LSAs are flooded to the interface. The following configuration prevents

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

flooding of all OSPF LSAs out of a given interface, in this case R4's S0/0.45 interface.

On R4:

```
R4(config)#int s0/0.45
R4(config-subif)#ip ospf database-filter all out
```

For this filtering mechanism to work, the OSPF process must be cleared.

On R5:

```
R5#clear ip ospf proc
Reset ALL OSPF processes? [no]: y
```

Note: R4 and R5 are still maintaining their neighbor adjacency

```
R5#sh ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
4.4.4.4           0   FULL/  -        00:00:39    10.1.45.4       Serial0/0.54
```

You can see that R5 does not have any of the routes from the other routers; this includes R4's routes, which are advertised within the same area. Remember that the filtering is done on S0/0.45, meaning that all LSAs (regardless of their area) are filtered.

On R5:

```
R5#Show ip route ospf | inc O
R5#
```

Note: R1 has all the routes including the ones advertised by R5.

On R1:

```
R1#Show ip route ospf | inc O

O E2     50.5.5.0 [110/20] via 10.1.12.2, 00:03:25, Serial0/0.12
O        2.2.2.0 [110/65] via 10.1.12.2, 00:09:21, Serial0/0.12
O        3.3.3.0 [110/129] via 10.1.12.2, 00:08:57, Serial0/0.12
O IA     4.4.4.0 [110/193] via 10.1.12.2, 00:08:26, Serial0/0.12
O E2     55.5.5.0 [110/20] via 10.1.12.2, 00:03:25, Serial0/0.12
O E2     5.5.5.0 [110/20] via 10.1.12.2, 00:03:25, Serial0/0.12
O IA     40.4.4.0 [110/193] via 10.1.12.2, 00:08:16, Serial0/0.12
O        10.1.23.0 [110/128] via 10.1.12.2, 00:31:12, Serial0/0.12
O IA     10.1.45.0 [110/256] via 10.1.12.2, 00:31:12, Serial0/0.12
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Surinderjeet Singhh

On R4:

```
R4#Clear ip ospf proc
Reset ALL OSPF processes? [no]: y
```

To verify the configuration:

On R4:

```
R4#Sh ip rou ospf | inc O

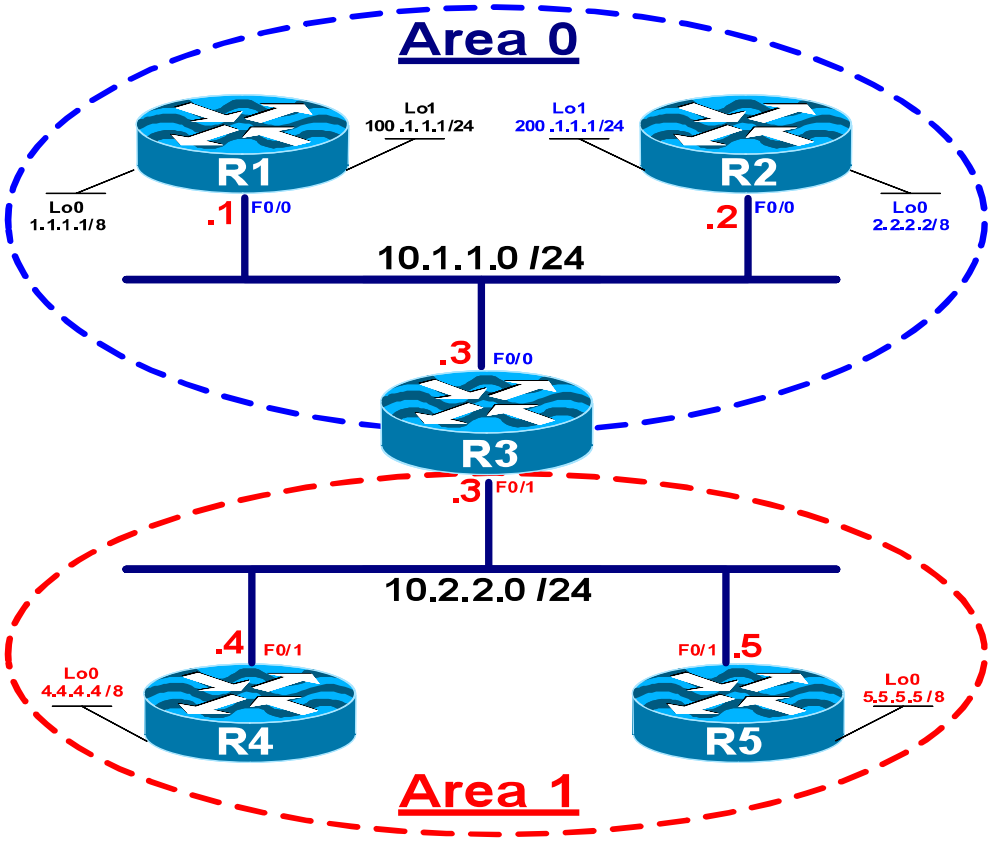
O E2     50.5.5.0 [110/20] via 10.1.45.5, 00:00:03, Serial0/0.45
O E2     55.5.5.0 [110/20] via 10.1.45.5, 00:00:03, Serial0/0.45
O E2     5.5.5.0 [110/20] via 10.1.45.5, 00:00:03, Serial0/0.45
```

Task 17

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 9
Additional OSPF Filtering



Lab Setup:

- Configure the F0/0 interface of R1, R2, and R3 in VLAN 100.
- Configure the F0/1 interface of R3, R4, and R5 in VLAN 200.
- Configure the IP addressing based on the above diagram.

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Task 1

Configure the F0/0 interface of R1, R2 and R3 in OSPF area 0. R1, R2 and R3 should be configured with 1.1.1.1, 2.2.2.2 and 3.3.3.3 as their router ids respectively.

On R1:

```
R1(config)#router ospf 1
R1(config-router)#router-id 1.1.1.1
R1(config-router)#netw 10.1.1.1 0.0.0.0 area 0
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#router-id 2.2.2.2
R2(config-router)#netw 10.1.1.2 0.0.0.0 area 0
```

On R3:

```
R3(config)#router ospf 1
R3(config-router)#router-id 3.3.3.3
R3(config-router)#netw 10.1.1.3 0.0.0.0 area 0
```

To verify the configuration:

On R1:

```
R1#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
2.2.2.2           1   FULL/BDR        00:00:31    10.1.1.2        FastEthernet0/0
3.3.3.3           1   FULL/DR         00:00:30    10.1.1.3        FastEthernet0/0
```

Task 2

Configure the F0/1 interface of R3, R4 and R5 in OSPF area 1. R4 should use 4.4.4.4 and R5 should use 5.5.5.5 as their router-ids.

On R3:

```
R3(config)#router ospf 1
R3(config-router)#netw 10.2.2.3 0.0.0.0 area 1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4:

```
R4(config)#router ospf 1
R4(config-router)#router-id 4.4.4.4
R4(config-router)#netw 10.2.2.4 0.0.0.0 area 1
```

On R5:

```
R5(config)#router ospf 1
R5(config-router)#router-id 5.5.5.5
R5(config-router)#netw 10.2.2.5 0.0.0.0 area 1
```

To verify the configuration:

On R4:

```
R4#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
3.3.3.3           1   FULL/DROTHER    00:00:37    10.2.2.3        FastEthernet0/1
5.5.5.5           1   FULL/DR         00:00:37    10.2.2.5        FastEthernet0/1
```

Task 3

Configure the Lo0 interface of R1 in OSPF Area0, and redistribute the Lo1 interface of this router as external type 1. Lo0 interface of R1 should be advertised with its correct mask.

On R1:

```
R1(config)#int lo0
R1(config-if)#ip ospf netw point-to-point

R1(config)#route-map TST permit 10
R1(config-route-map)#match int lo1

R1(config)#router ospf 1
R1(config-router)#netw 1.1.1.1 0.0.0.0 area 0
R1(config-router)#redistribute connected route-map TST metric-type 1 sub
```

To verify the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3:

```
R3#Show ip route ospf | inc O

O    1.0.0.0/8 [110/2] via 10.1.1.1, 00:01:03, FastEthernet0/0
O E1    100.1.1.0 [110/21] via 10.1.1.1, 00:00:47, FastEthernet0/0
```

Task 4

Configure the Lo0 interface of R2 in OSPF area 0 and redistribute the Lo1 interface of this router as external type 2. Lo0 interface of R2 should be advertised with its correct mask.

On R2:

```
R2(config)#int lo0
R2(config-if)#ip ospf netw point-to-point

R2(config)#route-map TST permit 10
R2(config-route-map)#match interface lo1

R2(config)#router ospf 1
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 0
R2(config-router)#redistribute connected route-map TST sub
```

To verify the configuration:

On R3:

```
R3#Show ip route ospf | inc O

O    1.0.0.0/8 [110/2] via 10.1.1.1, 00:04:09, FastEthernet0/0
O    2.0.0.0/8 [110/2] via 10.1.1.2, 00:01:01, FastEthernet0/0
O E1    100.1.1.0 [110/21] via 10.1.1.1, 00:03:53, FastEthernet0/0
O E2 200.1.1.0/24 [110/20] via 10.1.1.2, 00:00:51, FastEthernet0/0
```

Task 5

Lo0 interface of R4 should be redistributed and tagged with a value of 44, whereas, the Lo0 interface of R5 should be redistributed and tagged with a value of 55.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4:

```
R4(config)#route-map TST permit 10
R4(config-route-map)#match interface lo0
R4(config-route-map)#set tag 44

R4(config)#router ospf 1
R4(config-router)#redistribute connected route-map TST subnets
```

To verify the configuration:

On R5:

```
R5#Show ip ospf da external 4.0.0.0

            OSPF Router with ID (5.5.5.5) (Process ID 1)

                Type-5 AS External Link States

  Routing Bit Set on this LSA
  LS age: 157
  Options: (No TOS-capability, DC)
  LS Type: AS External Link
  Link State ID: 4.0.0.0 (External Network Number )
  Advertising Router: 4.4.4.4
  LS Seq Number: 80000001
  Checksum: 0x560A
  Length: 36
  Network Mask: /8
        Metric Type: 2 (Larger than any link state path)
        TOS: 0
        Metric: 20
        Forward Address: 0.0.0.0
        External Route Tag: 44

R5#Show ip route ospf | inc O

O IA 1.0.0.0/8 [110/3] via 10.2.2.3, 00:07:52, FastEthernet0/1
O    2.0.0.0/8 [110/3] via 10.2.2.3, 00:04:44, FastEthernet0/1
O E1    100.1.1.0 [110/22] via 10.2.2.3, 00:07:31, FastEthernet0/1
O E2 4.0.0.0/8 [110/20] via 10.2.2.4, 00:01:05, FastEthernet0/1
O E2 200.1.1.0/24 [110/20] via 10.2.2.3, 00:04:29, FastEthernet0/1
O IA    10.1.1.0 [110/2] via 10.2.2.3, 00:10:38, FastEthernet0/1
```

On R5:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R5(config)#route-map TST permit 10
R5(config-route-map)#match interface lo0
R5(config-route-map)#set tag 55

R5(config)#router ospf 1
R5(config-router)#redistribute connected route-map TST sub
```

To verify the configuration:

On R3:

```
R3#Show ip ospf da exter 5.0.0.0

            OSPF Router with ID (3.3.3.3) (Process ID 1)

                Type-5 AS External Link States

  Routing Bit Set on this LSA
  LS age: 161
  Options: (No TOS-capability, DC)
  LS Type: AS External Link
  Link State ID: 5.0.0.0 (External Network Number )
  Advertising Router: 5.5.5.5
  LS Seq Number: 80000001
  Checksum: 0xF15E
  Length: 36
  Network Mask: /8
        Metric Type: 2 (Larger than any link state path)
        TOS: 0
        Metric: 20
        Forward Address: 0.0.0.0
        External Route Tag: 55

R3#Show ip route ospf | inc O

O    1.0.0.0/8 [110/2] via 10.1.1.1, 00:09:55, FastEthernet0/0
O IA 2.0.0.0/8 [110/2] via 10.1.1.2, 00:06:47, FastEthernet0/0
O E1    100.1.1.0 [110/21] via 10.1.1.1, 00:09:39, FastEthernet0/0
O E2 4.0.0.0/8 [110/20] via 10.2.2.4, 00:03:07, FastEthernet0/1
O E2 5.0.0.0/8 [110/20] via 10.2.2.5, 00:00:39, FastEthernet0/1
O E2 200.1.1.0/24 [110/20] via 10.1.1.2, 00:06:37, FastEthernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 6

Configure R1 to filter existing and future prefixes that have a tag of 44. DO NOT use an access-list or a prefix-list to accomplish this task.

Let's check the routing table of R1:

On R1:

```
R1#Show ip route ospf | inc O

O    2.0.0.0/8 [110/2] via 10.1.1.2, 00:07:53, FastEthernet0/0
O E2 4.0.0.0/8 [110/20] via 10.1.1.3, 00:04:09, FastEthernet0/0
O E2 5.0.0.0/8 [110/20] via 10.1.1.3, 00:01:41, FastEthernet0/0
O E2 200.1.1.0/24 [110/20] via 10.1.1.2, 00:07:43, FastEthernet0/0
O IA    10.2.2.0 [110/2] via 10.1.1.3, 00:14:57, FastEthernet0/0
```

Let's configure R1 to filter network 4.0.0.0/8 based on its tag value:

```
R1(config)#route-map T44 deny 10
R1(config-route-map)#match tag 44

R1(config)#route-map T44 permit 90

R1(config)#router ospf 1
R1(config-router)#distribute-list route-map T44 in
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf | inc O

O    2.0.0.0/8 [110/2] via 10.1.1.2, 00:00:30, FastEthernet0/0
O E2 5.0.0.0/8 [110/20] via 10.1.1.3, 00:00:30, FastEthernet0/0
O E2 200.1.1.0/24 [110/20] via 10.1.1.2, 00:00:30, FastEthernet0/0
O IA    10.2.2.0 [110/2] via 10.1.1.3, 00:00:30, FastEthernet0/0
```

Task 7

Configure R2 to filter all networks that carry a tag of 55.

On R2:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2#Show ip route ospf | inc O

O    1.0.0.0/8 [110/2] via 10.1.1.1, 00:11:43, FastEthernet0/0
O E1    100.1.1.0 [110/21] via 10.1.1.1, 00:11:43, FastEthernet0/0
O E2 4.0.0.0/8 [110/20] via 10.1.1.3, 00:07:53, FastEthernet0/0
O E2 5.0.0.0/8 [110/20] via 10.1.1.3, 00:05:25, FastEthernet0/0
O IA    10.2.2.0 [110/2] via 10.1.1.3, 00:11:43, FastEthernet0/0
```

To configure the task:

```
R2(config)#route-map T55 deny 10
R2(config-route-map)#match tag 55

R2(config)#route-map T55 permit 90

R2(config)#router ospf 1
R2(config-router)#distribute-list route-map T55 in
```

To verify the configuration:

On R2:

```
R2#Show ip route ospf | inc O

O       1.1.1.0 [110/2] via 10.1.1.1, 00:01:14, FastEthernet0/0
O E1    100.1.1.0 [110/21] via 10.1.1.1, 00:01:14, FastEthernet0/0
O E2 4.0.0.0/8 [110/20] via 10.1.1.3, 00:01:14, FastEthernet0/0
O IA    10.2.2.0 [110/2] via 10.1.1.3, 00:01:14, FastEthernet0/0
```

Task 8

Configure R4 to filter all existing and future OSPF "E1" routes. **Do not** use a prefix-list or an access-list to accomplish this task.

On R4:

```
R4#Show ip route ospf | inc O

O IA 1.0.0.0/8 [110/3] via 10.2.2.3, 00:17:24, FastEthernet0/1
O    2.0.0.0/8 [110/3] via 10.2.2.3, 00:14:16, FastEthernet0/1
O E1    100.1.1.0 [110/22] via 10.2.2.3, 00:17:03, FastEthernet0/1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O E2 5.0.0.0/8 [110/20] via 10.2.2.5, 00:08:09, FastEthernet0/1
O E2 200.1.1.0/24 [110/20] via 10.2.2.3, 00:14:01, FastEthernet0/1
O IA    10.1.1.0 [110/2] via 10.2.2.3, 00:20:10, FastEthernet0/1

R4(config)#route-map E1 deny 10
R4(config-route-map)#match route-type external type-1

R4(config)#route-map E1 permit 90

R4(config)#router ospf 1
R4(config-router)#distribute-list route-map E1 in
```

To verify the configuration:

On R4:

```
R4#Show ip route ospf | inc O

O IA 1.0.0.0/8 [110/3] via 10.2.2.3, 00:00:05, FastEthernet0/1
O    2.0.0.0/8 [110/3] via 10.2.2.3, 00:00:05, FastEthernet0/1
O E2 5.0.0.0/8 [110/20] via 10.2.2.5, 00:00:05, FastEthernet0/1
O E2 200.1.1.0/24 [110/20] via 10.2.2.3, 00:00:05, FastEthernet0/1
O IA    10.1.1.0 [110/2] via 10.2.2.3, 00:00:05, FastEthernet0/1
```

Task 9

Configure R5 to filter all existing and future OSPF "E2" routes. **Do not** use a prefix-list or an access-list to accomplish this task.

On R5:

```
R5#Show ip route ospf | inc O

O IA 1.0.0.0/8 [110/3] via 10.2.2.3, 00:20:06, FastEthernet0/1
O    2.0.0.0/8 [110/3] via 10.2.2.3, 00:16:58, FastEthernet0/1
O E1    100.1.1.0 [110/22] via 10.2.2.3, 00:19:45, FastEthernet0/1
O E2 4.0.0.0/8 [110/20] via 10.2.2.4, 00:13:19, FastEthernet0/1
O E2 200.1.1.0/24 [110/20] via 10.2.2.3, 00:16:43, FastEthernet0/1
O IA    10.1.1.0 [110/2] via 10.2.2.3, 00:22:52, FastEthernet0/1

R5(config)#route-map E2 deny 10
R5(config-route-map)#match route-type external type-2
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R5(config)#route-map E2 permit 90

R5(config)#router ospf 1
R5(config-router)#distribute-list route-map E2 in
```

To verify the configuration:

On R5:

```
R5#Show ip route ospf | inc O

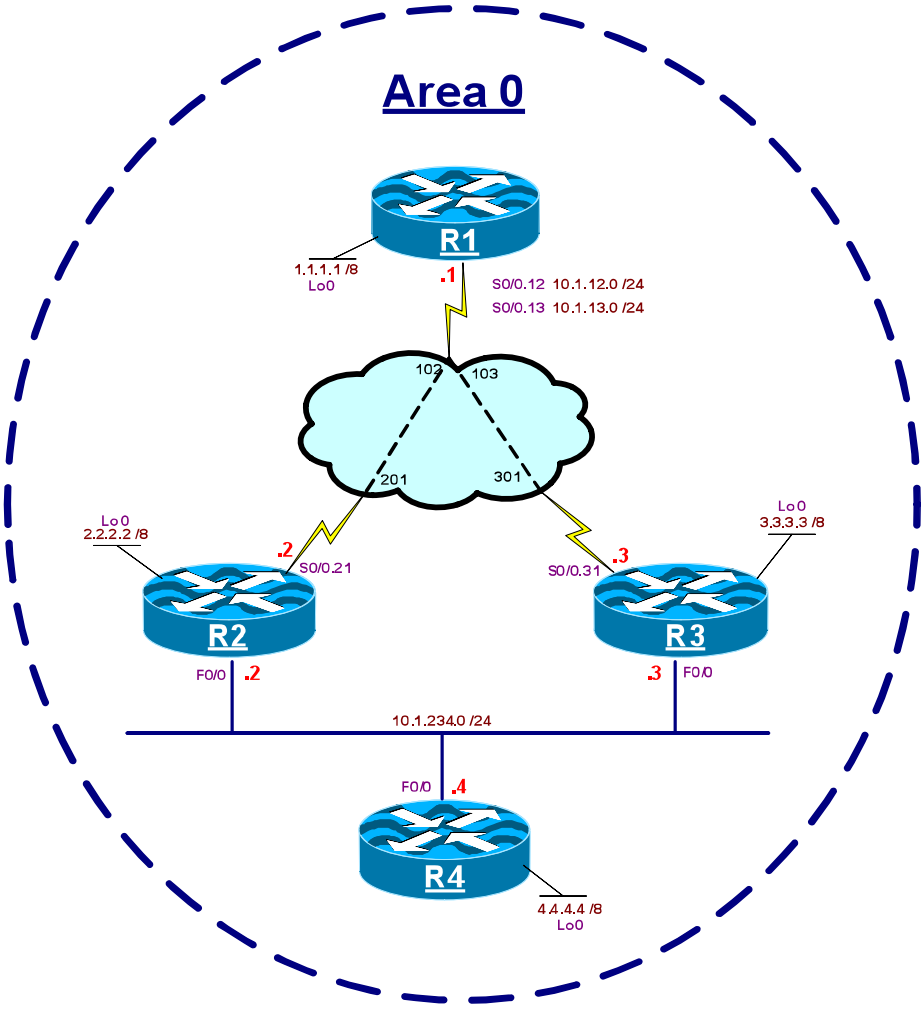
O IA 1.0.0.0/8 [110/3] via 10.2.2.3, 00:20:06, FastEthernet0/1
O    2.0.0.0/8 [110/3] via 10.2.2.3, 00:16:58, FastEthernet0/1
O E1    100.1.1.0 [110/22] via 10.2.2.3, 00:19:45, FastEthernet0/1
O IA    10.1.1.0 [110/2] via 10.2.2.3, 00:22:52, FastEthernet0/1
```

Task 10

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 10
Redirecting Traffic in OSPF



Lab Setup:

- R1 should be configured with two point-to-point links, one connecting R1 to R2 and the other connecting R1 to R3. R2 and R3 should also be configured in a point-to-point manner.
- R2, R3 and R4's F0/0 interfaces should be configured to be in VLAN 234.
- Configure IP addressing based on the above diagram.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 1

Configure OSPF on all routers, advertise their directly connected interfaces in area 0, and ensure that these routers can reach all the advertised networks. Loopback interfaces **must** be advertised with their correct mask. Router ID of these routers should be based on their Lo0 interface.

On All Routers:

```
Rx(config-if)#router ospf 1
Rx(config-router)#netw 0.0.0.0 0.0.0.0 are 0

Rx(config-router)#interface Lo0
Rx(config-if)#ip ospf network point-to-point
```

On R1:

```
R1(config)#router ospf 1
R1(config-router)#router-id 1.1.1.1
```

On R2:

```
R1(config)#router ospf 1
R1(config-router)#router-id 2.2.2.2
```

On R3:

```
R1(config)#router ospf 1
R1(config-router)#router-id 3.3.3.3
```

On R4:

```
R1(config)#router ospf 1
R1(config-router)#router-id 4.4.4.4
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf | inc O

O       2.2.2.2 [110/65] via 10.1.12.2, 00:00:51, Serial0/0.12
O       3.3.3.3 [110/65] via 10.1.13.3, 00:01:01, Serial0/0.13
O       4.4.4.4 [110/66] via 10.1.13.3, 00:00:11, Serial0/0.13
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O       10.1.234.0 [110/65] via 10.1.13.3, 00:00:21, Serial0/0.13
```

On R4:

```
R4#Show ip route ospf

     1.0.0.0/32 is subnetted, 1 subnets
O       1.1.1.1 [110/66] via 10.1.234.3, 00:01:14, FastEthernet0/0
                [110/66] via 10.1.234.2, 00:01:14, FastEthernet0/0
     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/2] via 10.1.234.2, 00:01:14, FastEthernet0/0
     3.0.0.0/32 is subnetted, 1 subnets
O       3.3.3.3 [110/2] via 10.1.234.3, 00:01:14, FastEthernet0/0
     10.0.0.0/24 is subnetted, 3 subnets
O       10.1.13.0 [110/65] via 10.1.234.3, 00:01:14, FastEthernet0/0
O       10.1.12.0 [110/65] via 10.1.234.2, 00:01:14, FastEthernet0/0
```

Task 2

R1 has two ways to reach network 4.0.0.0 /8; configure the appropriate router such that R1 uses R2 to reach this network. R1 should go directly to R3 to reach network 3.0.0.0. However, if R2 goes down, R3 should be used as a transit router to reach network 4.0.0.0 /8.

DO **NOT** USE THE FOLLOWING COMMANDS:
Bandwidth, any global config command, "IP OSPF cost" command, PBR or the distance command.

The routing table of R1 reveals that R1 can reach network 4.0.0.0 via R2 and R3.

On R1:

```
R1#Show ip route ospf

     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/65] via 10.1.12.2, 00:05:12, Serial0/0.12
     3.0.0.0/32 is subnetted, 1 subnets
O       3.3.3.3 [110/65] via 10.1.13.3, 00:05:22, Serial0/0.13
     4.0.0.0/32 is subnetted, 1 subnets
O       4.4.4.4 [110/66] via 10.1.13.3, 00:04:32, Serial0/0.13
                [110/66] via 10.1.12.2, 00:04:32, Serial0/0.12
     10.0.0.0/24 is subnetted, 3 subnets
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O       10.1.234.0 [110/65] via 10.1.13.3, 00:04:42, Serial0/0.13
                   [110/65] via 10.1.12.2, 00:04:42, Serial0/0.12
```

On R3:

```
R3(config)#router ospf 1
R3(config-router)#max-metric router-lsa
```

The above command will cause R3 to originate LSAs with a maximum metric of 0xffff (LSInfinity). This is done so that other routers do not prefer R3 as a transit hop in their path to a given network.

To Verify the configuration:

On R1:

```
R1#Show ip route ospf

     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/65] via 10.1.12.2, 00:21:10, Serial0/0.12
     3.0.0.0/32 is subnetted, 1 subnets
O       3.3.3.3 [110/65] via 10.1.13.3, 00:21:20, Serial0/0.13
     4.0.0.0/32 is subnetted, 1 subnets
O       4.4.4.4 [110/66] via 10.1.12.2, 00:20:31, Serial0/0.12
     10.0.0.0/24 is subnetted, 3 subnets
O       10.1.234.0 [110/65] via 10.1.12.2, 00:20:41, Serial0/0.12
```

To test this configuration:

On R2:

```
R2(config)#int s0/0.21
R2(config-subif)#shut
```

On R1:

```
R1#Show ip route ospf

     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/65600] via 10.1.13.3, 00:00:11, Serial0/0.13
     3.0.0.0/32 is subnetted, 1 subnets
O       3.3.3.3 [110/65] via 10.1.13.3, 00:23:03, Serial0/0.13
     4.0.0.0/32 is subnetted, 1 subnets
O       4.4.4.4 [110/65600] via 10.1.13.3, 00:00:11, Serial0/0.13
     10.0.0.0/24 is subnetted, 3 subnets
O       10.1.234.0 [110/65599] via 10.1.13.3, 00:00:11, Serial0/0.13
```

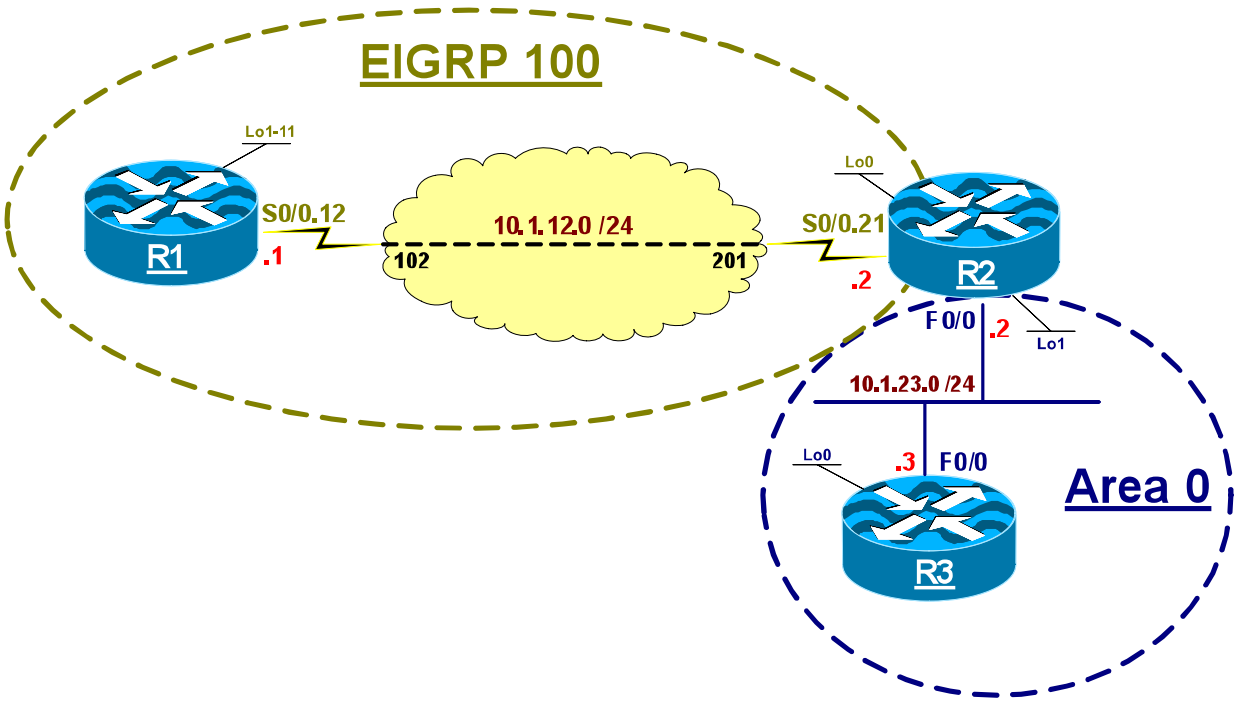

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 3

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 11
Database Overload Protection



Lab Setup:

- The frame-relay connection between R1 and R2 must be configured in a point-to-point manner.

- R2 and R3's F0/0 interface should be configured in VLAN 23.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP Addressing:

Router	Interface	IP address	Area
R1	Lo1	1.0.0.1 /8	Eigrp 100
	Lo2	2.0.0.1 /8	Eigrp 100
	Lo3	3.0.0.1 /8	Eigrp 100
	Lo4	4.0.0.1 /8	Eigrp 100
	Lo5	5.0.0.1 /8	Eigrp 100
	Lo6	6.0.0.1 /8	Eigrp 100
	Lo7	7.0.0.1 /8	Eigrp 100
	Lo8	8.0.0.1 /8	Eigrp 100
	Lo9	9.0.0.1 /8	Eigrp 100
	Lo10	10.0.0.1 /24	Eigrp 100
	Lo11	11.0.0.1 /8	Eigrp 100
	S0/0.12	10.1.12.1 /24	Eigrp 100
R2	Lo0	20.2.2.2 /8	Eigrp 100
	Lo1	22.2.2.2 /8	OSPF area 0
	S0/0.21	10.1.12.2 /24	Eigrp 100
	F0/0	10.1.23.2 /24	OSPF area 0
R3	Lo0	3.3.3.3 /8	OSPF area 0
	F0/0	10.1.23.3 /24	OSPF area 0

Task 1

Configure the routers as follows:

- On R1:, configure EIGRP 100 and advertise networks 1.0.0.0 /8 – 5.0.0.0 /8 and the frame-relay interface to R2. Disable auto summarization.
- On R2:, configure EIGRP 100 and advertise networks 20.0.0.0 /8 and the frame-relay interface link to R1. Disable auto summarization.
- On R2:, configure OSPF and advertise network 22.0.0.0 /8 and the F0/0 interface to R3 in area 0.
- On R3:, configure OSPF and advertise all of its interfaces in area 0.

On R1:

```
R1(config)#router eigrp 100
R1(config-router)#no au
R1(config-router)#netw 1.0.0.0
R1(config-router)#netw 2.0.0.0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config-router)#netw 3.0.0.0
R1(config-router)#netw 4.0.0.0
R1(config-router)#netw 5.0.0.0
R1(config-router)#netw 10.1.12.1 0.0.0.0
```

On R2:

```
R2(config-router)#router eigrp 100
R2(config-router)#netw 10.1.12.2 0.0.0.0
R2(config-router)#netw 20.0.0.0
R2(config-router)#no auto

R2(config-if)#router ospf 1
R2(config-router)#netw 22.2.2.2 0.0.0.0 are 0
R2(config-router)#netw 10.1.23.2 0.0.0.0 are 0
```

On R3:

```
R3(config-if)#router ospf 1
R3(config-router)#netw 0.0.0.0 0.0.0.0 are 0
```

To verify the configuration:

On R2:

```
R2#Sh ip route eigrp

D    1.0.0.0/8 [90/2297856] via 10.1.12.1, 00:04:04, Serial0/0.21
D    2.0.0.0/8 [90/2297856] via 10.1.12.1, 00:01:14, Serial0/0.21
     3.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
D       3.0.0.0/8 [90/2297856] via 10.1.12.1, 00:04:04, Serial0/0.21
D    4.0.0.0/8 [90/2297856] via 10.1.12.1, 00:04:04, Serial0/0.21
D    5.0.0.0/8 [90/2297856] via 10.1.12.1, 00:04:04, Serial0/0.21

R2#Show ip route ospf

     3.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
O       3.3.3.3/32 [110/2] via 10.1.23.3, 00:03:48, FastEthernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 2

On R2:, perform a mutual redistribution between OSPF and EIGRP 100. Use any metric appropriate.

On R2:

```
R2(config)#router eigrp 100
R2(config-router)#redistribute ospf 1 metric 1 1 1 1 1

R2(config-router)#router ospf 1
R2(config-router)#redistribute eigrp 100 subnets
```

When redistributing routes into OSPF, they will be redistributed with a metric of 20 and metric-type of 2.

To verify the configuration:

On R3:

```
R3#Show ip route ospf

O E2 1.0.0.0/8 [110/20] via 10.1.23.2, 00:01:18, FastEthernet0/0
O E2 2.0.0.0/8 [110/20] via 10.1.23.2, 00:01:18, FastEthernet0/0
O E2 4.0.0.0/8 [110/20] via 10.1.23.2, 00:01:18, FastEthernet0/0
O E2 20.0.0.0/8 [110/20] via 10.1.23.2, 00:01:18, FastEthernet0/0
O E2 5.0.0.0/8 [110/20] via 10.1.23.2, 00:01:18, FastEthernet0/0
     22.0.0.0/32 is subnetted, 1 subnets
O       22.2.2.2 [110/2] via 10.1.23.2, 00:05:38, FastEthernet0/0
     10.0.0.0/24 is subnetted, 2 subnets
O E2    10.1.12.0 [110/20] via 10.1.23.2, 00:01:18, FastEthernet0/0
```

Six routes were redistributed into OSPF routing domain.

Task 3

Configure R2 such that the maximum number of prefixes that can be redistributed into OSPF routing protocol is 10. R2 should generate two warning messages. The first message should occur when the number of redistributed prefixes reaches 70% of the configured threshold (10). The second message should occur when the 10th prefix is redistributed.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2:

```
R2(config-router)#router ospf 1
R2(config-router)#redistribute maximum-prefix 10 70 warning-only
```

The above command limits the number of prefixes that can be redistributed into OSPF routing domain. In this case, the router will generate two warning messages, the first one will be generated when 70% of the configured threshold (10) is reached. The second message will be generated when the configured threshold (10) is exceeded.

The initial config file has created 11 loopback interfaces for testing purpose. Test this policy by advertising these loopback interfaces one at a time and observe the console warning messages.

Task 4

The administrator of R1 is constantly violating the above policy, in order to safeguard against this, you should configure R2 such that only 10 prefixes are allowed to be redistributed into OSPF, if R1 advertises more than 10 prefixes in EIGRP 100, and R2 should ignore the extra prefixes.

On R2:

```
R2(config)#router ospf 1
R2(config-router)#no redistribute maximum-prefix 10 70 warning-only
R2(config-router)#redistribute maximum-prefix 10 70
```

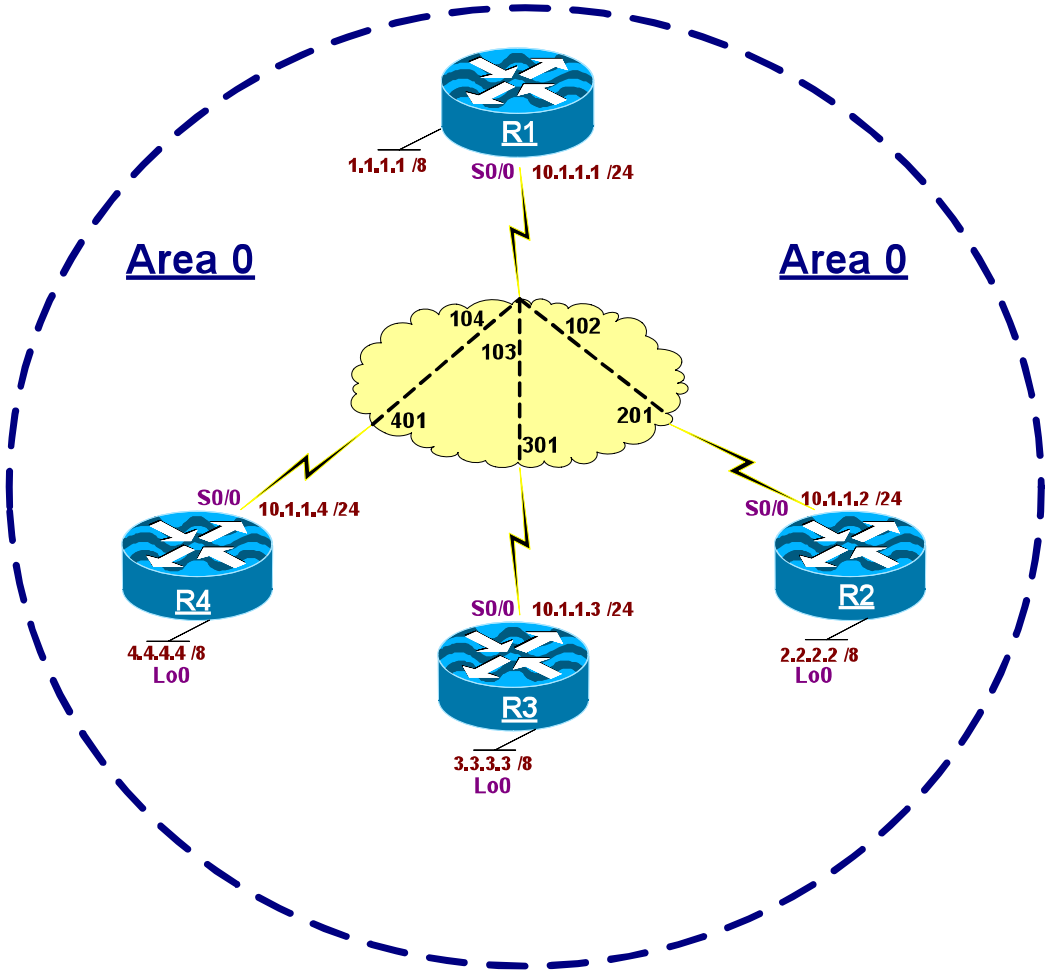
Since the "warning-only" keyword is not used, R2 will ignore any advertisement above the set threshold. You should test this by advertising more loopback interfaces on R1.

Task 5

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 12 – OSPF Non-Broadcast Networks



Lab Setup:

- Configure R1 as the hub and R2, R3 and R4 as spokes.
- Configure all routers in a frame-relay multipoint manner. **Do not** configure sub-interfaces on any of the routers. Use the broadcast keyword when configuring the "frame-relay map" statements.
- Use the IP addressing chart below for IP assignment.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP addressing:

Router	Interface / IP address
R1	S0/0 = 10.1.1.1 /24 Loopback0 = 1.1.1.1 /8
R2	**S0/0 = 10.1.1.2 /24** **Loopback0 = 2.2.2.2 /8**
R3	S0/0 = 10.1.1.3 /24 Loopback0 = 3.3.3.3 /8
R4	**S0/0 = 10.1.1.4 /24** **Loopback0 = 4.4.4.4 /8**

Task 1

Configure OSPF on all routers and advertise their directly connected interfaces in area 0. Ensure that loopback 0 interface of these routers are advertised with their correct mask.

On R1:

```
R1(config)#router ospf 1
R1(config-router)#netw 10.1.1.1 0.0.0.0 area 0
R1(config-router)#netw 1.1.1.1 0.0.0.0 area 0

R1(config-router)#int lo0
R1(config-if)#ip ospf network point-to-point
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#netw 10.1.1.2 0.0.0.0 area 0
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 0

R2(config-router)#int lo0
R2(config-if)#ip ospf network point-to-point
```

On R3:

```
R3(config)#router ospf 1
R3(config-router)#netw 10.1.1.3 0.0.0.0 area 0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0

R3(config-router)#int lo0
R3(config-if)#ip ospf network point-to-point
```

On R4:

```
R4(config)#router ospf 1
R4(config-router)#netw 10.1.1.4 0.0.0.0 area 0
R4(config-router)#netw 4.4.4.4 0.0.0.0 area 0

R4(config-router)#int lo0
R4(config-if)#ip ospf network point-to-point
```

Note: The following command is required so the spokes will not participate in DR/BDR election.

On R2, R3 and R4

```
Rx(config)#interface serial0/0
Rx(config-if)#ip ospf priority 0
```

To verify the configuration:

On R1:

```
R1#Sh ip ospf neighbor

R1#Show ip route ospf
```

Note: There are no neighbor adjacencies established, as a result of that; there will not be any routes in the OSPF routing table.

The problem is the OSPF network type, the default OSPF network type on multipoint frame-relay interface is non broadcast, the following show command reveals the OSPF network type.

On R1:

```
R1#Show ip ospf interface s0/0

Serial0/0 is up, line protocol is up
  Internet Address 10.1.1.1/24, Area 0
  Process ID 1, Router ID 1.1.1.1, Network Type NON_BROADCAST, Cost: 64
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 1.1.1.1, Interface address 10.1.1.1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
  No backup designated router on this network
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
    oob-resync timeout 120
    Hello due in 00:00:11
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 1/1, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 0, maximum is 0
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 0, Adjacent neighbor count is 0
  Suppress hello for 0 neighbor(s)
```

In OSPF "NON-BROADCAST" network type, the routers do not have broadcast or multicast capability, and since OSPF uses multicast to form neighbor adjacency, these routers have not established a neighbor adjacency.

Since the task states that the network type cannot be changed, then, the "neighbor" router configuration command can be used to accomplish this task. In this case, the "neighbor" command needs to be configured on the hub router only. With the "neighbor" command, OSPF's behavior is changed to unicast instead of multicast.

In a hub and spoke topology, since the spokes do not have layer two reachability to each other, they should never participate in the DR/BDR election. The priority of their interface should always be set to zero. Since the priority of the interface is set to zero, IOS will discard the neighbor command on that router.

On R1:

```
R1(config)#router ospf 1
R1(config-router)#neighbor 10.1.1.2
R1(config-router)#neighbor 10.1.1.3
R1(config-router)#neighbor 10.1.1.4
```

You should see the following console messages:

%OSPF-5-ADJCHG: Process 1, Nbr 2.2.2.2 on Serial0/0 from LOADING to FULL, Loading Done
%OSPF-5-ADJCHG: Process 1, Nbr 3.3.3.3 on Serial0/0 from LOADING to FULL, Loading Done
%OSPF-5-ADJCHG: Process 1, Nbr 4.4.4.4 on Serial0/0 from LOADING to FULL, Loading Done

To verify the configuration:

On R1:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
2.2.2.2           0   FULL/DROTHER    00:01:46    10.1.1.2        Serial0/0
3.3.3.3           0   FULL/DROTHER    00:01:36    10.1.1.3        Serial0/0
4.4.4.4           0   FULL/DROTHER    00:01:42    10.1.1.4        Serial0/0

R1#Show ip route ospf

O    2.0.0.0/8 [110/65] via 10.1.1.2, 00:00:52, Serial0/0
O    3.0.0.0/8 [110/65] via 10.1.1.3, 00:00:52, Serial0/0
O    4.0.0.0/8 [110/65] via 10.1.1.4, 00:00:52, Serial0/0
```

Note: Once the "neighbor" command is configured on the hub router, the routers will transition into full state and exchange routes.

Remember, when the "neighbor" command is configured, all OSPF packets will use unicast instead of multicast. Therefore, the "frame-relay map" commands do not need to be configured with the "broadcast" keyword.

On R1:

```
R1(config)#int s0/0
R1(config-if)#No frame-relay map ip 10.1.1.2 102 broadcast
R1(config-if)#No frame-relay map ip 10.1.1.3 103 broadcast
R1(config-if)#No frame-relay map ip 10.1.1.4 104 broadcast

R1(config-if)#frame-relay map ip 10.1.1.2 102
R1(config-if)#frame-relay map ip 10.1.1.3 103
R1(config-if)#frame-relay map ip 10.1.1.4 104
```

On R2:

```
R2(config)#int s0/0
R2(config-if)#No frame-relay map ip 10.1.1.1 201 broadcast

R2(config-if)#frame-relay map ip 10.1.1.1 201
```

On R3:

```
R3(config)#int s0/0
R3(config-if)#No frame-relay map ip 10.1.1.1 301 broadcast

R3(config-if)#frame-relay map ip 10.1.1.1 301
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4:

```
R4(config)#int s0/0
R4(config-if)#No frame-relay map ip 10.1.1.1 401 broadcast

R4(config-if)#frame-relay map ip 10.1.1.1 401
```

To test the configuration:

On R1:

```
R1#Show ip route ospf

O    2.0.0.0/8 [110/65] via 10.1.1.2, 00:00:10, Serial0/0
O    3.0.0.0/8 [110/65] via 10.1.1.3, 00:00:10, Serial0/0
O    4.0.0.0/8 [110/65] via 10.1.1.4, 00:00:10, Serial0/0
```

On R2:

```
R2#Show ip route ospf

O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:07:21, Serial0/0
O    3.0.0.0/8 [110/65] via 10.1.1.3, 00:07:21, Serial0/0
O    4.0.0.0/8 [110/65] via 10.1.1.4, 00:07:21, Serial0/0
```

On R3:

```
R3#Show ip route ospf

O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:07:57, Serial0/0
O    2.0.0.0/8 [110/65] via 10.1.1.2, 00:07:57, Serial0/0
O    4.0.0.0/8 [110/65] via 10.1.1.4, 00:07:57, Serial0/0
```

On R4:

```
R4#Show ip route ospf

O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:08:30, Serial0/0
O    2.0.0.0/8 [110/65] via 10.1.1.2, 00:08:30, Serial0/0
O    3.0.0.0/8 [110/65] via 10.1.1.3, 00:08:30, Serial0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 2

Ensure that every router has NLRI to the loopback interfaces advertised in OSPF, you should use ping to test and verify reachability.

To test the reachability:

On R2:

```
R2#Show ip route ospf

O     1.0.0.0/8 [110/65] via 10.1.1.1, 08:27:26, Serial0/0
O     3.0.0.0/8 [110/65] via 10.1.1.3, 08:27:26, Serial0/0
O     4.0.0.0/8 [110/65] via 10.1.1.4, 08:27:26, Serial0/0
```

Note: R2 does not have reachability to networks 3.0.0.0/8 or 4.0.0.0/8, because it does not have layer 2 mapping for the next hop IP address.

In OSPF NON_BROADCAST network type the next hop IP address is the IP address of the router that originated the route and not the router that advertised it. Therefore, the spokes will not have NLRI to networks advertised by other spokes, this problem can be resolved by configuring the following frame-relay map commands:

On each spoke a "frame-relay map" command is configured for the frame-relay interface IP address of the other spokes using their only DLCI pointing to the hub, as follows:

On R2:

```
R2(config)#interface s0/0
R2(config-if)#frame-relay map ip 10.1.1.3 201
R2(config-if)#frame-relay map ip 10.1.1.4 201
```

On R3:

```
R3(config)#interface s0/0
R3(config-if)#frame-relay map ip 10.1.1.2 301
R3(config-if)#frame-relay map ip 10.1.1.4 301
```

On R4:

```
R4(config)#interface s0/0
R4(config-if)#frame-relay map ip 10.1.1.2 401
R4(config-if)#frame-relay map ip 10.1.1.3 401
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To test the configuration:

On R2:

```
R2#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/113/117 ms

R2#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/115/120 ms
```

On R3:

```
R3#Ping 2.2.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/125/168 ms

R3#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/113/116 ms
```

On R4:

```
R4#Ping 2.2.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/121/148 ms

R4#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

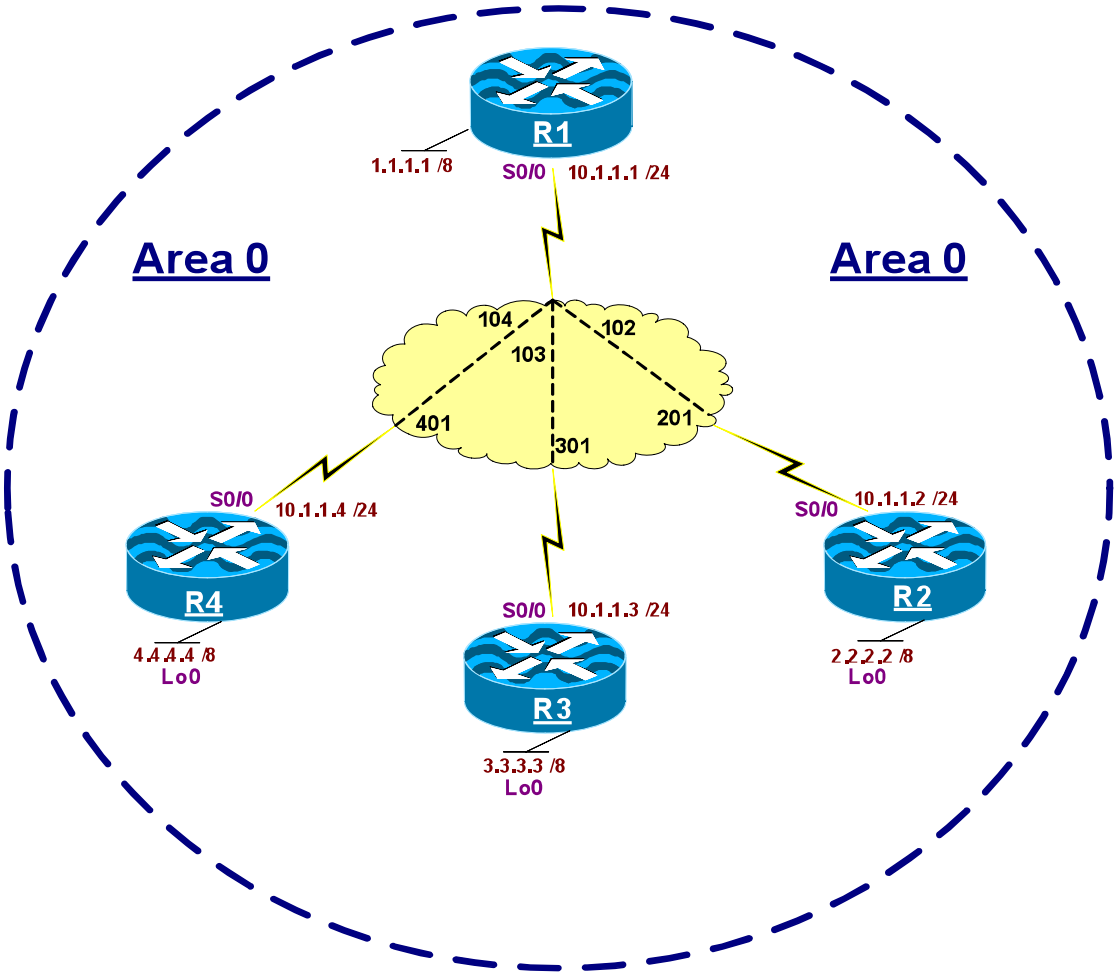
```
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/113/116 ms
```

Task 3

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 13 – OSPF Broadcast Networks



Lab Setup:

- Configure R1 as the hub and R2, R3 and R4 as spokes.
- Configure all routers in a frame-relay multipoint manner. **Do not** configure sub-interfaces on any of the routers. Use the "broadcast" keyword when configuring the "frame-relay map" statements.
- Use the IP addressing chart below for IP assignment.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP addressing:

Router	Interface / IP address
R1	S0/0 = 10.1.1.1 /24 Loopback0 = 1.1.1.1 /8
R2	**S0/0 = 10.1.1.2 /24** **Loopback0 = 2.2.2.2 /8**
R3	S0/0 = 10.1.1.3 /24 Loopback0 = 3.3.3.3 /8
R4	**S0/0 = 10.1.1.4 /24** **Loopback0 = 4.4.4.4 /8**

Task 1

Configure OSPF on all routers and advertise their directly connected interfaces in area 0. Ensure that loopback 0 interface of these routers are advertised with their correct mask. You should use OSPF "BROADCAST" network type to accomplish this task.

On R1:

```
R1(config)#router ospf 1
R1(config-router)#netw 10.1.1.1 0.0.0.0 area 0
R1(config-router)#netw 1.1.1.1 0.0.0.0 area 0

R1(config-router)#int lo0
R1(config-if)#ip ospf network point-to-point

R1(config)#interface serial0/0
R1(config-if)#ip ospf network broadcast
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#netw 10.1.1.2 0.0.0.0 area 0
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 0

R2(config-router)#int lo0
R2(config-if)#ip ospf network point-to-point
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3:

```
R3(config)#router ospf 1
R3(config-router)#netw 10.1.1.3 0.0.0.0 area 0
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0

R3(config-router)#int lo0
R3(config-if)#ip ospf network point-to-point
```

On R4:

```
R4(config)#router ospf 1
R4(config-router)#netw 10.1.1.4 0.0.0.0 area 0
R4(config-router)#netw 4.4.4.4 0.0.0.0 area 0

R4(config-router)#int lo0
R4(config-if)#ip ospf network point-to-point
```

Note: The following command is required so the spokes will not participate in DR/BDR election; it also changes the network type to broadcast.

On R2, R3 and R4

```
Rx(config)#interface serial0/0
Rx(config-if)#ip ospf network broadcast
Rx(config-if)#ip ospf priority 0
```

To verify the configuration:

On R1:

```
R1#Sh ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
2.2.2.2           0   FULL/DROTHER    00:00:37    10.1.1.2        Serial0/0
3.3.3.3           0   FULL/DROTHER    00:00:39    10.1.1.3        Serial0/0
4.4.4.4           0   FULL/DROTHER    00:00:37    10.1.1.4        Serial0/0

R1#Sh ip route ospf

O    2.0.0.0/8 [110/65] via 10.1.1.2, 00:01:51, Serial0/0
O    3.0.0.0/8 [110/65] via 10.1.1.3, 00:01:51, Serial0/0
O    4.0.0.0/8 [110/65] via 10.1.1.4, 00:01:51, Serial0/0
```

On R2:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2#Show ip route ospf

O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:02:29, Serial0/0
O    3.0.0.0/8 [110/65] via 10.1.1.3, 00:02:29, Serial0/0
O    4.0.0.0/8 [110/65] via 10.1.1.4, 00:02:29, Serial0/0
```

On R3:

Note the next hop was NOT changed
Just like the previous lab

```
R3#Show ip route ospf

O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:03:09, Serial0/0
O    2.0.0.0/8 [110/65] via 10.1.1.2, 00:03:09, Serial0/0
O    4.0.0.0/8 [110/65] via 10.1.1.4, 00:03:09, Serial0/0
```

On R4:

```
R4#Show ip route ospf

O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:03:46, Serial0/0
O    2.0.0.0/8 [110/65] via 10.1.1.2, 00:03:46, Serial0/0
O    3.0.0.0/8 [110/65] via 10.1.1.3, 00:03:46, Serial0/0
```

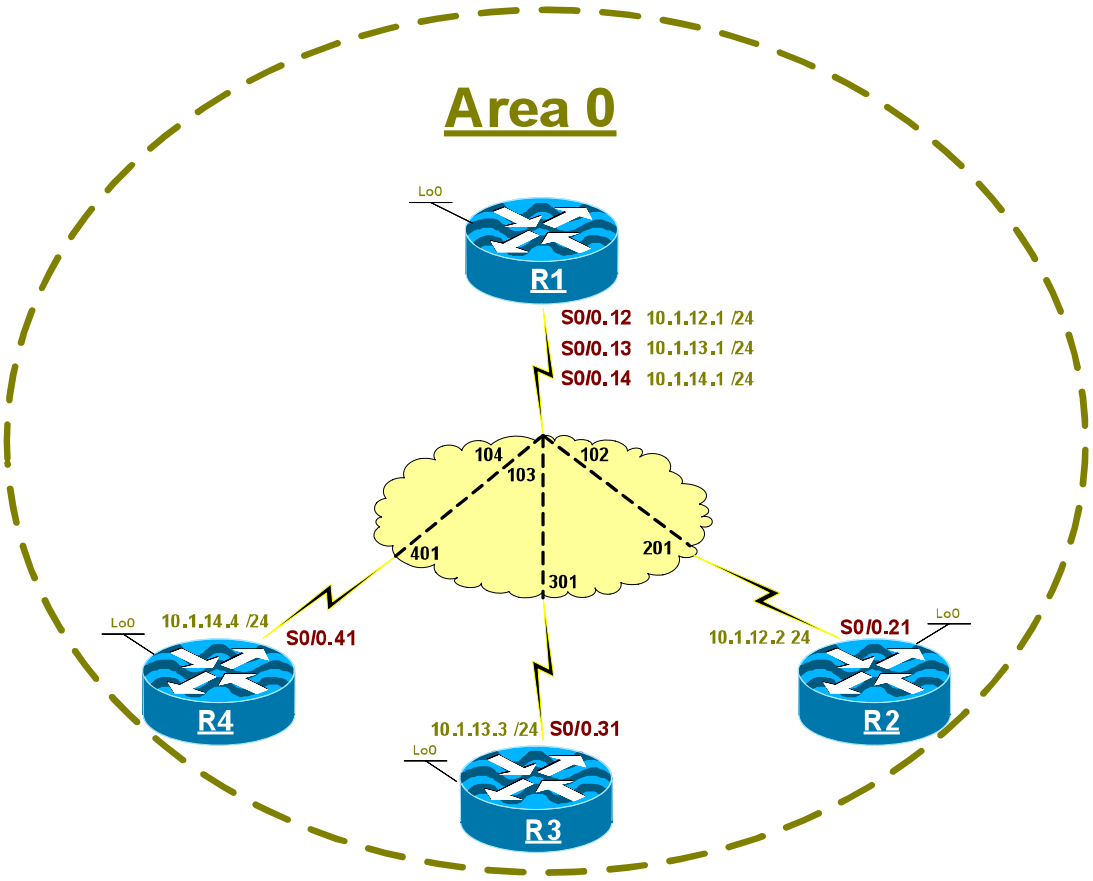
Note: Once again the next hop IP address is pointing to the router that originated the route. In this case the frame-relay solution from the previous lab can also be used as the solution to this problem. Remember that the "broadcast" keyword should not be used when configuring the "frame-relay map" statements on the spokes pointing to the frame-relay interface IP address of the other spokes.

Task 2

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 14 – OSPF Point-to-Point Networks



Lab Setup:

- Configure R1 as the hub and R2, R3 and R4 as spokes.
- Configure all routers in a frame-relay point-to-point manner.
- Use the IP addressing chart below for IP assignment.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP addressing:

Router	Interface / IP address	DLCI / Router
R1	**S0/0.12 = 10.1.12.1 /24** **S0/0.13 = 10.1.13.1 /24** **S0/0.14 = 10.1.14.1 /24** **Loopback0 = 1.1.1.1 /8**	**102 / R2** **103 / R3** **104 / R4**
R2	S0/0.21 = 10.1.12.2 /24 Loopback0 = 2.2.2.2 /8	201 / R1
R3	**S0/0.31 = 10.1.13.3 /24** **Loopback0 = 3.3.3.3 /8**	**301 / R1**
R4	S0/0.41 = 10.1.14.4 /24 Loopback0 = 4.4.4.4 /8	401 / R1

Task 1

Configure OSPF on all routers and advertise their directly connected interfaces in area 0. Ensure that loopback 0 interface of these routers are advertised with their correct mask.

On R1:

```
R1(config-if)#router ospf 1
R1(config-router)#netw 1.1.1.1 0.0.0.0 area 0
R1(config-router)#netw 10.1.12.1 0.0.0.0 area 0
R1(config-router)#netw 10.1.13.1 0.0.0.0 area 0
R1(config-router)#netw 10.1.14.1 0.0.0.0 area 0

R1(config-router)#int lo0
R1(config-if)#ip ospf network point-to-point
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 0
R2(config-router)#netw 10.1.12.2 0.0.0.0 area 0

R2(config-router)#int lo0
R2(config-if)#ip ospf network point-to-point
```

On R3:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config)#router ospf 1
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.13.3 0.0.0.0 area 0

R3(config-router)#int lo0
R3(config-if)#ip ospf network point-to-point
```

On R4:

```
R4(config)#router ospf 1
R4(config-router)#netw 4.4.4.4 0.0.0.0 area 0
R4(config-router)#netw 10.1.14.4 0.0.0.0 area 0

R4(config-router)#int lo0
R4(config-if)#ip ospf network point-to-point
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf

O    2.0.0.0/8 [110/65] via 10.1.12.2, 00:02:33, Serial0/0.12
O    3.0.0.0/8 [110/65] via 10.1.13.3, 00:02:33, Serial0/0.13
O    4.0.0.0/8 [110/65] via 10.1.14.4, 00:02:33, Serial0/0.14
```

On R2:

```
R2#Show ip route ospf

O    1.0.0.0/8 [110/65] via 10.1.12.1, 00:03:07, Serial0/0.21
O    3.0.0.0/8 [110/129] via 10.1.12.1, 00:03:07, Serial0/0.21
O    4.0.0.0/8 [110/129] via 10.1.12.1, 00:03:07, Serial0/0.21
     10.0.0.0/24 is subnetted, 3 subnets
O       10.1.14.0 [110/128] via 10.1.12.1, 00:03:07, Serial0/0.21
O       10.1.13.0 [110/128] via 10.1.12.1, 00:03:07, Serial0/0.21
```

On R3:

```
R3#Show ip route ospf

O    1.0.0.0/8 [110/65] via 10.1.13.1, 00:04:15, Serial0/0.31
O    2.0.0.0/8 [110/129] via 10.1.13.1, 00:04:15, Serial0/0.31
O    4.0.0.0/8 [110/129] via 10.1.13.1, 00:04:15, Serial0/0.31
     10.0.0.0/24 is subnetted, 3 subnets
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O       10.1.14.0 [110/128] via 10.1.13.1, 00:04:15, Serial0/0.31
O       10.1.12.0 [110/128] via 10.1.13.1, 00:04:15, Serial0/0.31
```

On R4:

```
R4#Show ip route ospf

O    1.0.0.0/8 [110/65] via 10.1.14.1, 00:05:04, Serial0/0.41
O    2.0.0.0/8 [110/129] via 10.1.14.1, 00:05:04, Serial0/0.41
O    3.0.0.0/8 [110/129] via 10.1.14.1, 00:05:04, Serial0/0.41
     10.0.0.0/24 is subnetted, 3 subnets
O       10.1.13.0 [110/128] via 10.1.14.1, 00:05:04, Serial0/0.41
O       10.1.12.0 [110/128] via 10.1.14.1, 00:05:04, Serial0/0.41
```

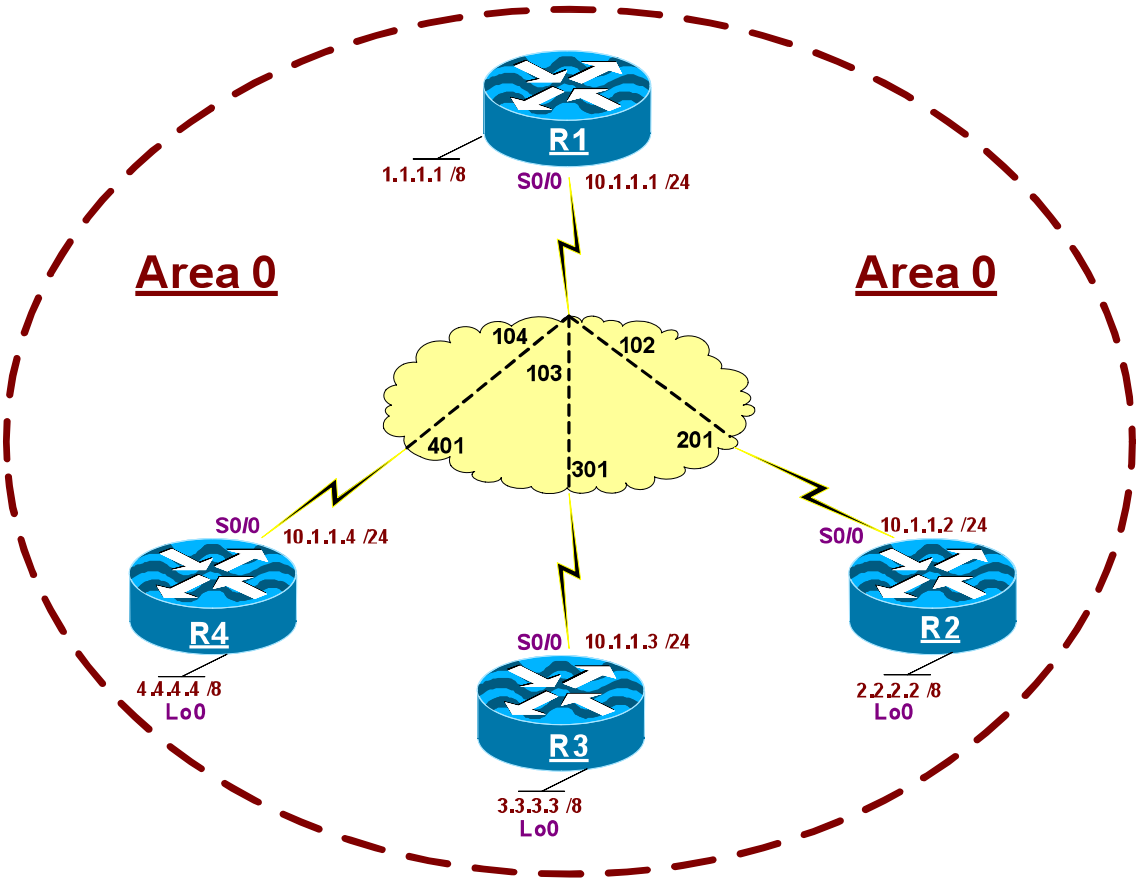
Note: The next hop is changed, this is because of the OSPF network type, in OSPF point-to-point network type. The next hop IP address is no longer the router that originated the route, and it is the router that advertised the route.

Task 2

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 15 – OSPF Point-to-Multipoint Networks-I



Lab Setup:

- Configure R1 as the hub and R2, R3 and R4 as spokes.
- Configure all routers in a frame-relay multipoint manner. **Do not** configure sub-interface/s on any of the routers. Use the "broadcast" keyword when configuring the "frame-relay map" statements.
- Use the IP addressing chart below for IP assignment.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP addressing:

Router	Interface / IP address
R1	S0/0 = 10.1.1.1 /24 Loopback0 = 1.1.1.1 /8
R2	**S0/0 = 10.1.1.2 /24** **Loopback0 = 2.2.2.2 /8**
R3	S0/0 = 10.1.1.3 /24 Loopback0 = 3.3.3.3 /8
R4	**S0/0 = 10.1.1.4 /24** **Loopback0 = 4.4.4.4 /8**

Task 1

Configure OSPF on all routers and advertise their directly connected interfaces in area 0. You should use OSPF BROADCAST network type for the frame-relay interfaces to accomplish this task. You may **not** have reachability to the advertised networks. Ensure that loopback 0 interface of these routers are advertised with their correct mask.

On R1:

```
R1(config-if)#router ospf 1
R1(config-router)#netw 1.1.1.1 0.0.0.0 area 0
R1(config-router)#netw 10.1.1.1 0.0.0.0 area 0

R1(config-router)#int lo0
R1(config-if)#ip ospf network point-to-point

R1(config-router)#int s0/0
R1(config-if)#ip ospf network Broadcast
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 0
R2(config-router)#netw 10.1.1.2 0.0.0.0 area 0

R2(config-router)#int lo0
R2(config-if)#ip ospf network point-to-point
R2(config-if)#ip ospf priority 0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config-router)#int s0/0
R2(config-if)#ip ospf network Broadcast
```

On R3:

```
R3(config)#router ospf 1
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.1.3 0.0.0.0 area 0

R3(config-router)#int lo0
R3(config-if)#ip ospf network point-to-point

R3(config-router)#int s0/0
R3(config-if)#ip ospf network Broadcast
R3(config-if)#ip ospf priority 0
```

On R4:

```
R4(config)#router ospf 1
R4(config-router)#netw 4.4.4.4 0.0.0.0 area 0
R4(config-router)#netw 10.1.1.4 0.0.0.0 area 0

R4(config-router)#int lo0
R4(config-if)#ip ospf network point-to-point

R4(config-router)#int s0/0
R4(config-if)#ip ospf network Broadcast
R4(config-if)#ip ospf priority 0
```

To verify the configuration:

On R1:

```
R1#Sh ip route ospf

O    2.0.0.0/8 [110/65] via 10.1.1.2, 00:00:46, Serial0/0
O    3.0.0.0/8 [110/65] via 10.1.1.3, 00:00:46, Serial0/0
O    4.0.0.0/8 [110/65] via 10.1.1.4, 00:00:46, Serial0/0
```

On R2:

```
R2#Show ip route ospf

O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:01:07, Serial0/0
O    3.0.0.0/8 [110/65] via 10.1.1.3, 00:01:07, Serial0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O    4.0.0.0/8 [110/65] via 10.1.1.4, 00:01:07, Serial0/0
```

On R3:

```
R3#Show ip route ospf

O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:01:15, Serial0/0
O    2.0.0.0/8 [110/65] via 10.1.1.2, 00:01:15, Serial0/0
O    4.0.0.0/8 [110/65] via 10.1.1.4, 00:01:15, Serial0/0
```

On R4:

```
R4#Show ip route ospf

O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:01:22, Serial0/0
O    2.0.0.0/8 [110/65] via 10.1.1.2, 00:01:22, Serial0/0
O    3.0.0.0/8 [110/65] via 10.1.1.3, 00:01:22, Serial0/0
```

Task 2

Ensure that the routers have reachability to every loopback interface advertised in OSPF routing protocol, **do not** use the "frame-relay map" command or any global configuration command as part of the solution to accomplish this task.

On All Routers:

```
Rx(config)#int s0/0
Rx(config-if)#ip ospf net point-to-multipoint
```

To verify the configuration:

On R1:

```
R1#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
4.4.4.4           0   FULL/  -        00:01:39    10.1.1.4        Serial0/0
3.3.3.3           0   FULL/  -        00:01:38    10.1.1.3        Serial0/0
2.2.2.2           0   FULL/  -        00:01:38    10.1.1.2        Serial0/0
```

On R2:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
1.1.1.1           0   FULL/  -        00:01:37    10.1.1.1        Serial0/0

R2#Show ip route ospf | inc O

O       1.1.1.1 [110/65] via 10.1.1.1, 00:05:09, Serial0/0
O       3.3.3.3 [110/129] via 10.1.1.1, 00:05:09, Serial0/0
O       4.4.4.4 [110/129] via 10.1.1.1, 00:05:09, Serial0/0
O       10.1.1.3/32 [110/128] via 10.1.1.1, 00:05:09, Serial0/0
O       10.1.1.1/32 [110/64] via 10.1.1.1, 00:05:09, Serial0/0
O       10.1.1.4/32 [110/128] via 10.1.1.1, 00:05:09, Serial0/0
```

On R3:

```
R3#Show ip route ospf | inc O

O       1.1.1.1 [110/65] via 10.1.1.1, 00:04:23, Serial0/0
O       2.2.2.2 [110/129] via 10.1.1.1, 00:04:23, Serial0/0
O       4.4.4.4 [110/129] via 10.1.1.1, 00:04:23, Serial0/0
O       10.1.1.2/32 [110/128] via 10.1.1.1, 00:04:23, Serial0/0
O       10.1.1.1/32 [110/64] via 10.1.1.1, 00:04:23, Serial0/0
O       10.1.1.4/32 [110/128] via 10.1.1.1, 00:04:23, Serial0/0
```

Note: The hosts routes are advertised to the local router because of the OSPF network type.

On R4:

```
R4#Show ip route ospf

O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:01:06, Serial0/0
O    2.0.0.0/8 [110/129] via 10.1.1.1, 00:01:06, Serial0/0
O    3.0.0.0/8 [110/129] via 10.1.1.1, 00:01:06, Serial0/0
     10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O       10.1.1.2/32 [110/128] via 10.1.1.1, 00:01:06, Serial0/0
O       10.1.1.3/32 [110/128] via 10.1.1.1, 00:01:06, Serial0/0
O       10.1.1.1/32 [110/64] via 10.1.1.1, 00:01:06, Serial0/0
```

With OSPF point-to-multipoint network types, each router creates and advertises a host route for the IP address of its interface connected to the frame-relay cloud. Because of this behavior, the spoke routers can now have NLRI to all the other spoke routers. The next hop IP address of the advertised prefixes is set based on the advertising router and not the router that originated the route.

This fixes the spoke-to-spoke reachability from layer 3's perspective.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To Test the configuration:

On R2:

```
R2#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/113/117 ms

R2#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/114/124 ms
```

On R3:

```
R3#Ping 2.2.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/114/120 ms

R3#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/113/117 ms
```

On R4:

```
R4#Ping 2.2.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/113/117 ms

R4#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

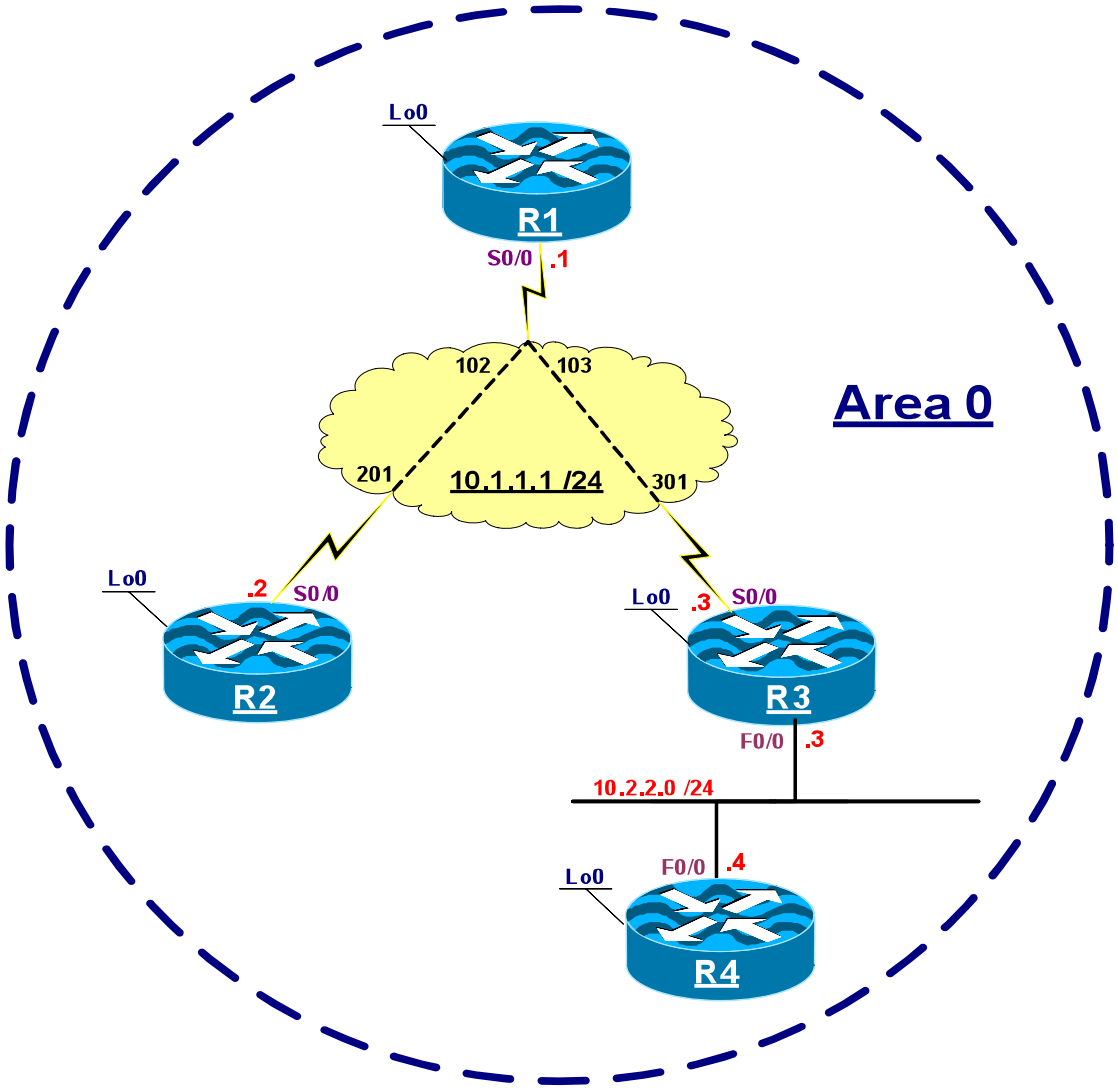
```
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/113/116 ms
```

Task 3

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 16 – OSPF Point-to-Multipoint Networks-II



©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Lab Setup:

- Configure R1 as the hub and R2, R3 as spokes.
- Configure all routers in a frame-relay multipoint manner. **Do not** configure sub-interface/s.
- F0/0 interface of R3 and R4 should be configured in VLAN 200.
- These routers should use the "broadcast" keyword when configuring the "frame-relay map" statements.
- Use the IP addressing chart below for IP assignment.

IP addressing:

Router	Interface / IP address
R1	S0/0 = 10.1.1.1 /24 Loopback0 = 1.1.1.1 /8
R2	**S0/0 = 10.1.1.2 /24** **Loopback0 = 2.2.2.2 /8**
R3	S0/0 = 10.1.1.3 /24 F0/0 = 10.2.2.3 /24 Loopback0 = 3.3.3.3 /8
R4	**F0/0 = 10.2.2.4 /24** **Loopback0 = 4.4.4.4 /8**

Task 1

Configure OSPF on all routers and advertise their directly connected interfaces in area 0. Ensure that loopback 0 interface of these routers are advertised with their correct mask. Ensure the OSPF BROADCAST network type is configured on the OSPF enabled interfaces except the loopbacks. You may **not** have reachability to the advertised networks.

On All Routers

```
Rx(config-if)#int lo0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Rx(config-if)#ip ospf net point-to-point
```

On R1:, R2 and R3

```
Rx(config)#int s0/0
Rx(config-if)#ip ospf priority 0
Rx(config-if)#ip ospf net broadcast
```

On R1:

```
R1(config)#router ospf 1
R1(config-router)#netw 1.1.1.1 0.0.0.0 are 0
R1(config-router)#netw 10.1.1.1 0.0.0.0 area 0
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 0
R2(config-router)#netw 10.1.1.2 0.0.0.0 are 0
```

On R3:

```
R3(config)#router ospf 1
R3(config-router)#netw 3.3.3.3 0.0.0.0 are 0
R3(config-router)#netw 10.2.2.3 0.0.0.0 are 0
R3(config-router)#netw 10.1.1.3 0.0.0.0 are 0
```

On R4:

```
R4(config)#router ospf 1
R4(config-router)#netw 4.4.4.4 0.0.0.0 are 0
R4(config-router)#netw 10.2.2.4 0.0.0.0 are 0
```

To verify the configuration:

On R1:

```
R1#Sh ip route ospf | inc O

O    2.0.0.0/8 [110/65] via 10.1.1.2, 00:00:48, Serial0/0
O    3.0.0.0/8 [110/65] via 10.1.1.3, 00:00:48, Serial0/0
O    4.0.0.0/8 [110/66] via 10.1.1.3, 00:00:48, Serial0/0
O       10.2.2.0 [110/65] via 10.1.1.3, 00:00:48, Serial0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2:

```
R2#Sh ip route ospf | inc O

O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:00:48, Serial0/0
O    3.0.0.0/8 [110/65] via 10.1.1.3, 00:00:38, Serial0/0
O    4.0.0.0/8 [110/66] via 10.1.1.3, 00:00:38, Serial0/0
O       10.2.2.0 [110/65] via 10.1.1.3, 00:00:38, Serial0/0
```

On R3:

```
R3#Sh ip route ospf | inc O

O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:00:48, Serial0/0
O    2.0.0.0/8 [110/65] via 10.1.1.2, 00:00:38, Serial0/0
O    4.0.0.0/8 [110/2] via 10.2.2.4, 00:01:32, FastEthernet0/0
```

On R4:

```
R4#Sh ip route ospf | inc O

O    1.0.0.0/8 [110/66] via 10.2.2.3, 00:00:38, FastEthernet0/0
O    2.0.0.0/8 [110/66] via 10.2.2.3, 00:00:38, FastEthernet0/0
O    3.0.0.0/8 [110/2] via 10.2.2.3, 00:01:32, FastEthernet0/0
O       10.1.1.0 [110/65] via 10.2.2.3, 00:00:48, FastEthernet0/0
```

Task 2

Ensure that these routers can ping every loopback interface advertised in this routing domain. **Do not** use frame-relay map, static routes, PPP or any global configuration command as part of the solution to accomplish this task.

On R1:, R2 and R3

```
Rx(config)#int s0/0
Rx(config-if)#ip ospf network point-to-multipoint
```

On R3: and R4

```
Rx(config)#int f0/0
Rx(config-if)#ip ospf network point-to-multipoint
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the configuration:

On R1:

```
R1#Sh ip route ospf | inc O

O    2.0.0.0/8 [110/65] via 10.1.1.2, 00:02:24, Serial0/0
O    3.0.0.0/8 [110/65] via 10.1.1.3, 00:02:24, Serial0/0
O    4.0.0.0/8 [110/66] via 10.1.1.3, 00:01:13, Serial0/0
O       10.1.1.2/32 [110/64] via 10.1.1.2, 00:02:24, Serial0/0
O       10.2.2.3/32 [110/64] via 10.1.1.3, 00:01:52, Serial0/0
O       10.1.1.3/32 [110/64] via 10.1.1.3, 00:02:24, Serial0/0
O       10.2.2.4/32 [110/65] via 10.1.1.3, 00:01:13, Serial0/0
```

On R2:

```
R2#Sh ip route ospf | inc O

O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:02:24, Serial0/0
O    3.0.0.0/8 [110/129] via 10.1.1.1, 00:02:24, Serial0/0
O    4.0.0.0/8 [110/130] via 10.1.1.1, 00:01:13, Serial0/0
O       10.2.2.3/32 [110/128] via 10.1.1.1, 00:01:52, Serial0/0
O       10.1.1.3/32 [110/128] via 10.1.1.1, 00:02:24, Serial0/0
O       10.1.1.1/32 [110/64] via 10.1.1.1, 00:02:24, Serial0/0
O       10.2.2.4/32 [110/129] via 10.1.1.1, 00:01:13, Serial0/0
```

On R3:

```
R3#Sh ip route ospf | inc O

O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:01:52, Serial0/0
O    2.0.0.0/8 [110/129] via 10.1.1.1, 00:01:52, Serial0/0
O    4.0.0.0/8 [110/2] via 10.2.2.4, 00:01:23, FastEthernet0/0
O       10.1.1.2/32 [110/128] via 10.1.1.1, 00:01:52, Serial0/0
O       10.1.1.1/32 [110/64] via 10.1.1.1, 00:01:52, Serial0/0
O       10.2.2.4/32 [110/1] via 10.2.2.4, 00:01:23, FastEthernet0/0
```

On R4:

```
R4#Sh ip route ospf | inc O

O    1.0.0.0/8 [110/66] via 10.2.2.3, 00:01:23, FastEthernet0/0
O    2.0.0.0/8 [110/130] via 10.2.2.3, 00:01:23, FastEthernet0/0
O    3.0.0.0/8 [110/2] via 10.2.2.3, 00:01:23, FastEthernet0/0
O       10.1.1.2/32 [110/129] via 10.2.2.3, 00:01:23, FastEthernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O        10.2.2.3/32 [110/1] via 10.2.2.3, 00:01:23, FastEthernet0/0
O        10.1.1.3/32 [110/1] via 10.2.2.3, 00:01:23, FastEthernet0/0
O        10.1.1.1/32 [110/65] via 10.2.2.3, 00:01:23, FastEthernet0/0
```

To test the configuration:

On R2:

```
R2#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/116/124 ms

R2#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/116/124 ms
```

On R3:

```
R3#Ping 2.2.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/114/124 ms
```

On R4:

```
R4#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/58/61 ms

R4#Ping 2.2.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/114/116 ms
```

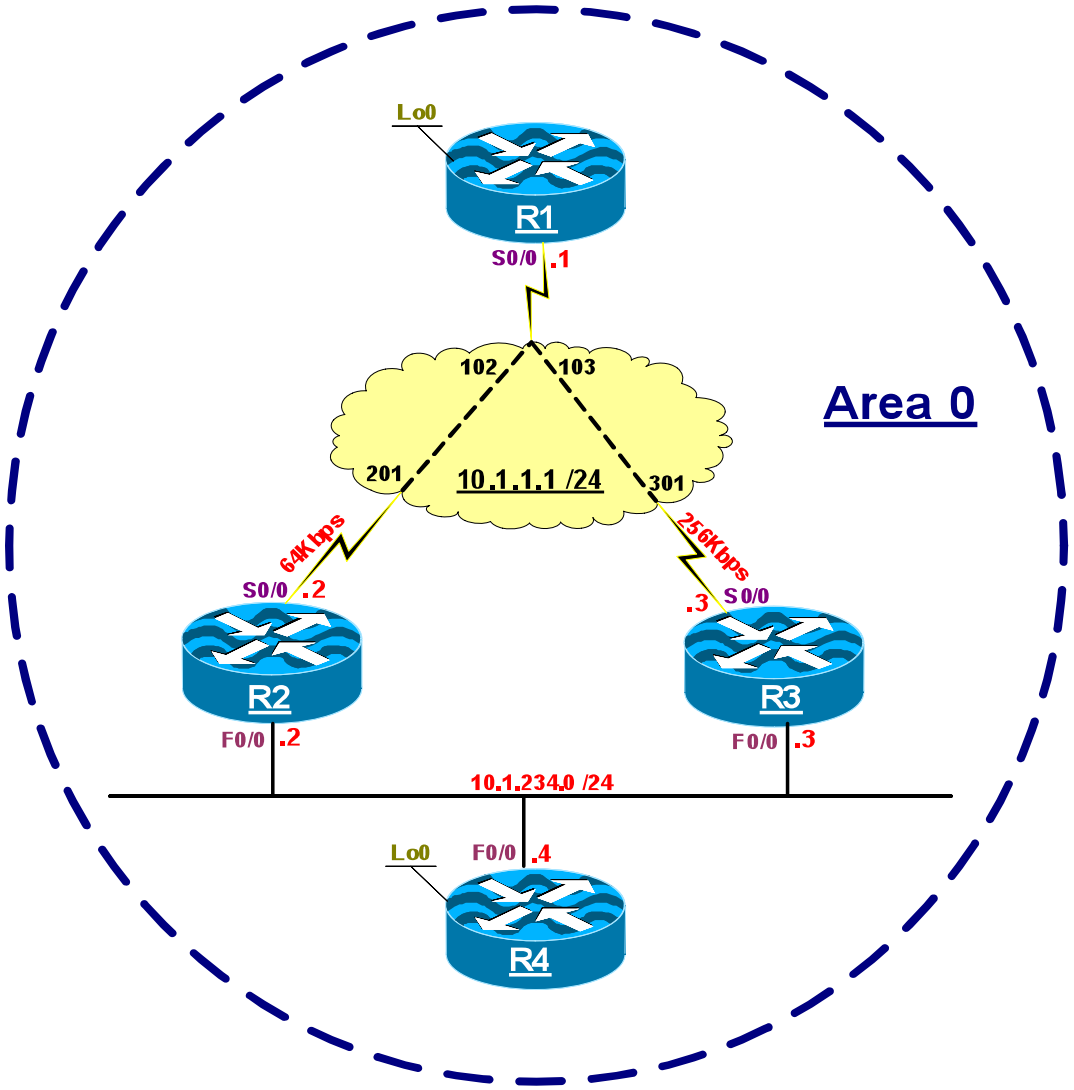

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 3

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 17 – OSPF Point-to-Multipoint NON-BROADCAST Networks



©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab Setup:

- Configure R1 as the hub and R2, and R3 as spokes.
- Configure all routers in a frame-relay multipoint manner. **Do not** configure sub-interfaces on any of the routers. Use the "broadcast" keyword when configuring the "frame-relay map" statements.
- Configure the F0/0 interface of R2, R3 and R4 in VLAN 200
- Use the IP addressing chart below for IP assignment.

IP addressing:

Router	Interface / IP address
R1	S0/0 = 10.1.1.1 /24 Loopback0 = 1.1.1.1 /8
R2	**S0/0 = 10.1.1.2 /24** **F0/0 = 10.1.234.2 /24**
R3	S0/0 = 10.1.1.3 /24 F0/0 = 10.1.234.3 /24
R4	**Loopback0 = 4.4.4.4 /8** **F0/0 = 10.1.234.4 /24**

Task 1

Configure OSPF on all routers and advertise their directly connected interfaces in area 0. You should use OSPF BROADCAST network type on the frame-relay interfaces to accomplish this task. Ensure that loopback 0 interface of R1 and R4 are advertised with their correct mask.

On R1:

```
R1(config-if)#router ospf 1
R1(config-router)#netw 1.1.1.1 0.0.0.0 area 0
R1(config-router)#netw 10.1.1.1 0.0.0.0 area 0

R1(config-router)#int lo0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config-if)#ip ospf network point-to-point

R1(config-router)#int s0/0
R1(config-if)#ip ospf network broadcast
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#netw 10.1.1.2 0.0.0.0 area 0
R2(config-router)#netw 10.1.234.2 0.0.0.0 area 0

R2(config-router)#int s0/0
R2(config-if)#ip ospf network broadcast
R2(config-if)#ip ospf priority 0
```

On R3:

```
R3(config)#router ospf 1
R3(config-router)#netw 10.1.1.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.234.3 0.0.0.0 area 0

R3(config-router)#int s0/0
R3(config-if)#ip ospf network broadcast
R3(config-if)#ip ospf priority 0
```

On R4:

```
R4(config)#router ospf 1
R4(config-router)#netw 4.4.4.4 0.0.0.0 area 0
R4(config-router)#netw 10.1.234.4 0.0.0.0 area 0

R4(config-router)#int lo0
R4(config-if)#ip ospf network point-to-point
```

To verify the configuration:

On R1:

```
R1#Sh ip route ospf

O    4.0.0.0/8 [110/66] via 10.1.1.3, 00:00:01, Serial0/0
               [110/66] via 10.1.1.2, 00:00:01, Serial0/0
     10.0.0.0/24 is subnetted, 2 subnets
O       10.1.234.0 [110/65] via 10.1.1.3, 00:00:01, Serial0/0
                   [110/65] via 10.1.1.2, 00:00:01, Serial0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2:

```
R2#Show ip route ospf

O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:01:52, Serial0/0
O    4.0.0.0/8 [110/2] via 10.1.234.4, 00:01:52, FastEthernet0/0
```

On R3:

```
R3#Show ip route ospf

O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:00:08, Serial0/0
O    4.0.0.0/8 [110/2] via 10.1.234.4, 00:00:08, FastEthernet0/0
```

On R4:

```
R4#Show ip route ospf

O    1.0.0.0/8 [110/66] via 10.1.234.3, 00:01:11, FastEthernet0/0
               [110/66] via 10.1.234.2, 00:01:11, FastEthernet0/0
     10.0.0.0/24 is subnetted, 2 subnets
O       10.1.1.0 [110/65] via 10.1.234.3, 00:01:21, FastEthernet0/0
                 [110/65] via 10.1.234.2, 00:01:21, FastEthernet0/0
```

Task 2

R2 has a frame-relay cir of 64Kbps and R3 has a frame-relay cir of 256Kbps as noted in the diagram, ensure that R1 traverses through R3 to reach the networks downstream to R2 and R3; R1 should **only** traverse through R2 if R3 is down. **Do not** use PBR, or OSPF cost, or any global configuration command to accomplish this task.

Note: Both R2 and R3 are advertising a cost of 1 (ref = 100,000,000 bps / bandwidth = 100,000,000 bps) for network 4.0.0.0 /8, R1 adds its cost of 64 through the frame-relay interface (ref = 100,000,000 / bandwidth = 1,544,000 bps) /to the cost that is advertised to it by R2 and R3. As a result of that, R1 performs equal cost load balancing. Note R2 or R3's frame-relay cost is not calculated by any of the routers.

One possible method of dealing with this scenario is to configure R1's frame-relay interface with OSPF "point-to-multipoint non-broadcast" network type. This network type allows R1 to associate a cost to each of its downstream neighbors; the neighbor with a lower cost will be chosen as the best route. Remember in order for 2 OSPF routers to exchange routes the network types must match, but there are 2 exceptions to this rule and they are as follows:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

1. A Point-to-Multipoint ←→ Point-to-Point
2. A Broadcast ←→ Non-Broadcast

In this case the first option is exercised as follows:

On R1:

```
R1(config)#int s0/0
R1(config-if)#ip ospf network point-to-multipoint non-broadcast
```

You should see the following console messages:

%OSPF-5-ADJCHG: Process 1, Nbr 10.1.234.2 on Serial0/0 from FULL to DOWN, Neighbor Down: Interface down or detached
%OSPF-5-ADJCHG: Process 1, Nbr 10.1.234.3 on Serial0/0 from FULL to DOWN, Neighbor Down: Interface down or detached

The reason the adjacency to R2 and R3 was torn down, is because of mismatch timers (hello & dead interval). The following command changes the OSPF hello-interval to match R2 and R3's hello-interval. Because the default timers for point-to-multipoint network type is 30/120, 30 for hello intervals and 120 for the deal interval.

```
R1(config-if)#ip ospf hello-interval 10
```

You should see the following console messages:

%OSPF-5-ADJCHG: Process 1, Nbr 10.1.234.2 on Serial0/0 from LOADING to FULL, Loading Done
%OSPF-5-ADJCHG: Process 1, Nbr 10.1.234.3 on Serial0/0 from LOADING to FULL, Loading Done

Note: Even though the routers are in fell state, they will not exchange any routes, unless they all agree whether they should have DR/BDR or not.

```
R1#Show ip route ospf
R1#
```

Let's change the OSPF network type on R1 such that the DR/BDR concept is no longer used by these routers. The following command changes the network type of R2 and R3 to point-to-point.

On R2: and R3

```
Rx(config)#int s0/0
Rx(config-if)#ip ospf network point-to-point
```

To verify the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1:

```
R1#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
10.1.234.3        0   FULL/  -        00:00:33    10.1.1.3        Serial0/0
10.1.234.2        0   FULL/  -        00:00:33    10.1.1.2        Serial0/0

R1#Show ip route ospf

O    4.0.0.0/8 [110/66] via 10.1.1.3, 00:01:53, Serial0/0
               [110/66] via 10.1.1.2, 00:01:53, Serial0/0
     10.0.0.0/24 is subnetted, 2 subnets
O       10.1.234.0 [110/65] via 10.1.1.3, 00:01:53, Serial0/0
                   [110/65] via 10.1.1.2, 00:01:53, Serial0/0
```

To configure the last step of this task, the neighbor commands are used with the cost keyword, in this case a lower cost is assigned to the preferred neighbor (R3):

```
R1(config-if)#router ospf 1
R1(config-router)#neighbor 10.1.1.2 cost 2
R1(config-router)#neighbor 10.1.1.3 cost 1
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf inc O

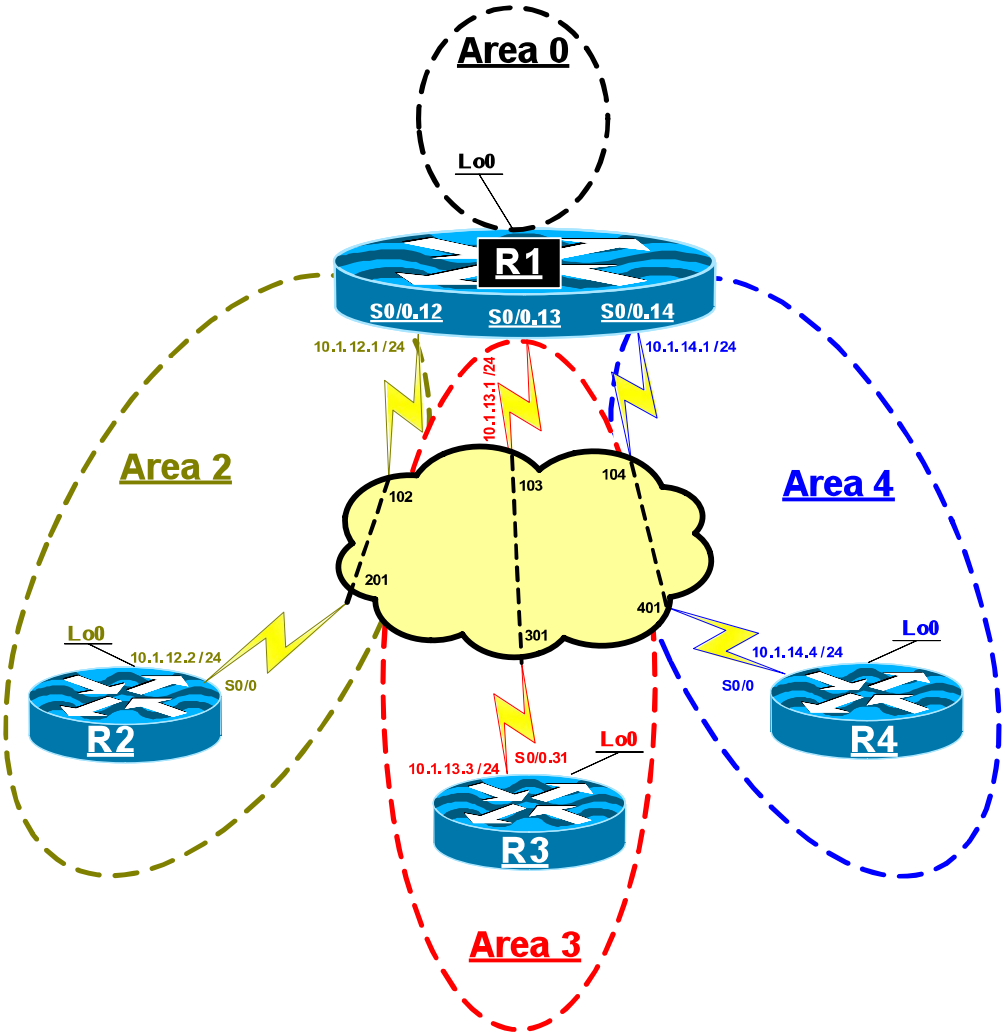
O    4.0.0.0/8 [110/3] via 10.1.1.3, 00:00:23, Serial0/0
O       10.1.234.0 [110/2] via 10.1.1.3, 00:00:23, Serial0/0
```

Task 3

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 18
OSPF and NBMA



Lab Setup:

- R1 should be configured with three sub-interfaces; the first sub-interface should be configured in a **point-to-point** manner connecting R1 to R2. R2 should not use a sub-interface for this connection.
- The second and the third sub-interface of R1 should be configured in a multipoint manner, one connecting R1 to R3 and the other one connecting R1 to R4.
- R3 should be configured in a point-to-point manner.
- R4 should **not** use a sub-interface for its connection to R1.

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

IP Addressing:

Router	Interface	IP address	Area
R1	Lo0 S0/0.12 S0/0.13 S0/0.14	1.1.1.1 /8 10.1.12.1 /24 10.1.13.1 /24 10.1.14.1 /24	Area 0 Area 2 Area 3 Area 4
R2	Lo0 S0/0	2.2.2.2 /8 10.1.12.2 /24	Area 2 Area 2
R3	Lo0 S0/0.31	3.3.3.3 /8 10.1.13.3 /24	Area 3 Area 3
R4	Lo0 S0/0	4.4.4.4 /8 10.1.14.4 /24	Area 4 Area 4

R1 should be the DR in all cases, if one is required.

Task 1

Configure OSPF on all routers and advertise their directly connected networks in their assigned area identified in the IP addressing chart. The routers may **not** have reachability or establish a neighbor adjacency.

On R1:

```
R1(config)#router ospf 1
R1(config-router)#netw 10.1.12.1 0.0.0.0 area 2
R1(config-router)#netw 10.1.13.1 0.0.0.0 area 3
R1(config-router)#netw 10.1.14.1 0.0.0.0 area 4
R1(config-router)#netw 1.1.1.1 0.0.0.0 area 0
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#netw 0.0.0.0 0.0.0.0 area 2
```

On R3:

```
R3(config)#router ospf 1
R3(config-router)#netw 0.0.0.0 0.0.0.0 area 3
```

On R4:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config-if)#router ospf 1
R4(config-router)#netw 0.0.0.0 0.0.0.0 area 4
```

Task 2

Ensure that when the routers in area 2 attempt to establish a neighbor adjacency they are successful (FULL STATE), but no routes are exchanged. **Do not** configure R2 to accomplish this task.

To see OSPF enabled interfaces and the network type on R2 and R1:

On R2:

```
R2#Sh ip ospf inter bri

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     2               2.2.2.2/8          1     LOOP  0/0
Se0/0        1     2               10.1.12.2/24       64    DR    0/0

R2#Show ip ospf int s0/0 | inc Network

  Process ID 1, Router ID 2.2.2.2, Network Type NON_BROADCAST, Cost: 64
```

On R1:

```
R1#Sh ip ospf inter bri

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     0               1.1.1.1/8          1     LOOP  0/0
Se0/0.12     1     2               10.1.12.1/24       64    P2P   0/0
Se0/0.13     1     3               10.1.13.1/24       64    DR    0/0
Se0/0.14     1     4               10.1.14.1/24       64    DR    0/0

R1#Show ip ospf inter s0/0.12 | inc Network

  Process ID 1, Router ID 1.1.1.1, Network Type POINT_TO_POINT, Cost: 64
```

R2's network type is NON_Broadcast, whereas, R1's network type is POINT_TO_POINT, they have different timers (hello and dead interval), if these interfaces are configured with identical timers, they will transition to full state, but they will not exchange routes.

Since R2 can not be configured, R1's hello and dead interval can be changed to match R2's, let's check the hello and dead interval of these two routers:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2:

```
R2#Show ip ospf int s0/0 | inc Hello

  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
    Hello due in 00:00:13
```

On R1:

```
R1#Show ip ospf int s0/0.12 | inc Hello

  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:08
```

On R1:

```
R1(config)#int s0/0.12
R1(config-subif)#ip ospf hello-interval 30
```

To verify the configuration:

On R1:

Note: Once the hello-interval is changed, the IOS will auto change the dead interval to be four times as much as the configured hello interval:

```
R1#Show ip ospf int s0/0.12 | inc Hello

 Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
    Hello due in 00:00:07
```

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 2.2.2.2 on Serial0/0.12 from LOADING to
FULL, Loading Done

R1#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
2.2.2.2           0   FULL/  -        00:01:35    10.1.12.2       Serial0/0.12

R1#Show ip route ospf
R1#
```

Note: The two routers are in full state but they have not exchanged routes. In order for two OSPF routers to exchange routes they have to be in full state first, and in order for the two routers to

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

transition into full state, the following must match:

- **Hello and Dead interval must match**
- **Area ID must match**
- **Authentication type and passwords must match**
- **Must be in the same subnet; well….as CCIEs we can get two routers to establish adjacency when they are on two different subnets.**
- **Area stub flag must match, meaning that the two routers have to be in the same area type, such as Normal, BB, Stub and/or NSSA.**
- **The MTUs must match**

Once the two routers transition into full state, they have to adhere to the following rule in order for the routes to exchange updates:

The network type must match. There is an exception to this rule, the network types that must have a DR/BDR can exchange updates, and the network types that do not have a requirement for DR/BDR can exchange updates.

In this case, the routers (R1 and R2) have transitioned into full state, but they are not exchanging routes, unless the network type is changed.

Task 3

Ensure that the routers in area 2 can exchange routes. R2 should not be configured at all to accomplish this task. **Do not** use a NON_BROADCAST network type to accomplish this task.

Since in this task R1 and R2 should exchange routes, and a network type of NON_BROADCAST can not be used, our only choice here is a BROADCAST network type:

On R1:

```
R1(config)#int s0/0.12
R1(config-subif)#ip ospf network broadcast
```

You should see the following console message stating that the routers have gone from loading state to full state:

```
%OSPF-5-ADJCHG: Process 1, Nbr 2.2.2.2 on Serial0/0.12 from LOADING to
FULL, Loading Done
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To Verify the configuration:

On R1:

Note: Once the network type is changed, the timers are automatically adjusted based on the configured network type.

```
R1#Show ip ospf interface s0/0.12 | inc Hello

  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
    Hello due in 00:00:02

R1#Show ip route ospf

     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/65] via 10.1.12.2, 00:02:09, Serial0/0.12
```

Task 4

Area 3 should be configured in a point-to-point network type, only one of the routers should be changed to accomplish this task.

On R3:

```
R3#Show ip ospf interface brief

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     3               3.3.3.3/8          1     LOOP  0/0
Se0/0.31     1     3               10.1.13.3/24       64    P2P   0/0
```

On R1:

```
R1#Show ip ospf inter brie

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     0               1.1.1.1/8          1     LOOP  0/0
Se0/0.12     1     2               10.1.12.1/24       64    BDR   1/1
Se0/0.13     1     3               10.1.13.1/24       64    DR    0/0
Se0/0.14     1     4               10.1.14.1/24       64    DR    0/0
```

Since the network type of R3 is point-to-point, the network of R1's S0/0.13 interface should be changed to point-to-point as well.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

In the earlier IOS releases when an interface was changed from "non-broadcast" to "point-to-point" we had to change the hello interval as well, because if the hello intervals did not match, the routers did not form neighbor adjacency. Nowadays the hello/dead intervals are automatically change when the network type is changed.

On R1:

```
R1(config)#int s0/0.13
R1(config-subif)#ip ospf net point-to-point
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf | inc O

O        2.2.2.2 [110/65] via 10.1.12.2, 00:14:35, Serial0/0.12
O        3.3.3.3 [110/65] via 10.1.13.3, 00:00:18, Serial0/0.13
```

Task 5

Area 4 should be configured with a totally different network type than the ones changed in the previous tasks; **do not** change the network type to accomplish this task.

On R4:

```
R4#Show ip ospf inter brie

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     4               4.4.4.4/8          1     LOOP  0/0
Se0/0        1     4               10.1.14.4/24       64    DR    0/0

R4#Show ip ospf int s0/0 | inc Network

   Process ID 1, Router ID 4.4.4.4, Network Type NON_BROADCAST, Cost: 64
```

On R1:

```
R1#Show ip ospf inter bri

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     0               1.1.1.1/8          1     LOOP  0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Se0/0.12      1      2                10.1.12.1/24        64     BDR    1/1
Se0/0.13      1      3                10.1.13.1/24        64     P2P    1/1
Se0/0.14      1      4                10.1.14.1/24        64     DR     0/0

R1#Show ip ospf inter s0/0.14 | inc Network

   Process ID 1, Router ID 1.1.1.1, Network Type NON_BROADCAST, Cost: 64
```

On R4:

Since R1 should be the DR, the priority of R4 is changed to zero:

```
R4(config)#int s0/0
R4(config-if)#ip ospf priority 0
```

R1, the hub router must be configured with a neighbor command pointing to R4:

On R1:

```
R1(config)#router ospf 1
R1(config-router)#neighbor 10.1.14.4
```

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 4.4.4.4 on Serial0/0.14 from LOADING to
FULL, Loading Done
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf | inc O

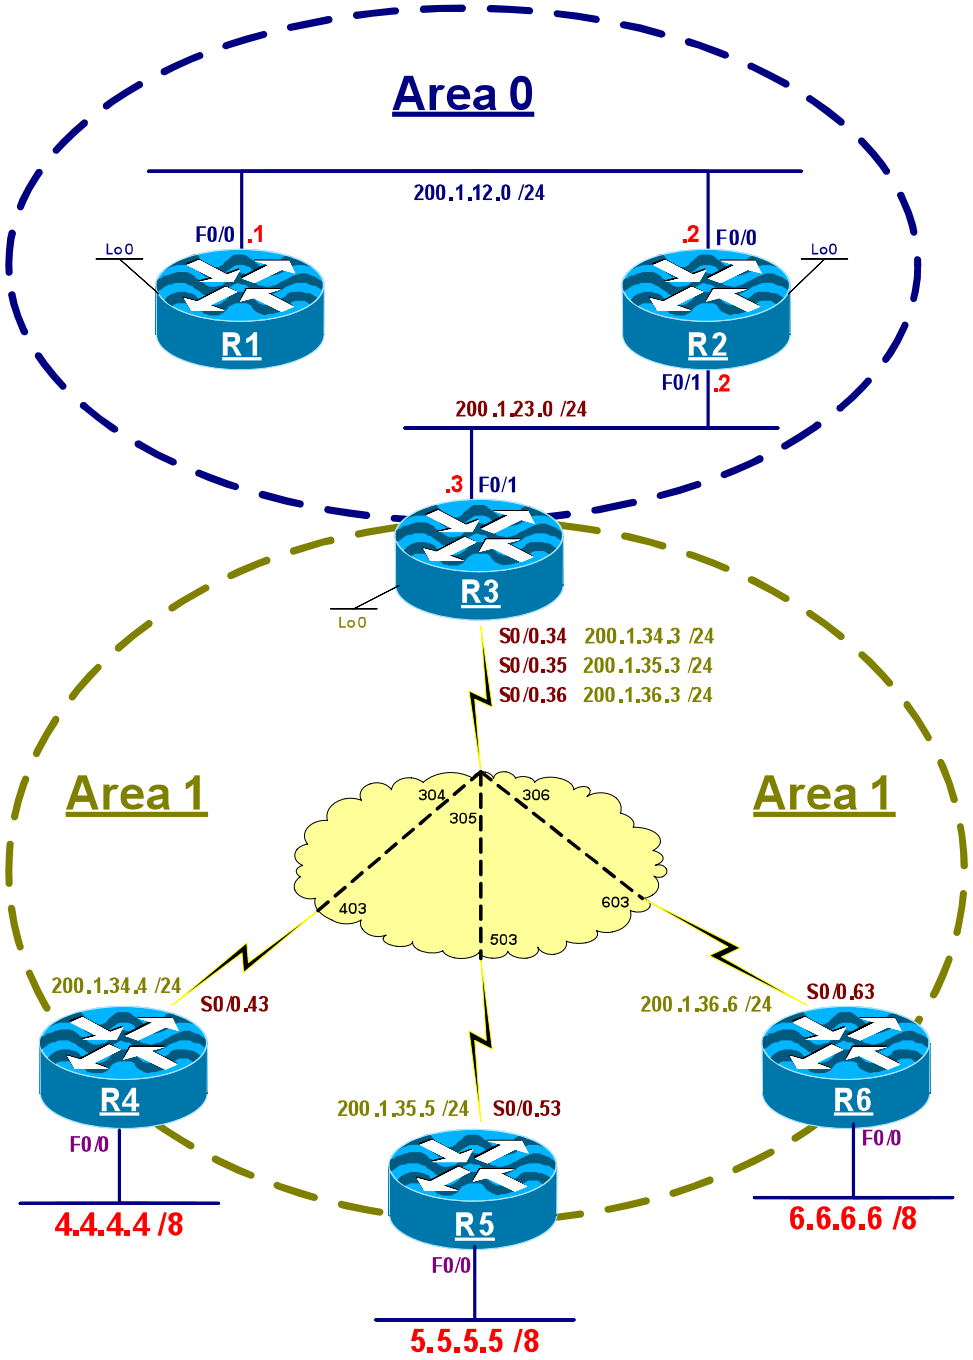
O        2.2.2.2 [110/65] via 10.1.12.2, 00:15:58, Serial0/0.12
O        3.3.3.3 [110/65] via 10.1.13.3, 00:08:55, Serial0/0.13
O        4.4.4.4 [110/65] via 10.1.14.4, 00:00:57, Serial0/0.14
```

Task 6

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 19
Forward Address Suppression



©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Lab Setup:

- Every frame-relay connection should be configured in a point-to-point manner
- Use the IP addressing scheme below for IP address assignment:
- Configure the F0/0 interface of R1 and R2 in VLAN 12 and the F0/1 interface of R2 and R3 in VLAN 23; F0/0 interface of R4, R5 and R6 should be configured in VLAN 4, 5 and 6 respectively.

IP addressing:

Router	Interface / IP address	DLCI	Connecting to
R1	F0/0 = 200.1.12.1 /24 Lo0 = 1.1.1.1 /8		R2
R2	**F0/0 = 200.1.12.2 /24** **F0/1 = 200.1.23.2 /24** **Lo0 = 2.2.2.2 /8**		**R1** **R3**
R3	F0/0 = 200.1.23.3 /24 S0/0.34 = 200.1.34.3 /24 S0/0.35 = 200.1.35.3 /24 S0/0.36 = 200.1.36.3 /24 Lo0 = 3.3.3.3 /8	 304 305 306	R2 R4 R5 R6
R4	**S0/0.43 = 200.1.34.4 /24** **F0/0 = 4.4.4.4 /8**	**403**	**R3**
R5	S0/0.53 = 200.1.35.5 /24 F0/0 = 5.5.5.5 /8	503	R3
R6	**S0/0.63 = 200.1.36.6 /24** **F0/0 = 6.6.6.6 /8**	**603**	**R3**

Task 1

Configure OSPF Area 0 on R1, R2 and advertise their directly connected interfaces in this area.

On R1: and R2

```
Rx(config)#router ospf 1
Rx(config-router)#netw 0.0.0.0 0.0.0.0 area 0
```

To verify the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1:

```
R1#Sh ip route ospf

     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/2] via 200.1.12.2, 00:00:02, FastEthernet0/0
O    200.1.23.0/24 [110/2] via 200.1.12.2, 00:00:02, FastEthernet0/0
```

Task 2

Configure R3's loopback interface and its frame-relay connection to R4, R5, and R6 in area 1, and its ethernet's connection to R2 in area 0.

On R3:

```
R3(config)#router ospf 1
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 1
R3(config-router)#netw 200.1.34.3 0.0.0.0 area 1
R3(config-router)#netw 200.1.35.3 0.0.0.0 area 1
R3(config-router)#netw 200.1.36.3 0.0.0.0 area 1
R3(config-router)#netw 200.1.23.3 0.0.0.0 area 0
```

To verify the configuration:

On R3:

```
R3#Sh ip route ospf

     1.0.0.0/32 is subnetted, 1 subnets
O       1.1.1.1 [110/3] via 200.1.23.2, 00:01:37, FastEthernet0/1
     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/2] via 200.1.23.2, 00:01:37, FastEthernet0/1
O    200.1.12.0/24 [110/2] via 200.1.23.2, 00:01:37, FastEthernet0/1
```

On R1:

```
R1#Sh ip route ospf

     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/2] via 200.1.12.2, 00:09:49, FastEthernet0/0
O IA 200.1.36.0/24 [110/66] via 200.1.12.2, 00:01:07, FastEthernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O    200.1.23.0/24 [110/2] via 200.1.12.2, 00:09:49, FastEthernet0/0
     3.0.0.0/32 is subnetted, 1 subnets
O IA    3.3.3.3 [110/3] via 200.1.12.2, 00:01:07, FastEthernet0/0
O IA 200.1.34.0/24 [110/66] via 200.1.12.2, 00:01:07, FastEthernet0/0
O IA 200.1.35.0/24 [110/66] via 200.1.12.2, 00:01:07, FastEthernet0/0
```

Task 3

Configure the frame-relay connection of R4, R5, and R6 to R3 in area 1. These routers should redistribute their F0/0 interface in OSPF routing protocol, you should **not** use an access-list or a prefix-list to accomplish this task.

On R4:

```
R4(config)#router ospf 1
R4(config-router)#netw 200.1.34.4 0.0.0.0 area 1
```

On R5:

```
R5(config)#router ospf 1
R5(config-router)#netw 200.1.35.5 0.0.0.0 area 1
```

On R6:

```
R6(config)#router ospf 1
R6(config-router)#netw 200.1.36.6 0.0.0.0 area 1
```

On R4, R5 and R6:

```
Rx(config)#route-map TST permit 10
Rx(config-route-map)#match interface F0/0

Rx(config)#router ospf 1
Rx(config-router)#redistribute connected route-map TST subnets
```

To verify the configuration:

On R1:

```
R1#Sh ip route ospf
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/2] via 200.1.12.2, 00:14:40, FastEthernet0/0
O IA 200.1.36.0/24 [110/66] via 200.1.12.2, 00:05:58, FastEthernet0/0
O    200.1.23.0/24 [110/2] via 200.1.12.2, 00:14:40, FastEthernet0/0
     3.0.0.0/32 is subnetted, 1 subnets
O IA    3.3.3.3 [110/3] via 200.1.12.2, 00:05:58, FastEthernet0/0
O E2 4.0.0.0/8 [110/20] via 200.1.12.2, 00:00:58, FastEthernet0/0
O IA 200.1.34.0/24 [110/66] via 200.1.12.2, 00:05:58, FastEthernet0/0
O E2 5.0.0.0/8 [110/20] via 200.1.12.2, 00:01:20, FastEthernet0/0
O IA 200.1.35.0/24 [110/66] via 200.1.12.2, 00:05:58, FastEthernet0/0
O E2 6.0.0.0/8 [110/20] via 200.1.12.2, 00:01:58, FastEthernet0/0
```

Task 4

Configure Area 1 as a NSSA.

On R3, R4, R5 and R6:

```
Rx(config)#router ospf 1
Rx(config-router)#area 1 nssa
```

To verify the configuration:

On R3:

```
R3#Sh ip route ospf

     1.0.0.0/32 is subnetted, 1 subnets
O       1.1.1.1 [110/3] via 200.1.23.2, 00:01:17, FastEthernet0/1
     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/2] via 200.1.23.2, 00:01:17, FastEthernet0/1
O N2 4.0.0.0/8 [110/20] via 200.1.34.4, 00:01:02, Serial0/0.34
O N2 5.0.0.0/8 [110/20] via 200.1.35.5, 00:01:02, Serial0/0.35
O N2 6.0.0.0/8 [110/20] via 200.1.36.6, 00:01:02, Serial0/0.36
O    200.1.12.0/24 [110/2] via 200.1.23.2, 00:01:17, FastEthernet0/1
```

On R1:

```
R1#Sh ip route ospf

     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/2] via 200.1.12.2, 00:20:04, FastEthernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O IA 200.1.36.0/24 [110/66] via 200.1.12.2, 00:11:22, FastEthernet0/0
O    200.1.23.0/24 [110/2] via 200.1.12.2, 00:20:04, FastEthernet0/0
     3.0.0.0/32 is subnetted, 1 subnets
O IA    3.3.3.3 [110/3] via 200.1.12.2, 00:11:22, FastEthernet0/0
O E2 4.0.0.0/8 [110/20] via 200.1.12.2, 00:03:31, FastEthernet0/0
O IA 200.1.34.0/24 [110/66] via 200.1.12.2, 00:11:22, FastEthernet0/0
O E2 5.0.0.0/8 [110/20] via 200.1.12.2, 00:03:31, FastEthernet0/0
O IA 200.1.35.0/24 [110/66] via 200.1.12.2, 00:11:22, FastEthernet0/0
O E2 6.0.0.0/8 [110/20] via 200.1.12.2, 00:03:31, FastEthernet0/0
```

On R2:

```
R2#Sh ip route ospf

     1.0.0.0/32 is subnetted, 1 subnets
O       1.1.1.1 [110/2] via 200.1.12.1, 00:21:38, FastEthernet0/0
O IA 200.1.36.0/24 [110/65] via 200.1.23.3, 00:12:56, FastEthernet0/1
     3.0.0.0/32 is subnetted, 1 subnets
O IA    3.3.3.3 [110/2] via 200.1.23.3, 00:12:56, FastEthernet0/1
O E2 4.0.0.0/8 [110/20] via 200.1.23.3, 00:05:06, FastEthernet0/1
O IA 200.1.34.0/24 [110/65] via 200.1.23.3, 00:12:56, FastEthernet0/1
O E2 5.0.0.0/8 [110/20] via 200.1.23.3, 00:05:06, FastEthernet0/1
O IA 200.1.35.0/24 [110/65] via 200.1.23.3, 00:12:56, FastEthernet0/1
O E2 6.0.0.0/8 [110/20] via 200.1.23.3, 00:05:06, FastEthernet0/1
```

To see an external route (4.0.0.0 /8)in the database:

```
R1#Sh ip ospf data external 4.0.0.0

            OSPF Router with ID (1.1.1.1) (Process ID 1)

                Type-5 AS External Link States

  Routing Bit Set on this LSA
  LS age: 397
  Options: (No TOS-capability, DC)
  LS Type: AS External Link
  Link State ID: 4.0.0.0 (External Network Number )
  Advertising Router: 3.3.3.3
  LS Seq Number: 80000001
  Checksum: 0xF9A6
  Length: 36
  Network Mask: /8
        Metric Type: 2 (Larger than any link state path)
        TOS: 0
        Metric: 20
```

The advertising router

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Forward Address: 200.1.34.4
External Route Tag: 0
```



Note forward address is NOT suppressed
It's basically the IP address of router that originated the route.

Task 5

Configure R3 to filter the following networks:
200.1.34.0 /24, 200.1.35.0 /24 and 200.1.36.0 /24

On R3:

```
R3(config)#router ospf 1
R3(config-router)#area 1 range 200.1.34.0 255.255.255.0 not-advertise
R3(config-router)#area 1 range 200.1.35.0 255.255.255.0 not-advertise
R3(config-router)#area 1 range 200.1.36.0 255.255.255.0 not-advertise
```

To verify the configuration:

On R2:

```
R2#Sh ip route ospf

     1.0.0.0/32 is subnetted, 1 subnets
O       1.1.1.1 [110/2] via 200.1.12.1, 00:26:45, FastEthernet0/0
     3.0.0.0/32 is subnetted, 1 subnets
O IA    3.3.3.3 [110/2] via 200.1.23.3, 00:18:03, FastEthernet0/1
```

On R1:

```
R1#Sh ip route ospf

     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/2] via 200.1.12.2, 00:27:08, FastEthernet0/0
O    200.1.23.0/24 [110/2] via 200.1.12.2, 00:27:08, FastEthernet0/0
     3.0.0.0/32 is subnetted, 1 subnets
O IA    3.3.3.3 [110/3] via 200.1.12.2, 00:18:26, FastEthernet0/0
```

Note: The routers in area 0 no longer have reachability to the prefixes from area 1.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 6

Configure R3 such that the routers in area 0 can reach the networks that were redistributed in task 3. Use minimum number of commands to accomplish this task; you should **not** use any global config commands as part of the solution to this task. **Do not** remove the commands from the previous steps.

On R3:

```
R3(config)#router ospf 1
R3(config-router)#area 1 nssa translate type7 suppress-fa
```

To verify the configuration:

On R2:

```
R2#Sh ip route ospf

     1.0.0.0/32 is subnetted, 1 subnets
O       1.1.1.1 [110/2] via 200.1.12.1, 00:29:06, FastEthernet0/0
     3.0.0.0/32 is subnetted, 1 subnets
O IA    3.3.3.3 [110/2] via 200.1.23.3, 00:20:25, FastEthernet0/1
O E2 4.0.0.0/8 [110/20] via 200.1.23.3, 00:00:27, FastEthernet0/1
O E2 5.0.0.0/8 [110/20] via 200.1.23.3, 00:00:27, FastEthernet0/1
O E2 6.0.0.0/8 [110/20] via 200.1.23.3, 00:00:27, FastEthernet0/1
```

On R1:

```
R1#Show ip route ospf

     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/2] via 200.1.12.2, 00:29:43, FastEthernet0/0
O    200.1.23.0/24 [110/2] via 200.1.12.2, 00:29:43, FastEthernet0/0
     3.0.0.0/32 is subnetted, 1 subnets
O IA    3.3.3.3 [110/3] via 200.1.12.2, 00:21:01, FastEthernet0/0
O E2 4.0.0.0/8 [110/20] via 200.1.12.2, 00:01:04, FastEthernet0/0
O E2 5.0.0.0/8 [110/20] via 200.1.12.2, 00:01:04, FastEthernet0/0
O E2 6.0.0.0/8 [110/20] via 200.1.12.2, 00:01:04, FastEthernet0/0
```

To test the configuration:

On R1:

```
R1#Ping 3.3.3.3
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/113/116 ms

R1#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 168/171/173 ms

R1#Ping 5.5.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 168/170/173 ms
```

The OSPF forwarding address suppression in translated type-5 LSAs feature causes an NSSA ABR to translate type-7 LSAs to type-5 LSAs, but use the 0.0.0.0 as the forwarding address instead of that specified in the type-7 LSA.

After the filters are applied and the "area 1 nssa translate type7 suppress-fa" command is configured, the output of the "show ip ospf da ext 4.0.0.0" will be changed to following:

```
R1#Sh ip ospf data external 4.0.0.0


            OSPF Router with ID (1.1.1.1) (Process ID 1)

                Type-5 AS External Link States

  Routing Bit Set on this LSA
  LS age: 297
  Options: (No TOS-capability, DC)
  LS Type: AS External Link
  Link State ID: 4.0.0.0 (External Network Number )
  Advertising Router: 3.3.3.3
  LS Seq Number: 80000002
  Checksum: 0x5738
  Length: 36
  Network Mask: /8
        Metric Type: 2 (Larger than any link state path)
        TOS: 0
        Metric: 20
        Forward Address: 0.0.0.0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
        External Route Tag: 0
```

Because the IP address of the router that originated the route/s are suppressed, area 0 routers no longer need to maintain extra prefixes in their routing table.

On R2:

```
R2#Sh ip route ospf

     1.0.0.0/32 is subnetted, 1 subnets
O       1.1.1.1 [110/2] via 200.1.12.1, 00:34:43, FastEthernet0/0
     3.0.0.0/32 is subnetted, 1 subnets
O IA    3.3.3.3 [110/2] via 200.1.23.3, 00:26:01, FastEthernet0/1
O E2 4.0.0.0/8 [110/20] via 200.1.23.3, 00:06:04, FastEthernet0/1
O E2 5.0.0.0/8 [110/20] via 200.1.23.3, 00:06:04, FastEthernet0/1
O E2 6.0.0.0/8 [110/20] via 200.1.23.3, 00:06:04, FastEthernet0/1
```

On R1:

```
R1#Show ip route ospf

     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/2] via 200.1.12.2, 00:35:35, FastEthernet0/0
O    200.1.23.0/24 [110/2] via 200.1.12.2, 00:35:35, FastEthernet0/0
     3.0.0.0/32 is subnetted, 1 subnets
O IA    3.3.3.3 [110/3] via 200.1.12.2, 00:26:53, FastEthernet0/0
O E2 4.0.0.0/8 [110/20] via 200.1.12.2, 00:06:56, FastEthernet0/0
O E2 5.0.0.0/8 [110/20] via 200.1.12.2, 00:06:56, FastEthernet0/0
O E2 6.0.0.0/8 [110/20] via 200.1.12.2, 00:06:56, FastEthernet0/0
```

Note: The backbone routers no longer need to maintain the extra prefixes for the links in area 1, but they have full reachability to the prefixes that were redistributed.

Task 7

Configure the appropriate router such that the routers in area 1 do not have the existing and future internal prefix/es that are advertised in area 0, but they should have reachability to all inter-area prefixes. Use minimum number of commands to accomplish this task. **Do not** use access-lists, prefix-lists or route-map/s.

Let's check the routing table of one of the routers in area 1:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R6:

```
R6#Show ip route ospf | inc O

O IA    1.1.1.1 [110/67] via 200.1.36.3, 00:04:02, Serial0/0.63
O IA    2.2.2.2 [110/66] via 200.1.36.3, 00:04:02, Serial0/0.63
O IA 200.1.23.0/24 [110/65] via 200.1.36.3, 00:04:02, Serial0/0.63
O       3.3.3.3 [110/65] via 200.1.36.3, 00:04:02, Serial0/0.63
O N2 4.0.0.0/8 [110/20] via 200.1.36.3, 00:04:02, Serial0/0.63
O    200.1.34.0/24 [110/128] via 200.1.36.3, 00:04:02, Serial0/0.63
O N2 5.0.0.0/8 [110/20] via 200.1.36.3, 00:04:02, Serial0/0.63
O    200.1.35.0/24 [110/128] via 200.1.36.3, 00:04:02, Serial0/0.63
O IA 200.1.12.0/24 [110/66] via 200.1.36.3, 00:04:02, Serial0/0.63
```

Note: The routers in area 1 can see inter-area routes. Let's configure this task:

On R3:

```
R3(config)#router ospf 1
R3(config-router)#area 1 nssa translate type7 suppress-fa No-summary
```

The above command filters all LSA type-3s and injects a default route in area 1 instead.

To verify the configuration:

On R6:

```
R6#Show ip route ospf | inc O

O       3.3.3.3 [110/65] via 200.1.36.3, 00:10:06, Serial0/0.63
O N2 4.0.0.0/8 [110/20] via 200.1.36.3, 00:10:06, Serial0/0.63
O    200.1.34.0/24 [110/128] via 200.1.36.3, 00:10:06, Serial0/0.63
O N2 5.0.0.0/8 [110/20] via 200.1.36.3, 00:10:06, Serial0/0.63
O    200.1.35.0/24 [110/128] via 200.1.36.3, 00:10:06, Serial0/0.63
O*IA 0.0.0.0/0 [110/65] via 200.1.36.3, 00:01:41, Serial0/0.63
```

To test the configuration:

On R6:

```
R6#Ping 1.1.1.1 source f0/0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 6.6.6.6
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

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

Note: When the "no-summary" keyword was added, the "area 1 NSSA no-summary" command was added to the OSPF configuration of R3:

To verify:

On R3:

```
R3#Show run | s router ospf

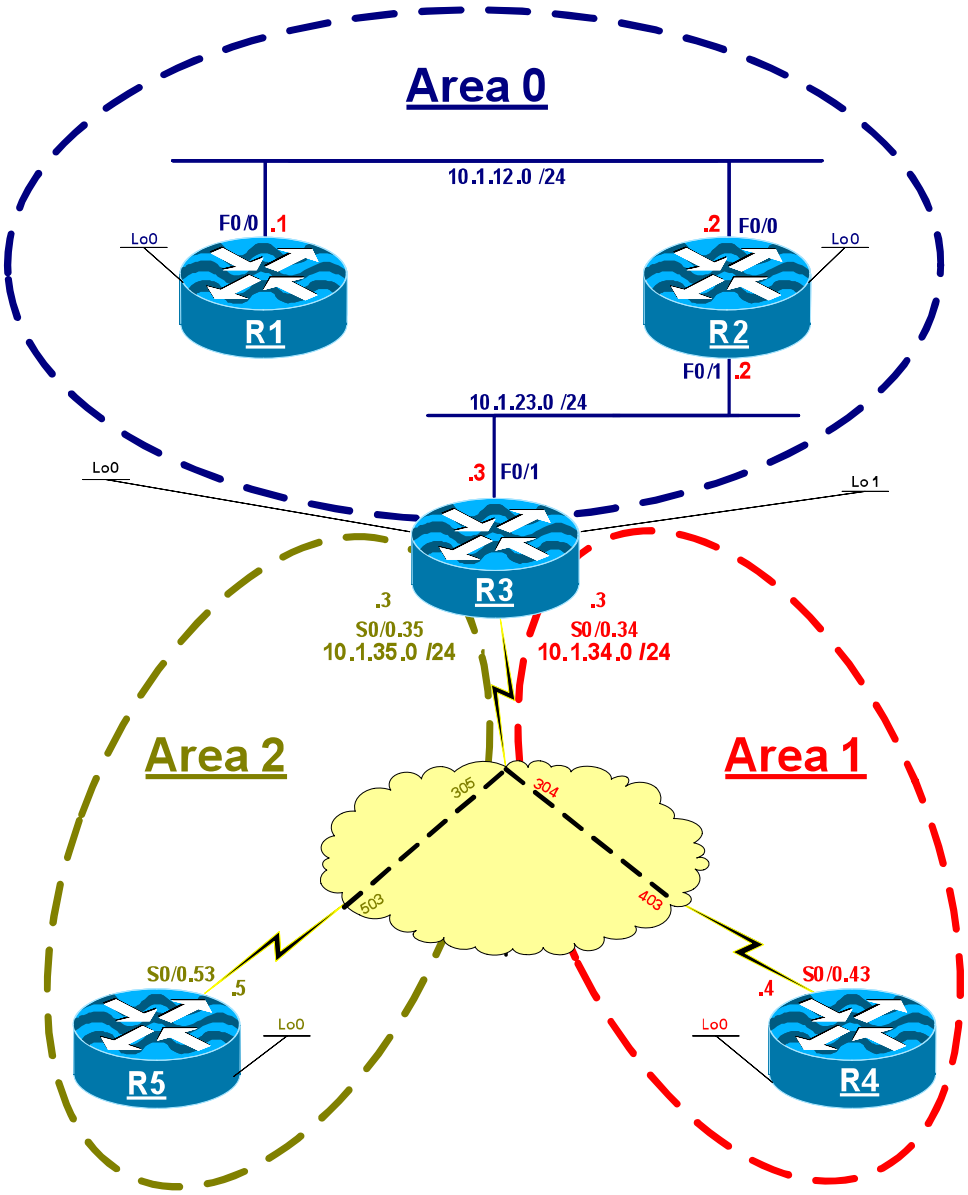
router ospf 1
 log-adjacency-changes
 area 1 nssa no-summary
 area 1 nssa translate type7 suppress-fa
 area 1 range 200.1.34.0 255.255.255.0 not-advertise
 area 1 range 200.1.35.0 255.255.255.0 not-advertise
 area 1 range 200.1.36.0 255.255.255.0 not-advertise
 network 3.3.3.3 0.0.0.0 area 1
 network 200.1.23.3 0.0.0.0 area 0
 network 200.1.34.3 0.0.0.0 area 1
 network 200.1.35.3 0.0.0.0 area 1
 network 200.1.36.3 0.0.0.0 area 1
```

Task 8

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 20 – OSPF NSSA No-Redistribution & Injection of Default Routes



©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Lab Setup:

- Configure all frame-relay connections in a point-to-point manner.
- Use the IP addressing and DLCI chart below.

IP addressing:

Router	Interface / IP address	DLCI assignment
R1	F0/0 = 10.1.12.1 /24 Loopback0 = 1.1.1.1 /8	
R2	**F0/0 = 10.1.12.2 /24** **F0/1 = 10.1.23.2 /24** **Loopback0 = 2.2.2.2 /8**	
R3	F0/1 = 10.1.23.3 /24 S0/0.34 = 10.1.34.3 /24 S0/0.35 = 10.1.35.3 /24 Loopback0 = 3.3.3.3 /8 Loopback1 = 33.3.3.3 /8	 304 305
R4	**S0/0.43 = 10.1.34.4 /24** **Loopback0 = 4.4.4.4 /8**	**403**
R5	S0/0.53 = 10.1.35.5 /24 Loopback0 = 5.5.5.5 /8	503

Task 1

Configure OSPF on the routers based on the following chart:

Router	Interface / Area
R1	F0/0 / **Area 0** Loopback0 / **Area 0**
R2	F0/0 / **Area 0** F0/1 / **Area 0** Loopback0 / **Area 0**
R3	F0/1 / **Area 0** S0/0.34 / **Area 1** S0/0.35 / **Area 2**
R4	S0/0.43 / **Area 1** Loopback0 / **Area 1**
R5	S0/0.53 / **Area 2** Loopback0 / **Area 2**

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 2

Configure R3 to redistribute its loopback 0 and 1 interfaces into this OSPF routing domain.

A route-map is configured to match on Lo0 or Lo1:

On R3:

```
R3(config)#route-map TST permit 10
R3(config-route-map)#match interface lo0 lo1
```

The connected interfaces are redistributed referencing the route-map:

```
R3(config)#router ospf 1
R3(config-router)#redistribute connected subnets route-map TST
```

Let's check the routing table of R1 and R5:

To verify the configuration:

On R1:

```
R1#Show ip route ospf | inc E2

O E2 33.0.0.0/8 [110/20] via 10.1.12.2, 00:00:44, FastEthernet0/0
O E2 3.0.0.0/8 [110/20] via 10.1.12.2, 00:00:44, FastEthernet0/0
```

On R5:

```
R5#Show ip route ospf | inc E2

O E2 33.0.0.0/8 [110/20] via 10.1.35.3, 00:01:07, Serial0/0.53
O E2 3.0.0.0/8 [110/20] via 10.1.35.3, 00:01:07, Serial0/0.53
```

Task 3

Configure area 1 and area 2 as NSSA, R3 should be configured such that the routers in these two areas receive a default route, this default route should be injected as an external route.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3:

```
R3(config)#router ospf 1
R3(config-router)#area 1 nssa default-information-originate
R3(config-router)#area 2 nssa default-information-originate
```

On R4:

```
R4(config)#router ospf 1
R4(config-router)#area 1 nssa
```

On R5:

```
R5(config)#router ospf 1
R5(config-router)#area 2 nssa
```

To verify the configuration:

On R4:

```
R4#Show ip route ospf | inc O

O IA    1.1.1.1 [110/67] via 10.1.34.3, 00:03:45, Serial0/0.43
O IA    2.2.2.2 [110/66] via 10.1.34.3, 00:03:45, Serial0/0.43
O N2 33.0.0.0/8 [110/20] via 10.1.34.3, 00:03:45, Serial0/0.43
O N2 3.0.0.0/8 [110/20] via 10.1.34.3, 00:03:45, Serial0/0.43
O IA    5.5.5.5 [110/129] via 10.1.34.3, 00:03:42, Serial0/0.43
O IA    10.1.12.0 [110/66] via 10.1.34.3, 00:03:45, Serial0/0.43
O IA    10.1.23.0 [110/65] via 10.1.34.3, 00:03:45, Serial0/0.43
O IA    10.1.35.0 [110/128] via 10.1.34.3, 00:03:45, Serial0/0.43
O*N2 0.0.0.0/0 [110/1] via 10.1.34.3, 00:03:45, Serial0/0.43
```

Note: The default route is injected as an external route.

On R5:

```
R5#Show ip route ospf | inc O

O IA    1.1.1.1 [110/67] via 10.1.35.3, 00:04:04, Serial0/0.53
O IA    2.2.2.2 [110/66] via 10.1.35.3, 00:04:04, Serial0/0.53
O N2 33.0.0.0/8 [110/20] via 10.1.35.3, 00:04:04, Serial0/0.53
O N2 3.0.0.0/8 [110/20] via 10.1.35.3, 00:04:04, Serial0/0.53
O IA    4.4.4.4 [110/129] via 10.1.35.3, 00:04:04, Serial0/0.53
O IA    10.1.12.0 [110/66] via 10.1.35.3, 00:04:04, Serial0/0.53
O IA    10.1.23.0 [110/65] via 10.1.35.3, 00:04:04, Serial0/0.53
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O IA    10.1.34.0 [110/128] via 10.1.35.3, 00:04:04, Serial0/0.53
O*N2 0.0.0.0/0 [110/1] via 10.1.35.3, 00:04:04, Serial0/0.53
```

Task 4

Configure area 1 such that it receives the default route injected by the ABR as an internal OSPF route.

On R3:

```
R3(config)#router ospf 1
R3(config-router)#area 1 nssa no-summary
```

To Verify the configuration:

On R4:

```
R4#Sh ip rou ospf | inc O

O N2 33.0.0.0/8 [110/20] via 10.1.34.3, 00:05:46, Serial0/0.43
O N2 3.0.0.0/8 [110/20] via 10.1.34.3, 00:05:46, Serial0/0.43
O*IA 0.0.0.0/0 [110/65] via 10.1.34.3, 00:00:07, Serial0/0.43
```

Note: The default route injected by the ABR of this area is an internal OSPF route.

Task 5

Configure R3 such that **only** area 2 receives the redistributed routes (3.0.0.0 /8 and 33.0.0.0 /8); you should **not** use any global configuration command or route-map as part of the solution to accomplish this task.

On R3:

```
R3(config)#router ospf 1
R3(config-router)#area 1 nssa no-redistribution
```

To verify the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4:

```
R4#Sh ip route ospf | inc O

O*IA 0.0.0.0/0 [110/65] via 10.1.34.3, 00:01:45, Serial0/0.43
```

You may have to "clear ip ospf proc" on R3 to see the above result. The "no-redistribution" is configured on the ABR, which happens to be an ASBR as well; this command stops redistribution of the external routes into the area specified.

Task 6

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Advanced
CCIE SERVICE PROVIDER
v3.0

www.MicronicsTraining.com

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

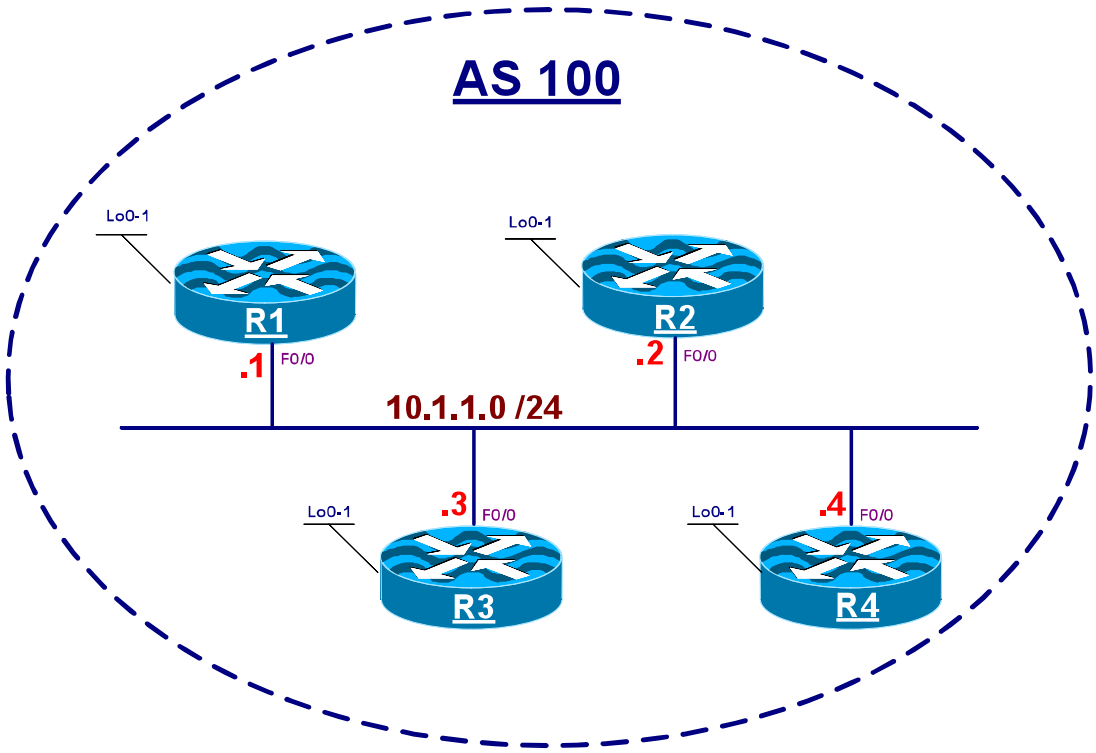
Paul Negron
CCIE #14856
SP

BGP

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Lab 1 – Establishing Neighbor Adjacency



Lab Setup:

- The F0/0 interface of these four routers should be configured in VLAN 100.
- Configure the routers according to the following IP addressing chart:

IP Addressing:

Router	Interface / IP Address	AS 100
R1	**F0/0 – 10.1.1.1 /24** **Lo0 – 1.1.1.1 /8** **Lo1 – 192.168.1.1 /24**	**100**
R2	F0/0 – 10.1.1.2 /24 Lo0 – 2.2.2.2 /8 Lo1 – 192.168.2.2 /24	100
R3	**F0/0 – 10.1.1.3 /24**	**100**

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

	Lo0 – 3.3.3.3 /8 **Lo1 – 192.168.3.3 /24**	
R4	F0/0 – 10.1.1.4 /24 Lo0 – 4.4.4.4 /8 Lo1 – 192.168.4.4 /24	100

Task 1

Configure routers R1 – R4 in AS 100, these routers should create an IBGP peer sessions between them, ensure that these routers advertise their loopback 0 interface in this AS.

On R1:

```
R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.1.2 remote-as 100
R1(config-router)#neighbor 10.1.1.3 remote-as 100
R1(config-router)#neighbor 10.1.1.4 remote-as 100
R1(config-router)#No syn
R1(config-router)#network 1.0.0.0
```

On R2:

```
R2(config)#router bgp 100
R2(config-router)#neighbor 10.1.1.1 remote-as 100
R2(config-router)#neighbor 10.1.1.3 remote-as 100
R2(config-router)#neighbor 10.1.1.4 remote-as 100
R2(config-router)#No syn
R2(config-router)#network 2.0.0.0
```

On R3:

```
R3(config)#router bgp 100
R3(config-router)#neighbor 10.1.1.1 remote-as 100
R3(config-router)#neighbor 10.1.1.2 remote-as 100
R3(config-router)#neighbor 10.1.1.4 remote-as 100
R3(config-router)#No syn
R3(config-router)#network 3.0.0.0
```

On R4:

```
R1(config)#router bgp 100
R4(config-router)#neighbor 10.1.1.1 remote-as 100
R4(config-router)#neighbor 10.1.1.2 remote-as 100
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config-router)#neighbor 10.1.1.3 remote-as 100
R4(config-router)#No syn
R4(config-router)#network 4.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip bgp

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*>i2.0.0.0          10.1.1.2                 0    100      0 i
*>i3.0.0.0          10.1.1.3                 0    100      0 i
```

The following explains the output of the show command:

- **The "*"** means that the entry in the table is valid.
- **The ">"** means that this is the best entry for the prefix.
- **The "i"** means that the entry was learned via an IBGP, this is the "i" to the left of the network column. The letter "i" under the path column, specifies the origin of the route.
- **The "network" column** identifies the prefix entry for the network, if the mask is omitted, the default mask is assumed.
- **The "next hop" column** identifies the next hop's IP address to receive the specified network address, if its "0.0.0.0" it identifies a prefix that the local router advertised.
- **The "metric" column** identifies the inter-as metric, or the MED attribute which is set to zero by default.
- **The "LocPrf" column** identifies the local preference attribute; this is used in the route selection process carried within the local AS ONLY. With the local-pref attribute, the higher value has more preference. The prefixes that are received from a peer AS are tagged with a local-pref value of 100; this value can be changed to influence the best path selection process. The changed value is only advertised to IBGP peers. When the local router advertises a prefix, no local-pref value is seen in the output of the "Show ip bgp" command. The default value of 100 can be changed by the **"BGP default local-preference"** command.
- **The "weight" column** identifies the weight attribute; the prefixes that are received via a neighbor (IBGP or EBGP) will have a weight attribute of 0, but the prefixes that are originated by the local router will have a weight value of 32768. This attribute overrides any other attribute for performing best path determination.
- **The "path" column** identifies the path attribute. If the prefixes were originated or learned via an IBGP neighbor, the path column will have the letter "i" without any ASN. If the prefix was learned

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

through another AS, then this column will have the AS number/s followed by the letter i, the ASNs indicate which AS advertised the prefix. The maximum number of ASes that a prefix can traverse through is 255.

Task 2

Reconfigure the routers as follows:

R2, R3 and R4 should be configured in AS 200, 300 and 400 respectively. Configure a full mesh peer session between these routers.

On R1:

```
R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.1.2 remote-as 200
R1(config-router)#neighbor 10.1.1.3 remote-as 300
R1(config-router)#neighbor 10.1.1.4 remote-as 400
R1(config-router)#No auto
R1(config-router)#network 1.0.0.0
```

On R2:

```
R2(config)#router bgp 200
R2(config-router)#neighbor 10.1.1.1 remote-as 100
R2(config-router)#neighbor 10.1.1.3 remote-as 300
R2(config-router)#neighbor 10.1.1.4 remote-as 400
R2(config-router)#No auto
R2(config-router)#network 2.0.0.0
```

On R3:

```
R3(config)#router bgp 300
R3(config-router)#neighbor 10.1.1.1 remote-as 100
R3(config-router)#neighbor 10.1.1.2 remote-as 200
R3(config-router)#neighbor 10.1.1.4 remote-as 400
R3(config-router)#No auto
R3(config-router)#network 3.0.0.0
```

On R4:

```
R1(config)#router bgp 400
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config-router)#neighbor 10.1.1.1 remote-as 100
R4(config-router)#neighbor 10.1.1.2 remote-as 200
R4(config-router)#neighbor 10.1.1.3 remote-as 300
R4(config-router)#No auto
R4(config-router)#network 4.0.0.0
```

To verify the configuration:

On R1:

```
R1#show ip bgp | b Net
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*  2.0.0.0          10.1.1.2                               0 400 200 i
*                   10.1.1.2                               0 300 200 i
*>                  10.1.1.2                 0             0 200 i
*  3.0.0.0          10.1.1.3                               0 400 300 i
*                   10.1.1.3                               0 200 300 i
*>                  10.1.1.3                 0             0 300 i
*  4.0.0.0          10.1.1.4                               0 200 400 i
*                   10.1.1.4                               0 300 400 i
*>                  10.1.1.4                 0             0 400 i
```

Note: The local-preference attribute is not assigned on any of the prefixes; this is because the prefixes are advertised by an EBGP peer. The best selection in the above output is based on the shortest ASN.

The MED value (metric column) is zero for some of the prefixes, and on others, it is not assigned, this is because when the prefix is advertised by the originating AS, the metric is set to "0", but when the same prefix is advertised by another AS, the MED value is removed.

Task 3

Reconfigure all the routers in AS 100; use the following policy for their IBGP peer sessions:

- Authentication must be enabled between the peers using "cisco" as the password.
- The peer session must be established based on the loopback 0's IP address.
- These routers should **only** advertise their loopback 1 in BGP.
- Provide NLRI to loopback0 interface using RIPv2.
- The peer session between the routers should only be established if they are running BGP version 4.
- Use peer-groups to accomplish this task.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Cisco's implementation of BGP in IOS 12.0(5)T or earlier releases supports BGP versions 2, 3, and 4, with dynamic negotiation down to Version 2. However, in IOS version 12.0(6)T or later, Cisco routers only support version 4 and they do not support dynamic negotiation down to Version 2. The reason you may see the "neighbor version" command configured on some Cisco routers is because may be the router is connecting and establishing a peer session with a Non-Cisco router, or the administrator is not aware of this fact.

On R1:

```
R1(config)#router bgp 100
R1(config-router)#No au
R1(config-router)#net 192.168.1.0

R1(config-router)#neighbor TST peer-group
R1(config-router)#neighbor TST remote-as 100
R1(config-router)#neighbor TST update-source lo0
R1(config-router)#neighbor TST version 4
R1(config-router)#neighbor TST password cisco

R1(config-router)#neighbor 2.2.2.2 peer-group TST
R1(config-router)#neighbor 3.3.3.3 peer-group TST
R1(config-router)#neighbor 4.4.4.4 peer-group TST

R1(config-router)#router rip
R1(config-router)#No au
R1(config-router)#ver 2
R1(config-router)#netw 10.0.0.0
R1(config-router)#netw 1.0.0.0
```

On R2:

```
R2(config)#router bgp 100
R2(config-router)#No au
R2(config-router)#netw 192.168.2.0

R2(config-router)#neighbor TST peer-group
R2(config-router)#neighbor TST remote-as 100
R2(config-router)#neighbor TST update-source lo0
R2(config-router)#neighbor TST version 4
R2(config-router)#neighbor TST password cisco
R2(config-router)#neighbor 1.1.1.1 peer-group TST
R2(config-router)#neighbor 3.3.3.3 peer-group TST
R2(config-router)#neighbor 4.4.4.4 peer-group TST

R2(config-router)#router rip
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config-router)#No au
R2(config-router)#ver 2
R2(config-router)#netw 10.0.0.0
R2(config-router)#netw 2.0.0.0
```

On R3:

```
R3(config)#router bgp 100
R3(config-router)#No au
R3(config-router)# network 192.168.3.0

R3(config-router)#neighbor TST peer-group
R3(config-router)#neighbor TST remote-as 100
R3(config-router)#neighbor TST update-source lo0
R3(config-router)#neighbor TST version 4
R3(config-router)#neighbor TST password cisco
R3(config-router)#neighbor 1.1.1.1 peer-group TST
R3(config-router)#neighbor 2.2.2.2 peer-group TST
R3(config-router)#neighbor 4.4.4.4 peer-group TST

R3(config-router)#router rip
R3(config-router)#No au
R3(config-router)#ver 2
R3(config-router)#netw 10.0.0.0
R3(config-router)#netw 3.0.0.0
```

On R4:

```
R4(config)#router bgp 100
R4(config-router)#No au
R4(config-router)# network 192.168.4.0

R4(config-router)#neighbor TST peer-group
R4(config-router)#neighbor TST remote-as 100
R4(config-router)#neighbor TST update-source lo0
R4(config-router)#neighbor TST version 4
R4(config-router)#neighbor TST password cisco
R4(config-router)#neighbor 1.1.1.1 peer-group TST
R4(config-router)#neighbor 2.2.2.2 peer-group TST
R4(config-router)#neighbor 3.3.3.3 peer-group TST

R4(config-router)#router rip
R4(config-router)#No au
R4(config-router)#ver 2
R4(config-router)#netw 10.0.0.0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config-router)#netw 4.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 192.168.1.0      0.0.0.0                  0         32768 i
*>i192.168.2.0      2.2.2.2                  0    100      0 i
*>i192.168.3.0      3.3.3.3                  0    100      0 i
*>i192.168.4.0      4.4.4.4                  0    100      0 i
```

To verify the configuration:

On R1:

```
R1#Show ip bgp peer-group TST

BGP peer-group is TST,  remote AS 100
  BGP version 4
  Default minimum time between advertisement runs is 0 seconds
 For address family: IPv4 Unicast
  BGP neighbor is TST, peer-group internal, members:
  2.2.2.2 3.3.3.3 4.4.4.4
  Index 0, Offset 0, Mask 0x0
  Update messages formatted 0, replicated 0
  Number of NLRIs in the update sent: max 0, min 0
```

Note: The output of the "show ip bgp peer-group TST" reveals the ip address of the members of the peer-group. One of the major benefits of a peer-group is that it can reduce the administrative overhead by cutting down redundant configuration on the routers.

Task 4

Remove the BGP configuration from the routers and reconfigure all four routers in AS 100 using **peer-session templates**; you should configure the following two templates to accomplish this task:

- Common Template: This template should contain the "**neighbor version 4**" and "**neighbor password**" command, this template should be applied to all neighbors.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

- IBGP Template: This template should contain the "**neighbor update-source**" and "**neighbor remote-as**" commands. This template should be applied to all IBGP neighbors.

You should advertise Loopback1 interface in BGP and loopback0 should be used as the IP address for establishing the peer sessions. **Do not** remove RIPv2's configuration.

On R1:

```
R1(config)#router bgp 100
R1(config-router)#No au
R1(config-router)#network 192.168.1.0

R1(config-router)#template peer-session Common
R1(config-router-stmp)#password cisco
R1(config-router-stmp)#version 4
R1(config-router-stmp)#exit-peer-session

R1(config-router)#template peer-session IBGP
R1(config-router-stmp)#inherit peer-session Common
R1(config-router-stmp)#update-source lo0
R1(config-router-stmp)#remote-as 100
R1(config-router-stmp)#exit-peer-session

R1(config-router)#neighbor 2.2.2.2 inherit peer-session IBGP
R1(config-router)#neighbor 3.3.3.3 inherit peer-session IBGP
R1(config-router)#neighbor 4.4.4.4 inherit peer-session IBGP
```

On R2:

```
R2(config)#router bgp 100
R2(config-router)#No au
R2(config-router)#network 192.168.2.0

R2(config-router)#template peer-session Common
R2(config-router-stmp)#password cisco
R2(config-router-stmp)#version 4
R2(config-router-stmp)#exit-peer-session

R2(config-router)#template peer-session IBGP
R2(config-router-stmp)#inherit peer-session Common
R2(config-router-stmp)#update-source lo0
R2(config-router-stmp)#remote-as 100
R2(config-router-stmp)#exit-peer-session
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config-router)#neighbor 1.1.1.1 inherit peer-session IBGP
R2(config-router)#neighbor 3.3.3.3 inherit peer-session IBGP
R2(config-router)#neighbor 4.4.4.4 inherit peer-session IBGP
```

On R3:

```
R3(config)#router bgp 100
R3(config-router)#No au
R3(config-router)#network 192.168.3.0

R3(config-router)#template peer-session Common
R3(config-router-stmp)#password cisco
R3(config-router-stmp)#version 4
R3(config-router-stmp)#exit-peer-session

R3(config-router)#template peer-session IBGP
R3(config-router-stmp)#inherit peer-session Common
R3(config-router-stmp)#update-source lo0
R3(config-router-stmp)#remote-as 100
R3(config-router-stmp)#exit-peer-session

R3(config-router)#neighbor 1.1.1.1 inherit peer-session IBGP
R3(config-router)#neighbor 2.2.2.2 inherit peer-session IBGP
R3(config-router)#neighbor 4.4.4.4 inherit peer-session IBGP
```

On R4:

```
R4(config)#router bgp 100
R4(config-router)#No au
R4(config-router)#network 192.168.4.0

R4(config-router)#template peer-session Common
R4(config-router-stmp)#password cisco
R4(config-router-stmp)#version 4

R4(config-router-stmp)#exit-peer-session

R4(config-router)#template peer-session IBGP
R4(config-router-stmp)#inherit peer-session Common
R4(config-router-stmp)#update-source lo0
R4(config-router-stmp)#remote-as 100
R4(config-router-stmp)#exit-peer-session

R4(config-router)#neighbor 1.1.1.1 inherit peer-session IBGP
R4(config-router)#neighbor 2.2.2.2 inherit peer-session IBGP
R4(config-router)#neighbor 3.3.3.3 inherit peer-session IBGP
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the configuration:

On R1:

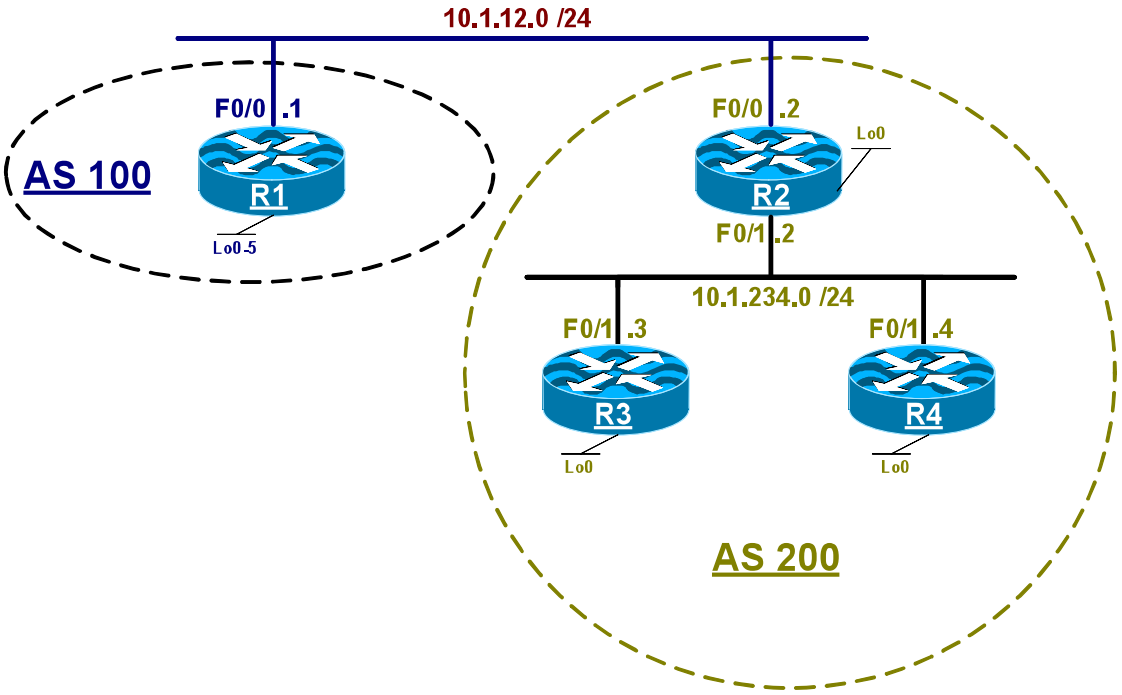
```
R1#Show ip bgp | b Network
   Network          Next Hop            Metric LocPrf Weight Path
*> 192.168.1.0      0.0.0.0                  0         32768 i
*>i192.168.2.0      2.2.2.2                  0    100      0 i
*>i192.168.3.0      3.3.3.3                  0    100      0 i
*>i192.168.4.0      4.4.4.4                  0    100      0 i
```

- **Peer-session template can be used to apply session specific configuration command/s to a group of neighbors that share a common session configuration.**
- **Peer-session templates can be reused and they support inheritance of another peer-session template/s, this means that nested peer-sessions can also be used.**
- **Peer-session templates support session specific commands only.**

Task 5

Reconfigure the routers based on the following IP address space and diagram.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh



Lab Setup:

- Configure F0/0 interface of R1 and R2 are in VLAN 12 and the F0/1 interface of R2, R3 and R4 in VLAN 234.
- Configure IP addressing on the routers using the IP addressing chart on the next page.

IP Addressing chart:

Router	AS number	Interface / IP address
R1	AS 100	Lo0 = 1.1.0.1 /24 Lo1 = 1.1.1.1 /24 Lo2 = 1.1.2.1 /24 Lo3 = 1.1.3.1 /24 Lo4 = 100.1.1.1 /24 Lo5 = 100.2.2.1 /24 F0/0 = 10.1.12.1 /24
R2	AS 200	Lo0 = 2.2.2.2 /8 F0/0 = 10.1.12.2 /24 F0/1 = 10.1.234.2 /24
R3	AS 200	Lo0 = 3.3.3.3 /8 F0/1 = 10.1.234.3 /24
R4	AS 200	Lo0 = 4.4.4.4 /8 F0/1 = 10.1.234.4 /24

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 6

- R1 in AS 100 should establish an EBGP peer session with R2 in AS 200. R1 should advertise all of its loopback interfaces in AS 100.

- R2, R3 and R4 should be configured in AS 200; these routers should establish IBGP peer sessions between them and advertise their loopback 0 interface in AS 200.
- Configure the router-ids of the routers as follows:

  R1 = 10.1.1.1, R2 = 10.2.2.2, R3 = 10.3.3.3 and R4 = 10.4.4.4

On R1:

```
R1(config)#router bgp 100
R1(config-router)#No synchronization
R1(config-router)#bgp router-id 10.1.1.1
R1(config-router)#network 1.1.0.0 mask 255.255.255.0
R1(config-router)#network 1.1.1.0 mask 255.255.255.0
R1(config-router)#network 1.1.2.0 mask 255.255.255.0
R1(config-router)#network 1.1.3.0 mask 255.255.255.0
R1(config-router)#network 100.1.1.0 mask 255.255.255.0
R1(config-router)#network 100.2.2.0 mask 255.255.255.0
R1(config-router)#neighbor 10.1.12.2 remote-as 200
R1(config-router)#No auto-summary
```

On R2:

```
R2(config)#router bgp 200
R2(config-router)#No synchronization
R2(config-router)#bgp router-id 10.2.2.2
R2(config-router)#network 2.0.0.0
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#neighbor 10.1.234.3 remote-as 200
R2(config-router)#neighbor 10.1.234.4 remote-as 200
R2(config-router)#No auto-summary
```

On R3:

```
R3(config)#router bgp 200
R3(config-router)#No synchronization
R3(config-router)#bgp router-id 10.3.3.3
R3(config-router)#network 3.0.0.0
R3(config-router)#neighbor 10.1.234.2 remote-as 200
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-router)#neighbor 10.1.234.4 remote-as 200
R3(config-router)#No auto-summary
```

On R4:

```
R4(config)#router bgp 200
R4(config-router)#No synchronization
R4(config-router)#bgp router-id 10.4.4.4
R4(config-router)#network 4.0.0.0
R4(config-router)#neighbor 10.1.234.2 remote-as 200
R4(config-router)#neighbor 10.1.234.3 remote-as 200
R4(config-router)#No auto-summary
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 1.1.0.0/24	0.0.0.0	0		32768	i
*> 1.1.1.0/24	0.0.0.0	0		32768	i
*> 1.1.2.0/24	0.0.0.0	0		32768	i
*> 1.1.3.0/24	0.0.0.0	0		32768	i
*> 2.0.0.0	10.1.12.2	0		0	200 i
*> 3.0.0.0	10.1.12.2			0	200 i
*> 4.0.0.0	10.1.12.2			0	200 i
*> 100.1.1.0/24	0.0.0.0	0		32768	i
*> 100.2.2.0/24	0.0.0.0	0		32768	i

On R2:

```
R2#Show ip bgp | b Network
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 1.1.0.0/24	10.1.12.1	0		0	100 i
*> 1.1.1.0/24	10.1.12.1	0		0	100 i
*> 1.1.2.0/24	10.1.12.1	0		0	100 i
*> 1.1.3.0/24	10.1.12.1	0		0	100 i
*> 2.0.0.0	0.0.0.0	0		32768	i
*>i3.0.0.0	10.1.234.3	0	100	0	i
*>i4.0.0.0	10.1.234.4	0	100	0	i
*> 100.1.1.0/24	10.1.12.1	0		0	100 i
*> 100.2.2.0/24	10.1.12.1	0		0	100 i

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
* i1.1.0.0/24       10.1.12.1                0    100      0 100 i
* i1.1.1.0/24       10.1.12.1                0    100      0 100 i
* i1.1.2.0/24       10.1.12.1                0    100      0 100 i
* i1.1.3.0/24       10.1.12.1                0    100      0 100 i
*>i2.0.0.0          10.1.234.2               0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          10.1.234.4               0    100      0 i
* i100.1.1.0/24     10.1.12.1                0    100      0 100 i
* i100.2.2.0/24     10.1.12.1                0    100      0 100 i
```

On R4:

```
R4#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
* i1.1.0.0/24       10.1.12.1                0    100      0 100 i
* i1.1.1.0/24       10.1.12.1                0    100      0 100 i
* i1.1.2.0/24       10.1.12.1                0    100      0 100 i
* i1.1.3.0/24       10.1.12.1                0    100      0 100 i
*>i2.0.0.0          10.1.234.2               0    100      0 i
*>i3.0.0.0          10.1.234.3               0    100      0 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
* i100.1.1.0/24     10.1.12.1                0    100      0 100 i
* i100.2.2.0/24     10.1.12.1                0    100      0 100 i
```

Note: R3 and R4 do not have NLRI to the next-hop IP address of 10.1.12.1, therefore, they will not have reachability to these addresses.

Task 7

Configure R2 to change the next hop IP address for all the networks advertised by R1 to the IP address of its F0/1 interface. You should use a template so the future policies can be installed once in that template and have it effect R3 and R4. **Do not** use peer-groups to accomplish this task.

On R2:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config)#router bgp 200
R2(config-router)# template peer-policy TST
R2(config-router-ptmp)# next-hop-self
R2(config-router-ptmp)#exit-peer-policy

R2(config-router)#neighbor 10.1.234.3 inherit peer-policy TST
R2(config-router)#neighbor 10.1.234.4 inherit peer-policy TST
```

To verify the configuration:

On R3:

```
R3#Sh ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.1.0.0/24       10.1.234.2               0    100      0 100 i
*>i1.1.1.0/24       10.1.234.2               0    100      0 100 i
*>i1.1.2.0/24       10.1.234.2               0    100      0 100 i
*>i1.1.3.0/24       10.1.234.2               0    100      0 100 i
*>i2.0.0.0          10.1.234.2               0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          10.1.234.4               0    100      0 i
*>i100.1.1.0/24     10.1.234.2               0    100      0 100 i
*>i100.2.2.0/24     10.1.234.2               0    100      0 100 i
```

On R4:

```
R4#Sh ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.1.0.0/24       10.1.234.2               0    100      0 100 i
*>i1.1.1.0/24       10.1.234.2               0    100      0 100 i
*>i1.1.2.0/24       10.1.234.2               0    100      0 100 i
*>i1.1.3.0/24       10.1.234.2               0    100      0 100 i
*>i2.0.0.0          10.1.234.2               0    100      0 i
*>i3.0.0.0          10.1.234.3               0    100      0 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*>i100.1.1.0/24     10.1.234.2               0    100      0 100 i
*>i100.2.2.0/24     10.1.234.2               0    100      0 100 i
```

Peer-policy templates are used to build a template of policy information that can be inherited by a given neighbor/s. The peer-policy template cannot be inherited by a peer-session template or a peer-group.

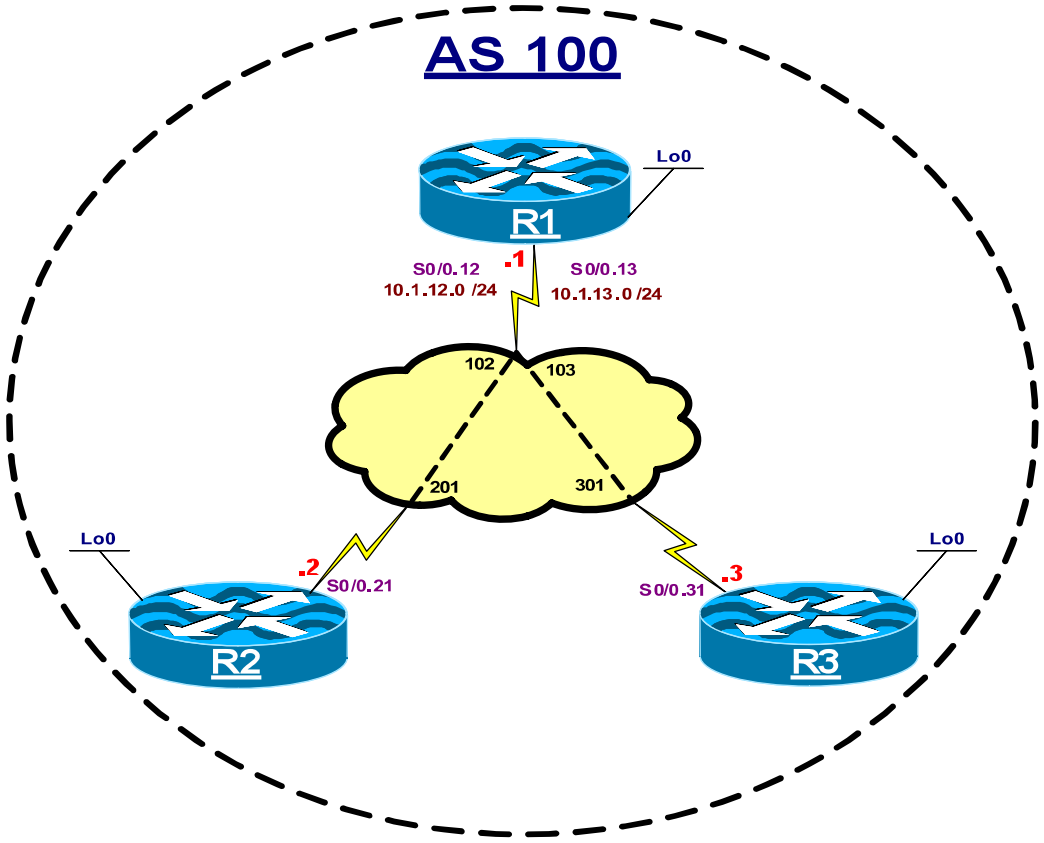
©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 8

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 2
Route Reflectors



Lab Setup:

- Configure R1 to have two point-to-point frame-relay connections, one connecting R1 to R2, and the other connecting R1 to R3.
- R2 and R3 should each be configured with a frame-relay point-to-point connection to R1
- Use the following IP address chart for IP address assignment.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP Addressing:

Router	Interface	IP address
R1	Lo0 The frame-relay connection to R2 The frame-relay connection to R3	1.1.1.1 /8 10.1.12.1 /24 10.1.13.1 /24
R2	**Lo0** **The frame-relay connection to R1**	**2.2.2.2 /8** **10.1.12.2 /24**
R3	Lo0 The frame-relay connection to R1	3.3.3.3 /8 10.1.13.3 /24

Task 1

Configure BGP AS 100 on all routers and ensure that the routers can successfully establish an IBGP peer session with each other. These routers should only advertise their loopback0 interface in BGP. To provide NLRI, the links between the routers should be advertised in RIPv2.

On All Routers

```
Rx(config-router)#router rip
Rx(config-router)#No au
Rx(config-router)#ver 2
Rx(config-router)#netw 10.0.0.0
```

On R1:

```
R1(config)#router bgp 100
R1(config-router)#netw 1.0.0.0
R1(config-router)#neighbor 10.1.12.2 remote-as 100
R1(config-router)#neighbor 10.1.13.3 remote-as 100
R1(config-router)#No auto
```

On R2:

```
R2(config)#router bgp 100
R2(config-router)#netw 2.0.0.0
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#neighbor 10.1.13.3 remote-as 100
R2(config-router)#No auto
```

On R3:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config)#router bgp 100
R3(config-router)#netw 3.0.0.0
R3(config-router)#neighbor 10.1.13.1 remote-as 100
R3(config-router)#neighbor 10.1.12.2 remote-as 100
R3(config-router)#No auto
```

To verify the configuration:

On R1:

```
R1#show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*>i2.0.0.0          10.1.12.2                0    100      0 i
*>i3.0.0.0          10.1.13.3                0    100      0 i

R1#show ip route bgp | inc B

B    2.0.0.0/8 [200/0] via 10.1.12.2, 00:01:39
B    3.0.0.0/8 [200/0] via 10.1.13.3, 00:01:45
```

On R2:

```
R2#show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.12.1                0    100      0 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          10.1.13.3                0    100      0 i

R2#show ip route bgp | inc B

B    1.0.0.0/8 [200/0] via 10.1.12.1, 00:02:21
B    3.0.0.0/8 [200/0] via 10.1.13.3, 00:02:42
```

On R3:

```
R3#show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.13.1                0    100      0 i
*>i2.0.0.0          10.1.12.2                0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3#show ip route bgp | inc B

B    1.0.0.0/8 [200/0] via 10.1.13.1, 00:03:26
B    2.0.0.0/8 [200/0] via 10.1.12.2, 00:03:40
```

Task 2

You received an e-mail from the management stating that within the next 12 months 20 additional routers will be added to this AS. In order to minimize the number of peer sessions within this AS, you decided to implement route reflectors. Configure R1 as a route reflector for this AS.

On R1:

```
R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.12.2 route-reflector-client
R1(config-router)#neighbor 10.1.13.3 route-reflector-client
```

On R2:

```
R2(config)#router bgp 100
R2(config-router)#No neighbor 10.1.13.3 remote-as 100
```

On R3:

```
R3(config)#router bgp 100
R3(config-router)#No neighbor 10.1.12.2 remote-as 100
```

In order for all IBGP speakers in an AS to exchange routes with one another, the IBGP speakers must be fully meshed (every router must establish a peer session to every other router). This is because of a rule called IBGP split-horizon, which states that when a BGP speaker receives and update from an IBGP peer, it will never advertise that update to another IBGP peer. Route-reflectors can be configured to reduce the number of peer sessions that must be established between the routers within a given AS. If a route-reflector is used, not all IBGP speakers need be fully meshed. In this model, the router that is configured to be the route-reflector must have a peer session established to every client, the clients must establish a peer session with the route reflector. The route reflector will reflect routes learned from one client to the other client/s.

To verify the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*>i2.0.0.0          10.1.12.2                0    100      0 i
*>i3.0.0.0          10.1.13.3                0    100      0 i
```

Note: The output of the following show commands reveal the fact that prefixes 2.0.0.0 and 3.0.0.0/8 was received from a route-reflector client:

```
R1#Show ip bgp 2.0.0.0/8

BGP routing table entry for 2.0.0.0/8, version 6
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to update-groups:
        2
  Local, (Received from a RR-client)
    10.1.12.2 from 10.1.12.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, valid, internal, best

R1#Show ip bgp 3.0.0.0/8

BGP routing table entry for 3.0.0.0/8, version 10
Paths: (1 available, best #1, table Default-IP-Routing-Table)
Flag: 0x820
  Advertised to update-groups:
        2
  Local, (Received from a RR-client)
    10.1.13.3 from 10.1.13.3 (3.3.3.3)
      Origin IGP, metric 0, localpref 100, valid, internal, best
```

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.12.1                0    100      0 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          10.1.13.3                0    100      0 i

R2#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/113/116 ms
```

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.13.1                0    100      0 i
*>i2.0.0.0          10.1.12.2                0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i

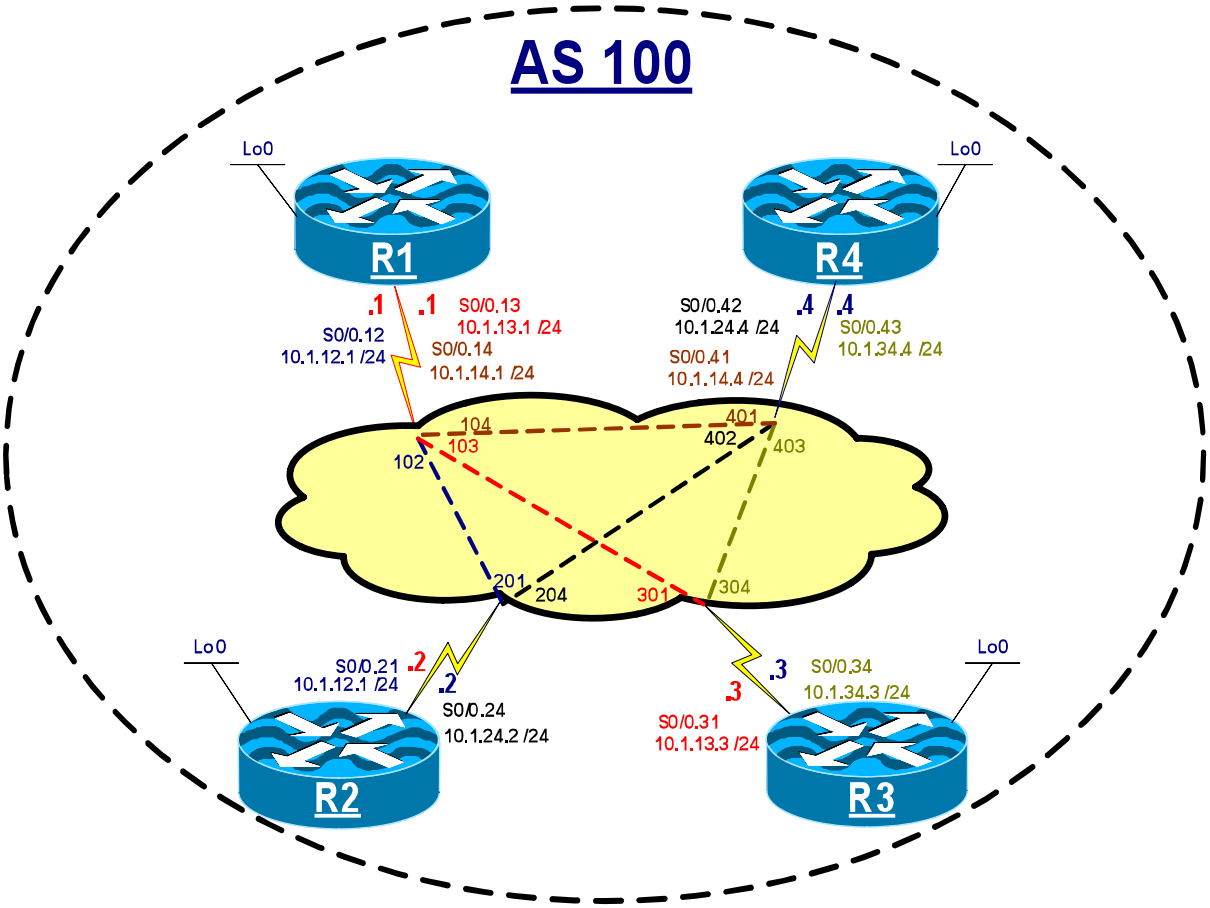
R3#Ping 2.2.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/112/116 ms
```

Task 3

After implementing the route reflector, you realized that if the route reflector is down, the entire network is dysfunctional; therefore, you decided to add R4 as the second route reflector for redundancy. Ensure that these routers can reach the advertised networks and the redundancy is operational.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh



DO NOT erase the existing configuration, the following configuration is added to the existing configuration.

Lab Setup:

Add the following configuration to the existing configuration:

- Configure R1 with an additional point-to-point frame-relay connection to R4, using the IP addressing and the DLCI information provided below.
- R2 and R3 should each be configured with an additional point-to-point frame-relay connection to R4. Use the IP addressing and the DLCI information provided below for these connections.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

- R4 should be configured with three point-to-point frame-relay connections, one to R1, the second one to R2 and the third one to R3. Use the following IP addressing and DLCI information for these connections.

IP Addressing:

Router	Interface	IP address	DLCI
R1	S0/0.14	10.1.14.1 /24	104
R2	**S0/0.24**	**10.1.24.2 /24**	**204**
R3	S0/0.34	10.1.34.3 /24	304
R4	**Lo0** **S0/0.42** **S0/0.43** **S0/0.41**	**4.4.4.4 /8** **10.1.24.4 /24** **10.1.34.4 /24** **10.1.14.4 /24**	 **402** **403** **401**

R4 is the secondary route-reflector. R4 should be configured as follows:

- **R4 should have a peer session with R1 – the route-reflectors should have full mesh peer sessions between them.**
- **R4 must have a peer session with R2 and R3.**
- **R4 must configure R2 and R3 as route-reflector clients.**

On R4:

```
R4(config)#router bgp 100
R4(config)#network 4.0.0.0
R4(config-router)#neighbor 10.1.14.1 remote-as 100
R4(config-router)#neighbor 10.1.24.2 remote-as 100
R4(config-router)#neighbor 10.1.34.3 remote-as 100
R4(config-router)#neighbor 10.1.24.2 route-reflector client
R4(config-router)#neighbor 10.1.34.3 route-reflector client

R4(config)#router rip
R4(config-router)#No au
R4(config-router)#ver 2
R4(config-router)#network 10.0.0.0
```

On R2:

```
R2(config)#router bgp 100
R2(config-router)#neighbor 10.1.24.4 remote-as 100
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3:

```
R3(config)#router bgp 100
R3(config-router)#neighbor 10.1.34.4 remote-as 100
```

On R1:

```
R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.14.4 remote-as 100
```

Having a single RR can introduce a single point of failure, it is best to have multiple RRs incase the RR fails, this redundancy is critical when there are many RR clients.

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
* i2.0.0.0          10.1.24.2                0    100      0 i
*>i                 10.1.12.2                0    100      0 i
* i3.0.0.0          10.1.34.3                0    100      0 i
*>i                 10.1.13.3                0    100      0 i
*>i4.0.0.0          10.1.14.4                0    100      0 i
```

On R2:

```
R2#Show ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
* i1.0.0.0          10.1.14.1                0    100      0 i
*>i                 10.1.12.1                0    100      0 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
* i3.0.0.0          10.1.34.3                0    100      0 i
*>i                 10.1.13.3                0    100      0 i
* i4.0.0.0          10.1.14.4                0    100      0 i
*>i                 10.1.24.4                0    100      0 i
```

On R3:

```
R3#Show ip bgp | b Network
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
   Network          Next Hop            Metric LocPrf Weight Path
* i1.0.0.0         10.1.14.1                0    100      0 i
*>i                10.1.13.1                0    100      0 i
* i2.0.0.0         10.1.24.2                0    100      0 i
*>i                10.1.12.2                0    100      0 i
*> 3.0.0.0         0.0.0.0                  0         32768 i
* i4.0.0.0         10.1.14.4                0    100      0 i
*>i                10.1.34.4                0    100      0 i
```

On R4:

```
R4#Show ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0         10.1.14.1                0    100      0 i
* i2.0.0.0         10.1.12.2                0    100      0 i
*>i                10.1.24.2                0    100      0 i
* i3.0.0.0         10.1.13.3                0    100      0 i
*>i                10.1.34.3                0    100      0 i
*> 4.0.0.0         0.0.0.0                  0         32768 i
```

To test the configuration:

On R1:

```
R1(config)#int s0/0
R1(config-if)#shut
```

On R2:

```
R2#Show ip bgp | b Network
```

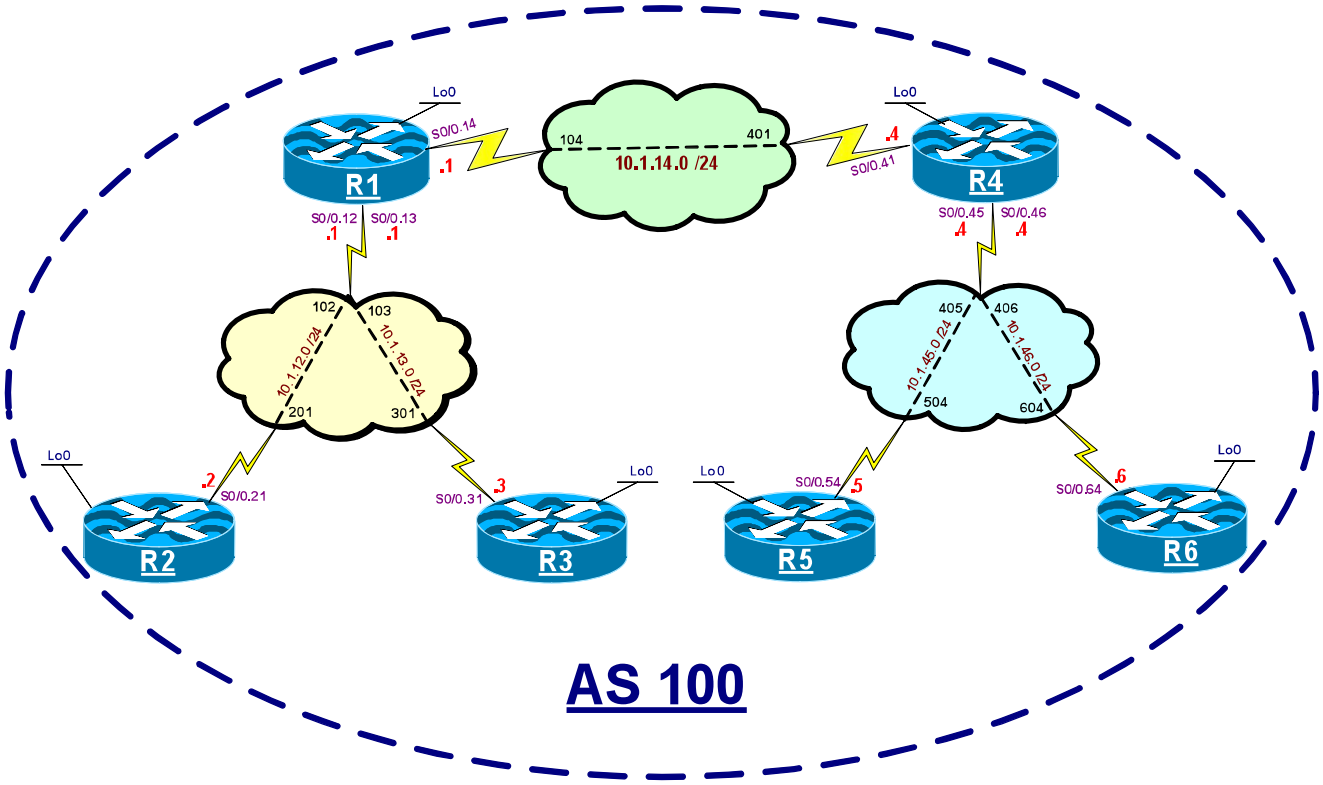
```
   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0         0.0.0.0                  0         32768 i
*>i3.0.0.0         10.1.34.3                0    100      0 i
*>i4.0.0.0         10.1.24.4                0    100      0 i

R2#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/114/116 ms
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Erase the startup configuration and reload the routers. Reconfigure the routers based on the following IP addressing and topology: You can copy and paste the configuration from the initial configuration called “Lab2-T4.txt”.



IP Addressing:

Router	Interface	IP address
R1	Lo0 The frame-relay connection to R4 The frame-relay connection to R2 The frame-relay connection to R3	1.1.1.1 /8 10.1.14.1 /24 10.1.12.1 /24 10.1.13.1 /24
R2	**Lo0** **The frame-relay connection to R1**	**2.2.2.2 /8** **10.1.12.2 /24**
R3	Lo0 The frame-relay connection to R1	3.3.3.3 /8 10.1.13.3 /24
R4	**Lo0** **The frame-relay connection to R1** **The frame-relay connection to R5** **The frame-relay connection to R6**	**4.4.4.4 /8** **10.1.14.4 /24** **10.1.45.4 /24** **10.1.46.4 /24**
R5	Lo0 The frame-relay connection to R1	5.5.5.5 /8 10.1.45.5 /24

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

R6	**Lo0** **The frame-relay connection to R1**	**6.6.6.6 /8** **10.1.46.6 /24**

Lab Setup:

- Configure R1 with three point-to-point frame-relay connections; these point-to-point connections should connect R1 to R2, R3 and R4.
- Configure R2 and R3 with a single point-to-point connection to R1.
- R4 should be configured with three point-to-point frame-relay connections; these point-to-point connections should connect R4 to R5, R6 and R1.
- Configure R5 and R6 with a single point-to-point connection to R4.

Task 4

Configure BGP on R1 – R6 based on the following policy:

- R1 should be configured as the route reflector for routers R2 and R3, whereas R4 should be configured to be the route-reflector for routers R5 and R6.
- R1 and R4 should be configured to have an IBGP peer session between them; these two routers should be configured in BGP AS 100.
- NLRI for the links should be provided through RIPv2.

On All Routers

```
Rx(config)#router rip
Rx(config-router)#No au
Rx(config-router)#ver 2
Rx(config-router)#netw 10.0.0.0
```

On R1:

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

R1(config-router)#netw 1.0.0.0
R1(config-router)#neighbor 10.1.12.2 remote-as 100
R1(config-router)#neighbor 10.1.13.3 remote-as 100
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config-router)#neighbor 10.1.14.4 remote-as 100
R1(config-router)#neighbor 10.1.12.2 route-reflector-client
R1(config-router)#neighbor 10.1.13.3 route-reflector-client
```

On R2:

```
R2(config)#router bgp 100
R2(config-router)#No au
R2(config-router)#network 2.0.0.0
R2(config-router)#neighbor 10.1.12.1 remote-as 100
```

On R3:

```
R3(config)#router bgp 100
R3(config-router)#No au
R3(config-router)#network 3.0.0.0
R3(config-router)#neighbor 10.1.13.1 remote-as 100
```

To verify the configuration:

On R3:

```
R3#Show ip bgp

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.13.1                0    100      0 i
*>i2.0.0.0          10.1.12.2                0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
```

On R4:

```
R4(config)#router bgp 100
R4(config-router)#No au
R4(config-router)#network 4.0.0.0
R4(config-router)#neighbor 10.1.14.1 remote-as 100
R4(config-router)#neighbor 10.1.45.5 remote-as 100
R4(config-router)#neighbor 10.1.46.6 remote-as 100
R4(config-router)#neighbor 10.1.45.5 route-reflector-client
R4(config-router)#neighbor 10.1.46.6 route-reflector-client
```

On R5:

```
R5(config)#router bgp 100
R5(config-router)#No au
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R5(config-router)#network 5.0.0.0
R5(config-router)#neighbor 10.1.45.4 remote-as 100
```

On R6:

```
R6(config)#router bgp 100
R6(config-router)#No au
R6(config-router)#network 6.0.0.0
R6(config-router)#neighbor 10.1.46.4 remote-as 100
```

To verify the configuration:

On R6:

```
R6#Show ip bgp b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.14.1                0    100      0 i
*>i2.0.0.0          10.1.12.2                0    100      0 i
*>i3.0.0.0          10.1.13.3                0    100      0 i
*>i4.0.0.0          10.1.46.4                0    100      0 i
*>i5.0.0.0          10.1.45.5                0    100      0 i
*> 6.0.0.0          0.0.0.0                  0         32768 i
```

On R3:

```
R3#Show ip route bgp | inc B

B    1.0.0.0/8 [200/0] via 10.1.13.1, 00:04:17
B    2.0.0.0/8 [200/0] via 10.1.12.2, 00:04:17
B    4.0.0.0/8 [200/0] via 10.1.14.4, 00:02:14
B    5.0.0.0/8 [200/0] via 10.1.45.5, 00:01:53
B    6.0.0.0/8 [200/0] via 10.1.46.6, 00:01:16
```

Note: AS 100 has two route reflectors, each route-reflector has its own clients, when a given RR receives an update from one of its clients. It advertises that prefix to the other RR/s, the other RR/s in turn advertise that prefix to their clients.

There are some additional optional non-transitive attributes that can be used when RRs are configured and they are originator-id, cluster-id and cluster-list.

- **Originator-id: This attribute is created by the RR; this is the router-id of the router that originated the prefix. It is created to avoid routing loops, a RR will not advertise a route back to the originator of the prefix and if the originator of a prefix receives an update with its own router-id, it will ignore that prefix.**

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

- **Cluster and cluster-id: A RR/s and its clients are collectively known as a cluster, each cluster must be uniquely identified, and the cluster-id is typically the router-id of the RR unless specifically configured.**

- **Cluster-list: This attribute is analogous to AS-path attribute, and it keeps track of the cluster-ids in the same way that the AS-path attribute keeps track of the AS numbers. When the RR advertises a prefix to a non-client it appends the cluster-id to that prefix's cluster-list, if a RR receives an update and sees its own cluster-id in the cluster-list, it will ignore that update.**

To see the attributes:

On R1:

```
R1#Show ip bgp 6.0.0.0

BGP routing table entry for 6.0.0.0/8, version 7
Paths: (1 available, best #1, table Default-IP-Routing-Table)
Flag: 0x820
  Advertised to update-groups:
        2
  Local
    10.1.46.6 (metric 1) from 10.1.14.4 (4.4.4.4)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Originator: 6.6.6.6, Cluster list: 4.4.4.4
```

Note: Prefix 6.0.0.0 was originated from 6.0.0.0 the originator of the prefix and it came from 4.4.4.4 (the cluster-list).

On R3:

```
R3#Show ip bgp 6.0.0.0

BGP routing table entry for 6.0.0.0/8, version 7
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  Local
    10.1.46.6 (metric 2) from 10.1.13.1 (1.1.1.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Originator: 6.6.6.6, Cluster list: 1.1.1.1, 4.4.4.4
```

Note: This prefix has gone through cluster-ids of 4.4.4.4 first and then it traversed through cluster-id 1.1.1.1 before it was received by the local router.

Note: The originator-id is the router-id of the router that originated that prefix, the output of the following "show" command reveals the router-id of the router that originated the route.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R6#Sh ip bgp

BGP table version is 7, local router ID is 6.6.6.6
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

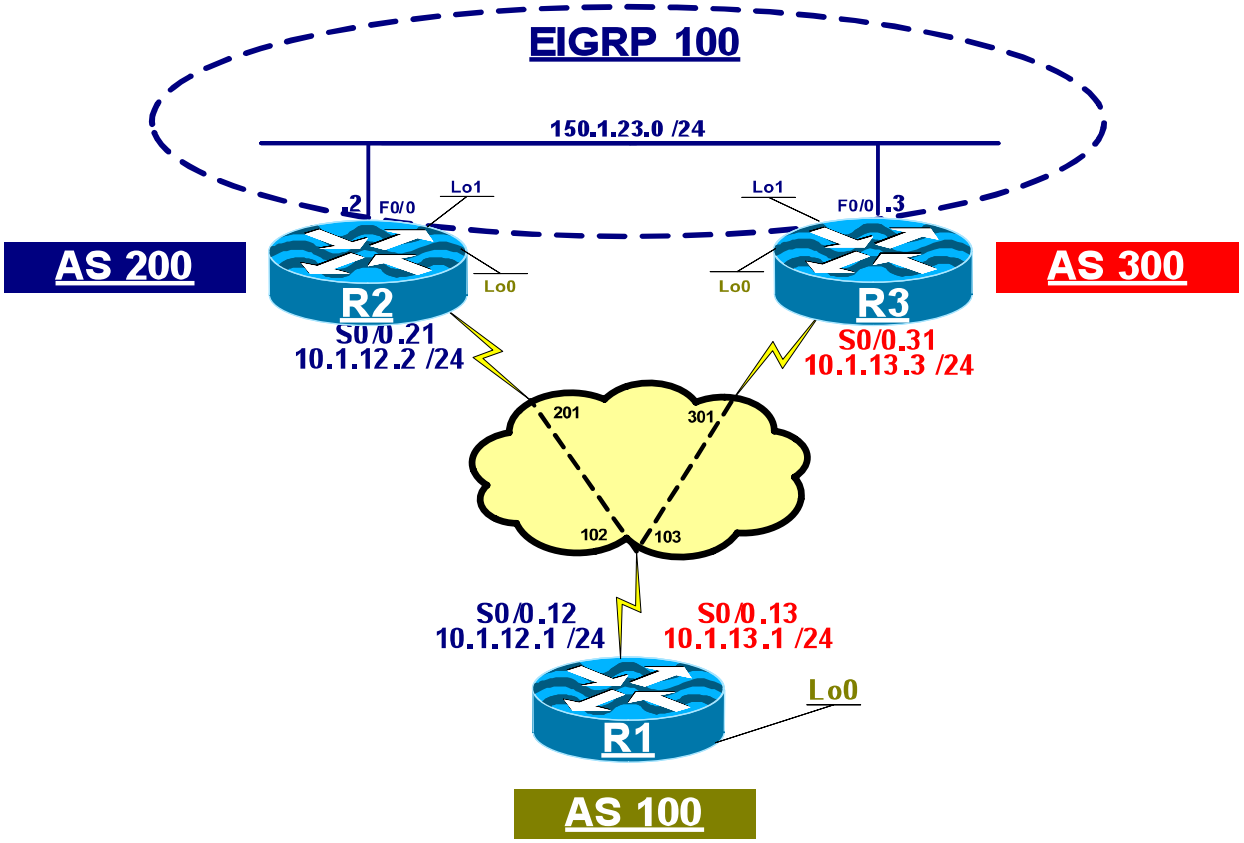
   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.14.1                0    100      0 i
*>i2.0.0.0          10.1.12.2                0    100      0 i
*>i3.0.0.0          10.1.13.3                0    100      0 i
*>i4.0.0.0          10.1.46.4                0    100      0 i
*>i5.0.0.0          10.1.45.5                0    100      0 i
*> 6.0.0.0          0.0.0.0                  0         32768 i
```

Task 5

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 3
Conditional Advertisement & BGP Backdoor



Lab Setup:

- Configure R1 to have two point-to-point frame-relay connections, one connecting R1 to R2, and the other connecting R1 to R3.
- R2 and R3 should be configured with a single frame-relay point-to-point connection to R1.
- Configure R2 and R3's F0/0 interface to be in VLAN 23.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP addressing:

Router	Interface	IP Address	AS number
R1	Lo0 S0/0.12 S0/0.13	1.1.1.1 /8 10.1.12.1 /24 10.1.13.1 /24	100
R2	**Lo0** **Lo1** **S0/0.21** **F0/0**	**2.2.2.2 /8** **150.1.2.2 /24** **10.1.12.2 /24** **150.1.23.2 /24**	**200**
R3	Lo0 Lo1 S0/0.31 F0/0	3.3.3.3 /8 150.1.3.3 /24 10.1.13.3 /24 150.1.23.3 /24	300

Task 1

Configure EBGP peer sessions as follows:

- R1 to have an EBGP peer session to R2 and R3
- R2 and R3 to have an EBGP peer session to R1

On R1:

```
R1(config)#router bgp 100
R1(config-router)#No auto
R1(config-router)#neighbor 10.1.12.2 remote-as 200
R1(config-router)#neighbor 10.1.13.3 remote-as 300
```

On R2:

```
R2(config)#router bgp 200
R2(config-router)#No au
R2(config-router)#neighbor 10.1.12.1 remote-as 100
```

On R3:

```
R3(config-if)#router bgp 300
R3(config-router)#No au
R3(config-router)#neighbor 10.1.13.1 remote-as 100
```

To verify the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1:

```
R1#Show ip bgp summary | b Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.12.2       4   200       3       3        1    0    0 00:00:39        0
10.1.13.3       4   300       3       3        1    0    0 00:00:32        0
```

Task 2

Configure R1, R2 and R3 to advertise their loopback0 interface in BGP.

On R1:

```
R1(config)#router bgp 100
R1(config-router)#network 1.0.0.0
```

On R2:

```
R2(config)#router bgp 200
R2(config-router)#network 2.0.0.0
```

On R3:

```
R3(config)#router bgp 300
R3(config-router)#network 3.0.0.0
```

To verify the configuration:

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.13.1                0             0 100 i
*> 2.0.0.0          10.1.13.1                               0 100 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 3

Configure RIPv2 and EIGRP 100 on the routers as follows:

- Configure RIPv2 on all routers to advertise network 10.0.0.0; disable auto summarization.
- R2 and R3 should advertise network 150.1.0.0 that will advertise their loopback1 and F0/0 interface in EIGRP AS#100. Disable auto summarization.

On R2: and R3

```
Rx(config)#router eigrp 100
Rx(config-router)#No au
Rx(config-router)#netw 150.1.0.0
```

On All Routers:

```
Rx(config)#router rip
Rx(config-router)#ver 2
Rx(config-router)#No au
Rx(config-router)#network 10.0.0.0
```

Task 4

Since network 150.1.0.0 is **not** advertised in BGP, if the link between R2 and R3 (the F0/0 interface) goes down, the loopback1 network of these two routers will not have reachability to each other. Even though there is a redundant link between these two routers through BGP, therefore, the administrator of R2 and R3 decided to advertise their loopback 1 interface in BGP for redundancy, configure these routers to accommodate this decision.

On R2:

```
R2(config-router)#router bgp 200
R2(config-router)#network 150.1.2.0 mask 255.255.255.0
```

On R3:

```
R3(config)#router bgp 300
R3(config-router)#network 150.1.3.0 mask 255.255.255.0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the configuration:

On R2:

```
R2#Show ip route bgp | inc B

B     1.0.0.0/8 [20/0] via 10.1.12.1, 00:07:33
B     3.0.0.0/8 [20/0] via 10.1.12.1, 00:07:03
B        150.1.3.0 [20/0] via 10.1.12.1, 00:00:35
```

On R3:

```
R3#Show ip route bgp | inc B

B     1.0.0.0/8 [20/0] via 10.1.13.1, 00:08:25
B     2.0.0.0/8 [20/0] via 10.1.13.1, 00:07:54
B        150.1.2.0 [20/0] via 10.1.13.1, 00:00:55
```

Task 5

After implementing the previous task, the administrators realized that the traffic between networks 150.1.2.0 /24 and 150.1.3.0 /24 is taking a sub-optimal path and it is not using the direct path between routers R2 and R3.

Implement a BGP solution to fix this problem; you should **not** use the distance or any global config mode command to accomplish this task.

Note: The BGP "backdoor" option is assigned to the network that is advertised to you, therefore, you should reference the network that is advertised to you and not the network that your local router is advertising:

On R2:

```
R2(config)#router bgp 200
R2(config-router)#network 150.1.3.0 mask 255.255.255.0 backdoor
```

On R3:

```
R3(config)#router bgp 300
R3(config-router)#network 150.1.2.0 mask 255.255.255.0 backdoor
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the configuration:

On R3:

```
R3#Show ip route eigrp | inc D

D       150.1.2.0 [90/156160] via 150.1.23.2, 00:00:20, FastEthernet0/0
```

On R2:

```
R2#Show ip route eigrp | inc D

D       150.1.3.0 [90/156160] via 150.1.23.3, 00:01:21, FastEthernet0/0
```

Note: R2 and R3 were receiving routing information for networks 150.1.2.0 /24 and 150.1.3.0 /24 from two different sources, BGP and EIGRP.

R2 and R3 were using the routing information from BGP because it had a lower administrative distance (20 versus 90).

The "network backdoor" command is a BGP solution to this problem; the BGP "backdoor" option assigns an administrative distance of 200 to networks 150.1.2.0 /24 and 150.1.3.0 /24, therefore, making the EIGRP more attractive.

Enter the following commands to actually see the changed administrative distance:

On R2: and R3

```
Rx(config)#No router eigrp 100
```

On R2:

```
R2#Sh ip route bgp | inc B

B    1.0.0.0/8 [20/0] via 10.1.12.1, 00:14:47
B    3.0.0.0/8 [20/0] via 10.1.12.1, 00:14:16
B       150.1.3.0 [200/0] via 10.1.12.1, 00:00:26
```

On R3:

```
R3#Show ip route bgp | inc B

B    1.0.0.0/8 [20/0] via 10.1.13.1, 00:15:36
B    2.0.0.0/8 [20/0] via 10.1.13.1, 00:15:05
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
B        150.1.2.0 [200/0] via 10.1.13.1, 00:01:15
```

Task 6

Remove the IP address from the F0/0 interface of R2 and R3 and ensure that the F0/0 interface of both routers is in administratively down state. You should also remove the loopback1 interface from these two routers.

On R2: and R3

```
Rx(config)#interface f0/0
Rx(config-if)#No ip address
Rx(config-if)#shutdown

Rx(config)#No int lo1
```

Task 7

Configure R1 as follows:

- If network 2.0.0.0 /8 is up and it is advertised to R1, R1 should **not** advertise its network 1.0.0.0 /8 to R3.
- R1 should advertise network 1.0.0.0 /8 to R3 **only** if network 2.0.0.0 /8 is down.

Before configuring this task you should verify the current BGP table of these routers:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.0.0.0          10.1.13.3                0             0 300 i


R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.12.1                0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
*> 3.0.0.0          10.1.12.1                              0 100 300 i

R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.13.1                0             0 100 i
*> 2.0.0.0          10.1.13.1                              0 100 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
```

To implement conditional advertisement of selected prefixes, the following MUST be used:

- **Advertise-map**
- **Non-exist-map**
- **Exist-map**

This situation calls for the use of the "advertise-map" and "non-exist-map" as follows:

On R1:

Step #1 – Identify the prefixes using two access-list:

```
R1(config)#access-list 1 permit 1.0.0.0 0.255.255.255

R1(config)#access-list 2 permit 2.0.0.0 0.255.255.255
```

Step #2 – Configure two route-maps, one to reference access-list 1 and the second one to reference access-list 2. To prevent confusion you should select meaningful names for the route-maps.

```
R1(config)#route-map ADV permit 10
R1(config-route-map)#match ip addr 1
R1(config-route-map)#exit

R1(config)#route-map NotThere permit 10
R1(config-route-map)#match ip addr 2
R1(config-route-map)#exit
```

Final Step – the route-maps are referenced by the "advertise-map" and "non-exist-map" options:

```
R1(config)#router bgp 100

R1(config-router)#neighbor 10.1.13.3 advertise-map ADV non-exist-map NotThere
```

The neighbor command has the following route-maps:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

- **The advertise-map – Specifies the name of the route-map that will be advertised if the condition of the non-exist-map is met.**

- **Non-exist-map – specifies the name of the route-map that will be compared to the advertise-map. If the condition is met and no match occurs, the route will be advertised. If a match occurs, then the condition is not met, and the route is withdrawn.**

Note: If network 2.0.0.0 is up, then network 1.0.0.0 should not be advertised to R3, since all the networks are up and advertised, R1 should withdraw its network (1.0.0.0 /8):

On R1:

Note: All the networks are up; therefore, network 1.0.0.0 /8 should not be advertised to R3.

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.0.0.0          10.1.13.3                0             0 300 i
```

The following show command reveals that R1 does not advertise its network (1.0.0.0 /8) to R3:

```
R1#Show ip bgp neighbors 10.1.13.3 advertised-routes | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.12.2                0             0 200 i
```

To verify this configuration:

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.13.1                              0 100 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
```

To test the condition:

On R2:

```
R2(config)#int lo0
R2(config-if)#shut
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The output of the following "show" command reveals that network 2.0.0.0 is down; therefore, R1 should advertise its network (1.0.0.0 /8) to R3. It may take few seconds for this policy to be implemented:

On R1:

```
R1#Show ip bgp neighbors 10.1.13.3 advertised-routes | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
```

To see the test on R3:

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.13.1                0             0 100 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
```

Task 8

Remove the configuration commands entered in Task 7 before you proceed to the next task. Ensure that the routers have the advertised networks in their BGP table.

On R1:

```
R1(config)#No access-list 1
R1(config)#No access-list 2

R1(config)#No route-map ADV

R1(config)#No route-map NotThere

R1(config)#router bgp 100
R1(config-router)#No neighbor 10.1.13.3 advertise-map ADV non-exist-map NotThere

R1#Clear ip bgp *
```

On R2:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config)#int lo0
R2(config-if)#No shut
```

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.0.0.0          10.1.13.3                0             0 300 i
```

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.12.1                0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.0.0.0          10.1.12.1                              0 100 300 i
```

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.13.1                0             0 100 i
*> 2.0.0.0          10.1.13.1                              0 100 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
```

Task 9

R1 should be configured according to the following policy:

1. If both networks (1.0.0.0 /8 and 2.0.0.0 /8) are up, then both networks should be advertised to R3.
2. If network 1.0.0.0 /8 is down, R1 should **not** advertise network 2.0.0.0 /8 to R3.
3. If network 2.0.0.0 /8 is down, then R1 should only advertise network 1.0.0.0 /8 to R3.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1:

Step #1 - The following two access-lists identify the two networks (1.0.0.0 /8 and 2.0.0.0 /8):

```
R1(config)#access-list 1 permit 1.0.0.0 0.255.255.255

R1(config)#access-list 2 permit 2.0.0.0 0.255.255.255
```

Step #2 - The following route-maps are configured to reference the two access-lists from the previous step:

```
R1(config)#route-map ADV permit 10
R1(config-route-map)#match ip addr 2

R1(config)#route-map EXIST permit 10
R1(config-route-map)#match ip addr 1
```

Final Step – With the following configuration, we are instructing BGP for the conditions of the task's requirements:

```
R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.13.3 advertise-map ADV exist-map EXIST
```

To test the first condition

If both networks (1.0.0.0 /8 and 2.0.0.0 /8) are up, then both networks should be advertised to R3.

Note: Both prefixes are up:

On R1:

```
R1#Show ip bgp | b Network
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.0.0.0          10.1.13.3                0             0 300 i
```

Let's see the prefixes that R1 is advertising to R3:

On R1:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1#Show ip bgp neighbor 10.1.13.3 advertised-routes | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
```

As you can see both prefixes are being advertised to R3, let's check R3's BGP table:

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.13.1                0             0 100 i
*> 2.0.0.0          10.1.13.1                              0 100 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
```

To test the second condition

If network 1.0.0.0 /8 is down, R1 should not advertise network 2.0.0.0 /8 to R3.

Let's shutdown R1's Lo0 interface:

On R1:

```
R1(config)#int lo0
R1(config-if)#shut
```

To force the change much faster:

On R1:

```
R1#Clear ip bgp * out
```

Let's see the prefixes that R1 is advertising to R3:

```
R1#Show ip bgp neighbor 10.1.13.3 advertised-routes | b Network
R1#
```

Note: The output of the above show command reveals that R1 is not advertising any prefixes to R3. Let's check R3's BGP table to verify:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 3.0.0.0          0.0.0.0                  0         32768 i
```

To test the third condition

If network 2.0.0.0 /8 is down, then R1 should only advertise network 1.0.0.0 /8 to R3.

Let's "no shut" R1's Lo0 and shutdown R2's Lo0:

On R1:

```
R1(config)#int lo0
R1(config-if)#no shut
```

On R2:

```
R2(config)#int lo0
R2(config-if)#shut
```

To force the change much faster:

On R1:

```
R1#Clear ip bgp * out
```

Let's see which prefixes are advertised to R3 by R1:

```
R1#Show ip bgp neighbor 10.1.13.3 advertised-routes | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
```

To verify the configuration on R3

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.13.1                0             0 100 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

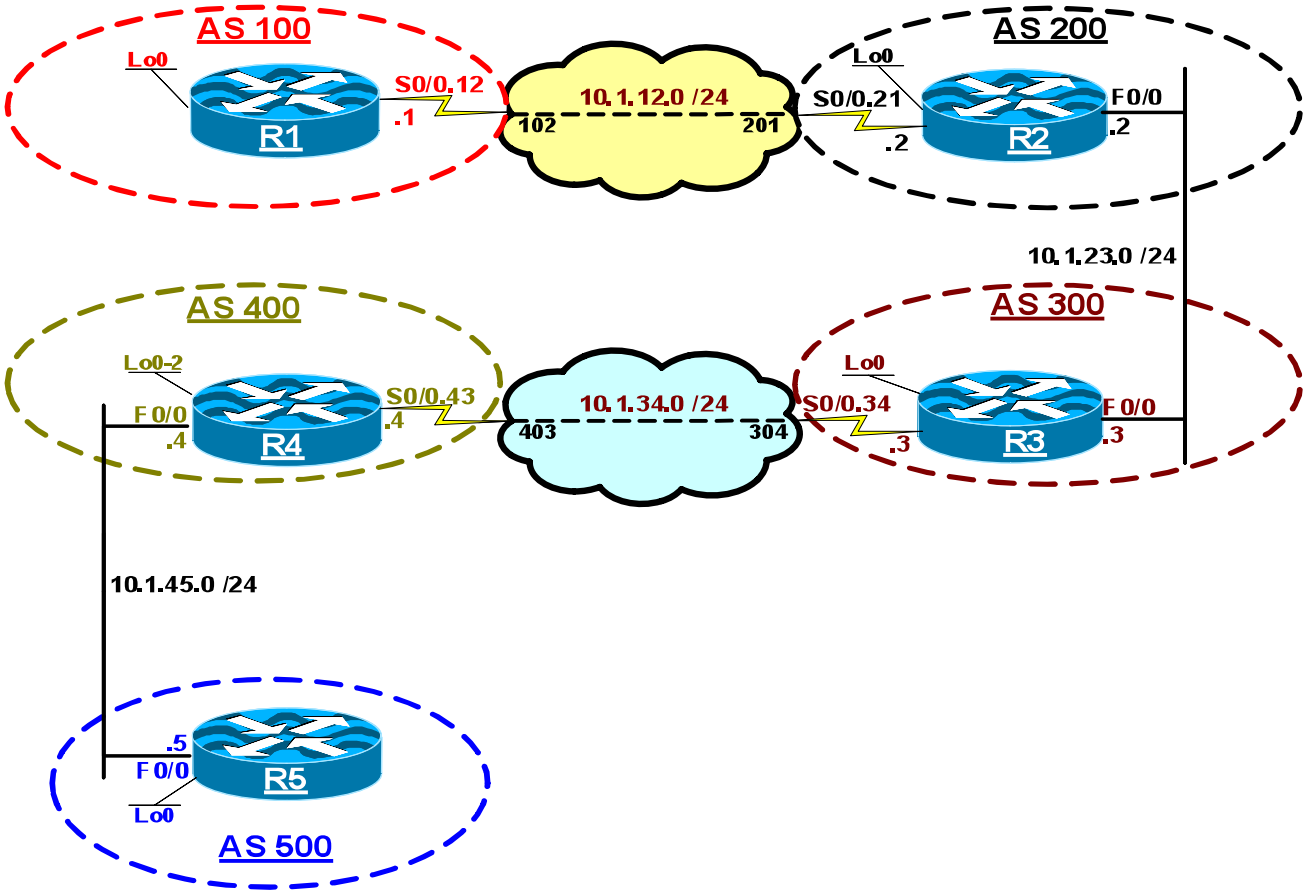
```
*> 3.0.0.0          0.0.0.0                  0         32768 i
```

Task 10

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 4
Route Dampening



Lab Setup:

- Configure the frame-relay connections between the routers in a point-to-point manner.
- Configure R2 and R3's F0/0 interface in VLAN 23.
- Configure R4 and R5's F0/0 interface in VLAN 45.
- The IP address assignment of the routers should be based on the following IP addressing chart:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP addressing:

Router	Interface	IP Address	AS number
R1	Lo0 S0/0.12	1.1.1.1 /8 10.1.12.1 /24	100
R2	Lo0 S0/0.21 F0/0	2.2.2.2 /8 10.1.12.2 /24 10.1.23.2 /24	200
R3	Lo0 S0/0.34 F0/0	3.3.3.3 /8 10.1.34.3 /24 10.1.23.3 /24	300
R4	Lo0 Lo1 Lo2 S0/0.43 F0/0	4.4.4.4 /8 40.4.4.4 /24 44.4.4.4 /24 10.1.34.4 /24 10.1.45.4 /24	400
R5	Lo0 F0/0	5.5.5.5 /8 10.1.45.5 /24	500

Task 1

Configure an EBGP peer session between R1 and R2 and only advertise their loopback0 interfaces in BGP. Ensure that these routers have NLRI to each other's advertised prefix. R1 should be in AS 100 and R2 should be in AS 200.

On R1:

```
R1(config)#router bgp 100
R1(config-router)#No au
R1(config-router)#network 1.0.0.0
R1(config-router)#neighbor 10.1.12.2 remote-as 200
```

On R2:

```
R2(config)#router bgp 200
R2(config-router)#No au
R2(config-router)#network 2.0.0.0
R2(config-router)#neighbor 10.1.12.1 remote-as 100
```

To verify the configuration:

On R1:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
```

Task 2

Configure route dampening on R1 using the default parameters.

On R1:

```
R1(config-router)#router bgp 100
R1(config-router)#bgp dampening
```

The parameters of BGP dampening are as follows:

- **Half time – Once a route has been assigned a penalty, the accumulated penalty is decreased every 5 seconds such that when the half-time period expires. The accumulated penalties are reduced by half. The default value of half time is 15 minutes and the range is 1 to 45 minutes.**
- **Reuse – If the penalties for a flapping route is decreased enough to fall below this value, the route is reusable. The default is 750 and the range is 1 to 20000.**
- **Suppress – Once the accumulated penalties reaches this value, the route is suppressed. The default value is 2000 and the range is 1 – 20000.**
- **Max-Suppress-Time – The maximum time in minutes that a route can be suppressed. The default value is 4 times the half-time value (60 minutes) and the range is 1 to 255.**

Therefore this configuration performs the following:

Half time = 15 minutes, reuse = 750, Suppress = 2000 and Max-Suppress-Time = 60.

To see the parameters for dampening:

On R1:

```
R1#Sh ip bgp dampening parameters

dampening 15 750 2000 60 (DEFAULT)
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Half-life time       : 15 mins        Decay Time          : 2320 secs
Max suppress penalty: 12000           Max suppress time: 60 mins
Suppress penalty     :  2000          Reuse penalty      : 750
```

If network 2.0.0.0 is shut down and then brought backup few times, the flap-statistics can be viewed in the output of the "show ip bgp dampening flap-statistics" command.

Task 3

Configure an EBGP peer session between R2 and R3, and advertise their loopback0 interface in BGP. Ensure that these routers have NLRI to each other's loopback interface/s. R3 should be configured in AS 300.

On R2:

```
R2(config)#router bgp 200
R2(config-router)#neighbor 10.1.23.3 remote-as 300
```

On R3:

```
R3(config)#router bgp 300
R3(config-router)#no au
R3(config-router)#network 3.0.0.0
R3(config-router)#neighbor 10.1.23.2 remote-as 200
```

To verify the configuration:

On R2:

```
R2#Show ip bgp | b Network
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.12.1                0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.0.0.0          10.1.23.3                0             0 300 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 4

Configure route dampening on R3 such that the half-life parameter is set to 30 minutes. This router should use the default parameters for suppress-limit, reuse, and maximum suppress time.

On R3:

```
R3(config-router)#bgp dampening 30 750 2000 60
```

To verify the configuration:

On R3:

```
R3#Sh ip bgp dampening parameters

dampening 30 750 2000 60
  Half-life time       : 30 mins       Decay Time        : 1045 secs
  Max suppress penalty:  3000          Max suppress time: 60 mins
  Suppress penalty     :  2000         Reuse penalty     : 750
```

Note: You may receive a "% dampening reconfiguration in progress for IPv4 unicast" message; you should wait few seconds and try again.

Task 5

Configure an EBGP session between R3 and R4. Advertise loopback0, loopback1 and loopback2 interfaces of R4 in BGP. R4 should be configured in AS 400.

On R4:

```
R4(config)#router bgp 400
R4(config-router)#netw 4.0.0.0
R4(config-router)#netw 40.4.4.0 mask 255.255.255.0
R4(config-router)#netw 44.4.4.0 mask 255.255.255.0
R4(config-router)#neighbor 10.1.34.3 remote-as 300
R4(config-router)#No au
```

On R3:

```
R3(config)#router bgp 300
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-router)#neighbor 10.1.34.4 remote-as 400
```

To verify the configuration:

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.23.2                              0 200 100 i
*> 2.0.0.0          10.1.23.2                0             0 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*> 4.0.0.0          10.1.34.4                0             0 400 i
*> 40.4.4.0/24      10.1.34.4                0             0 400 i
*> 44.4.4.0/24      10.1.34.4                0             0 400 i
```

On R4:

```
R4#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.34.3                              0 300 200 100 i
*> 2.0.0.0          10.1.34.3                              0 300 200 i
*> 3.0.0.0          10.1.34.3                0             0 300 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*> 40.4.4.0/24      0.0.0.0                  0         32768 i
*> 44.4.4.0/24      0.0.0.0                  0         32768 i
```

Task 6

Configure an EBGP peer session between R4 and R5. R5 should advertise its loopback 0 interface in BGP. Ensure that these routers have NLRI to each other's loopback interface. R5 should be configured in AS 500.

On R5:

```
R5(config)#router bgp 500
R5(config-router)#no au
R5(config-router)#neighbor 10.1.45.4 remote-as 400
R5(config-router)#netw 5.0.0.0
```

On R4:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config)#router bgp 400
R4(config-router)#neighbor 10.1.45.5 remote-as 500
```

To verify the configuration:

On R5:

```
R5#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.45.4                              0 400 300 200 100 i
*> 2.0.0.0          10.1.45.4                              0 400 300 200 i
*> 3.0.0.0          10.1.45.4                              0 400 300 i
*> 4.0.0.0          10.1.45.4                0             0 400 i
*> 5.0.0.0          0.0.0.0                  0         32768 i
*> 40.4.4.0/24      10.1.45.4                0             0 400 i
*> 44.4.4.0/24      10.1.45.4                0             0 400 i
```

Task 7

Configure route dampening On R4: as follows:

- Network 40.4.4.0 /24 should have the following dampening parameters applied: **max-suppress-time** of 90, **reuse** 800, **suppress** 2400 and a **half-time** of 20
- Network 44.4.4.0 /24 should have the following dampening parameters applied: **max-suppress-time** of 60, **reuse** 700, **suppress** 2000 and a **half-time** of 15

When route dampening is enabled, it affects all prefixes; but it can be configured to affect a specific route/s; this can be done through configuration of route-map/s. Route-map/s gives us this flexibility.

In this case, we have applied two different route-dampening parameters to different routes. The two networks are identified with access-lists. The route-map references the access-lists and sets the dampening parameters based on the networks.

On R4:

```
R4(config)#access-list 40 permit 40.4.4.0 0.0.0.255

R4(config)#access-list 44 permit 44.4.4.0 0.0.0.255
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config)#route-map TST permit 10
R4(config-route-map)#match ip addr 40
R4(config-route-map)#set dampening 20 800 2400 90

R4(config)#route-map TST permit 20
R4(config-route-map)#match ip addr 44
R4(config-route-map)#set dampening 15 700 2000 60

R4(config)#route-map TST permit 90

R4(config)#router bgp 400
R4(config-router)#bgp dampening route-map TST
```

To verify the configuration:

On R4:

```
R4#Sh ip bgp damp parameters

dampening 20 800 2400 90 (route-map TST 10)
  Half-life time       : 20 mins       Decay Time        : 3490 secs
  Max suppress penalty: 18075          Max suppress time: 90 mins
  Suppress penalty     :  2400         Reuse penalty     : 800

dampening 15 700 2000 60 (route-map TST 20)
  Half-life time       : 15 mins       Decay Time        : 2235 secs
  Max suppress penalty: 11200          Max suppress time: 60 mins
  Suppress penalty     :  2000         Reuse penalty     : 700
```

Task 8

Configure route dampening on R2 using the following policy:

- All the existing and future prefixes from AS 300 should have the following parameters applied:

  Max-Suppress-Time of 80
  Reuse 750
  Suppress 2200
  Half Time 30

On R2:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config)#ip as-path access-list 1 permit ^300$

R2(config)#route-map TST permit 10
R2(config-route-map)#match as-path 1
R2(config-route-map)#set dampening 30 750 2200 80
R2(config)#route-map TST permit 90

R2(config)#router bgp 200
R2(config-router)#bgp dampening route-map TST
```

The combination of "route-map" and the "as-path access-list" command can apply bgp dampening to an AS based on the AS number.

To verify the configuration:

On R2:

```
R2#Show ip bgp damp parameters

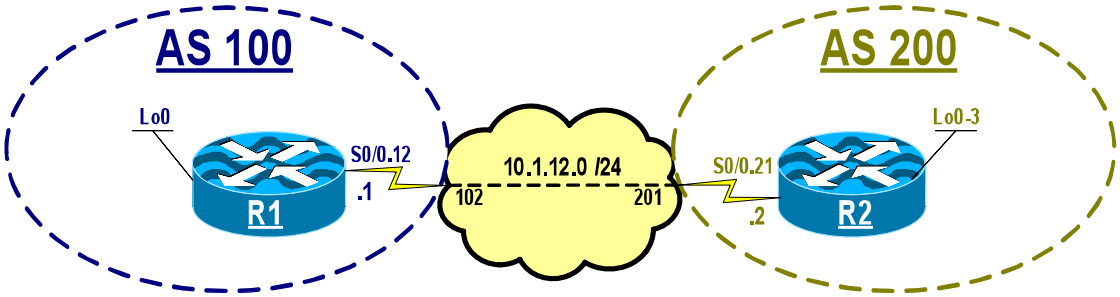
dampening 30 750 2200 80 (route-map TST 10)
  Half-life time       : 30 mins       Decay Time        : 1995 secs
  Max suppress penalty:  4755          Max suppress time: 80 mins
  Suppress penalty     :  2200         Reuse penalty     : 750
```

Task 9

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 5
Route Aggregation



Lab Setup:

- Configure frame-relay connection between the routers in a point-to-point manner.
- Use the following IP addressing chart for IP assignment.

Ip addressing;

Router	Interface	IP Address	AS number
R1	Lo0 Frame-Relay connection to R2	1.1.1.1 /8 10.1.12.1 /24	100
R2	Lo0 Lo1 Lo2 Lo3 Frame-Relay connection to R1	2.2.0.2 / 24 2.2.1.2 /24 2.2.2.2 /24 2.2.3.2 /24 10.1.12.2 /24	200

Task 1

Configure an EBGP session between the two routers and only advertise their loopback interface/s in BGP. R1 should be in AS 100 and R2 should be configured in AS 200. Diable auto-summarization on both routers.

On R1:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config)#router bgp 100
R1(config-router)#No au
R1(config-router)#netw 1.0.0.0
R1(config-router)#neighbor 10.1.12.2 remote-as 200
```

On R2:

```
R2(config)#router bgp 200
R2(config-router)#No au
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#network 2.2.0.0 mask 255.255.255.0
R2(config-router)#network 2.2.1.0 mask 255.255.255.0
R2(config-router)#network 2.2.2.0 mask 255.255.255.0
R2(config-router)#network 2.2.3.0 mask 255.255.255.0
```

To verify the configuration

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.2.0.0/24       10.1.12.2                0             0 200 i
*> 2.2.1.0/24       10.1.12.2                0             0 200 i
*> 2.2.2.0/24       10.1.12.2                0             0 200 i
*> 2.2.3.0/24       10.1.12.2                0             0 200 i
```

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.12.1                0             0 100 i
*> 2.2.0.0/24       0.0.0.0                  0         32768 i
*> 2.2.1.0/24       0.0.0.0                  0         32768 i
*> 2.2.2.0/24       0.0.0.0                  0         32768 i
*> 2.2.3.0/24       0.0.0.0                  0         32768 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 2

Configure R2 such that it summarizes its loopback interfaces and advertises a single summary to R1. R2 should **not** assign an "atomic-aggregate" attribute when it advertises the aggregate route. **Do not** configure a static route to accomplish this task.

Note: in BGP, an aggregate is only created if at least one of the specific routes of the aggregate exists in the BGP table. It is recommended to configure most if not all the specific routes with a network statement, because if only a single network is configured to satisfy the requirements. That particular network goes down, then, the aggregate will be removed.

There are many ways to advertise an aggregate, one way to advertise an aggregate is by creating a static route to "NULL0" that matches the aggregate route and then advertising the aggregate in BGP as follows:

```
(config)#ip route 2.2.0.0 255.255.252.0 null 0

(config)#router bgp 200
(config-router)#network 2.2.0.0 mask 255.255.252.0
```

However, since that is not an option here, every specific prefix of the aggregate must be advertised individually.

Let's advertise all the specific routes on R2:

On R2:

```
R2(config)#router bgp 200
R2(config-router)#aggregate-address 2.2.0.0 255.255.252.0
```

To verify the configuration:

On R1:

```
R1#Show ip bgp b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.2.0.0/24       10.1.12.2                0             0 200 i
*> 2.2.0.0/22       10.1.12.2                0             0 200 i
*> 2.2.1.0/24       10.1.12.2                0             0 200 i
*> 2.2.2.0/24       10.1.12.2                0             0 200 i
*> 2.2.3.0/24       10.1.12.2                0             0 200 i
```

On R2:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.12.1                0             0 100 i
*> 2.2.0.0/24       0.0.0.0                  0         32768 i
*> 2.2.0.0/22       0.0.0.0                            32768 i
*> 2.2.1.0/24       0.0.0.0                  0         32768 i
*> 2.2.2.0/24       0.0.0.0                  0         32768 i
*> 2.2.3.0/24       0.0.0.0                  0         32768 i
```

By default in BGP the aggregate and all the specific routes are advertised. A "summary-only" argument used with the aggregate-address will suppress the specific routes so that only the aggregate route is advertised. You can see from the above output that none of the specific prefixes are suppressed.

On R2:

```
R2(config)#router bgp 200
R2(config-router)#aggregate-address 2.2.0.0 255.255.252.0 summary-only
```

Note: The "summary-only" keyword will suppress all the specific prefixes under the aggregate address on R2, and because of this, R2's neighbor/s will not see the specific routes of the aggregate in their BGP table.

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.12.1                0             0 100 i
s> 2.2.0.0/24       0.0.0.0                  0         32768 i
*> 2.2.0.0/22       0.0.0.0                            32768 i
s> 2.2.1.0/24       0.0.0.0                  0         32768 i
s> 2.2.2.0/24       0.0.0.0                  0         32768 i
s> 2.2.3.0/24       0.0.0.0                  0         32768 i
```

Note: The letter "S" to the left of the ">" sign. The "S" means that these prefixes are suppressed.

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.2.0.0/22       10.1.12.2                0             0 200 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Note: The aggregate route is the only route advertised to R1. Let's see the attributes that are automatically attached to the aggregate route:

On R1:

```
R1#Show ip bgp 2.2.0.0/22

BGP routing table entry for 2.2.0.0/22, version 7
Paths: (1 available, best #1, table Default-IP-Routing-Table)
 Not advertised to any peer
 200, (aggregated by 200 2.2.3.2)
   10.1.12.2 from 10.1.12.2 (2.2.3.2)
    Origin IGP, metric 0, localpref 100, valid, external, atomic-aggregate, best
```

On R2:

```
R2#Show ip bgp 2.2.0.0/22

BGP routing table entry for 2.2.0.0/22, version 7
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to update-groups:
        1
  Local, (aggregated by 200 2.2.3.2)
    0.0.0.0 from 0.0.0.0 (2.2.3.2)
      Origin IGP, localpref 100, weight 32768, valid, aggregated, local, atomic-
aggregate, best
```

Note: The output of the "show ip bgp 2.2.0.0/22" command above, displays two different attributes, the "aggregator" and the "atomic-aggregate" attribute.

The "aggregator" attribute identifies the AS number that the aggregation was performed and it identifies the router-id of the router that performed the aggregation.

By default, when aggregation is configured in BGP, the "atomic-aggregate" attribute is attached to the aggregate address; this alarms the administrator that certain information could be hidden.

We know that the specific routes under an aggregate address are always suppressed/hidden when the summary or aggregation is performed under any routing protocol.

However, when aggregation is performed in BGP, certain information is lost. In BGP that information is not only the specific routes under that aggregate, but also it can also be the AS number/s that the prefixes traversed through to get to the router that is performing the aggregation. If it has not corrected a routing loop can occur. In order to prevent the routing loops from occurring, the "AS-SET" argument should be used when performing aggregation.

An atomic-aggregate – This attribute is assigned to the aggregate route automatically if the "as-set"

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

argument is not used in the "aggregate-address" command.

The "as-set" argument used in the aggregate-address command reveals the AS number/s that some, if not all the specific routes were originated from, once that information is revealed, the "atomic-aggregate" attribute is automatically removed.

On R2:

```
R2(config)#router bgp 200
R2(config-router)#aggregate-address 2.2.0.0 255.255.252.0 summary-only as-set
```

To verify the configuration:

On R1:

```
R1#Sh ip bgp 2.2.0.0/22

BGP routing table entry for 2.2.0.0/22, version 12
Paths: (1 available, best #1, table Default-IP-Routing-Table)
Flag: 0x820
  Not advertised to any peer
  200, (aggregated by 200 2.2.3.2)
    10.1.12.2 from 10.1.12.2 (2.2.3.2)
      Origin IGP, metric 0, localpref 100, valid, external, best
```

On R2:

```
R2#Sh ip bgp 2.2.0.0/22

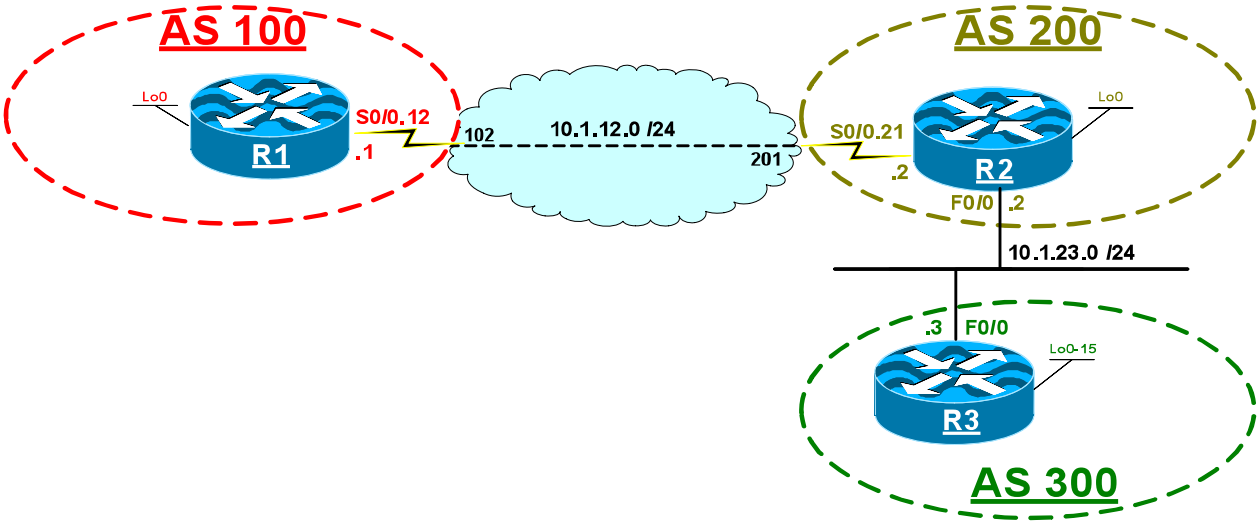
BGP routing table entry for 2.2.0.0/22, version 12
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to update-groups:
        1
Local, (aggregated by 200 2.2.3.2)
 0.0.0.0 from 0.0.0.0 (2.2.3.2)
  Origin IGP, localpref 100, weight 32768, valid, aggregated, local, best
```

Note: The atomic-aggregate is no longer attached to the aggregate-address.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Reconfigure the routers:

Reconfigure the routers using the follows diagram/IP addressing information and **only** advertise their loopback interfaces in BGP. You can use the initial config file for advertising and setting up the diagram.



Lab Setup:

- Configure a frame-relay point-to-point connection between routers R1 and R2.
- Configure the F0/0 interface of R2 and R3 in VLAN 23.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP addressing:

Router	Interface	IP Address	AS number
R1	Lo0 S0/0.12	3.1.11.1 /24 10.1.12.1 /24	100
R2	**Lo0** **S0/0.21** **F0/0**	**2.2.2.2 /8** **10.1.12.2 /24** **10.1.23.2 /24**	**200**
R3	Lo0 Lo1 Lo2 Lo3 Lo4 Lo5 Lo6 Lo7 Lo8 Lo9 Lo10 Lo12 Lo13 Lo14 Lo15 F0/0	3.1.0.3 /24 3.1.1.3 /24 3.1.2.3 /24 3.1.3.3 /24 3.1.4.3 /24 3.1.5.3 /24 3.1.6.3 /24 3.1.7.3 /24 3.1.8.3 /24 3.1.9.3 /24 3.1.10.3 /24 3.1.12.3 /24 3.1.13.3 /24 3.1.14.3 /24 3.1.15.3 /24 10.1.23.3 /24	300

Task 3

Configure R1 in AS 100 to establish an EBGP session with R2 in AS 200, and R2 in AS 200 should establish an EBGP peer session with R3 in AS 300. These routers should advertise their loopback interface/s in their AS and disable auto summarization.

On R1:

```
R1(config)#router bgp 100
R1(config-router)#No au
R1(config-router)#netw 3.1.11.0 mask 255.255.255.0
R1(config-router)#neighbor 10.1.12.2 remote-as 200
```

On R2:

```
R2(config)#router bgp 200
R2(config-router)#No au
R2(config-router)#neighbor 10.1.23.3 remote-as 300
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#network 2.0.0.0
```

R3's configuration can be downloaded from the initial config file provided with this work book

On R3:

```
R3(config)#router bgp 300
R3(config-router)#No au
R3(config-router)#neighbor 10.1.23.2 remote-as 200
R3(config-router)#netw 3.1.0.0 mask 255.255.255.0
R3(config-router)#netw 3.1.1.0 mask 255.255.255.0
R3(config-router)#netw 3.1.2.0 mask 255.255.255.0
R3(config-router)#netw 3.1.3.0 mask 255.255.255.0
R3(config-router)#netw 3.1.4.0 mask 255.255.255.0
R3(config-router)#netw 3.1.5.0 mask 255.255.255.0
R3(config-router)#netw 3.1.6.0 mask 255.255.255.0
R3(config-router)#netw 3.1.7.0 mask 255.255.255.0
R3(config-router)#netw 3.1.8.0 mask 255.255.255.0
R3(config-router)#netw 3.1.9.0 mask 255.255.255.0
R3(config-router)#netw 3.1.10.0 mask 255.255.255.0
R3(config-router)#netw 3.1.12.0 mask 255.255.255.0
R3(config-router)#netw 3.1.13.0 mask 255.255.255.0
R3(config-router)#netw 3.1.14.0 mask 255.255.255.0
R3(config-router)#netw 3.1.15.0 mask 255.255.255.0
```

To verify the configuration:

On R2:

```
R2#Sh ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.1.0.0/24       10.1.23.3                0             0 300 i
*> 3.1.1.0/24       10.1.23.3                0             0 300 i
*> 3.1.2.0/24       10.1.23.3                0             0 300 i
*> 3.1.3.0/24       10.1.23.3                0             0 300 i
*> 3.1.4.0/24       10.1.23.3                0             0 300 i
*> 3.1.5.0/24       10.1.23.3                0             0 300 i
*> 3.1.6.0/24       10.1.23.3                0             0 300 i
*> 3.1.7.0/24       10.1.23.3                0             0 300 i
*> 3.1.8.0/24       10.1.23.3                0             0 300 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
*> 3.1.9.0/24        10.1.23.3                0             0 300 i
*> 3.1.10.0/24       10.1.23.3                0             0 300 i
*> 3.1.11.0/24       10.1.12.1                0             0 100 i
*> 3.1.12.0/24       10.1.23.3                0             0 300 i
*> 3.1.13.0/24       10.1.23.3                0             0 300 i
*> 3.1.14.0/24       10.1.23.3                0             0 300 i
*> 3.1.15.0/24       10.1.23.3                0             0 300 i
```

On R1:

```
R1#Sh ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.1.0.0/24       10.1.12.2                               0 200 300 i
*> 3.1.1.0/24       10.1.12.2                               0 200 300 i
*> 3.1.2.0/24       10.1.12.2                               0 200 300 i
*> 3.1.3.0/24       10.1.12.2                               0 200 300 i
*> 3.1.4.0/24       10.1.12.2                               0 200 300 i
*> 3.1.5.0/24       10.1.12.2                               0 200 300 i
*> 3.1.6.0/24       10.1.12.2                               0 200 300 i
*> 3.1.7.0/24       10.1.12.2                               0 200 300 i
*> 3.1.8.0/24       10.1.12.2                               0 200 300 i
*> 3.1.9.0/24       10.1.12.2                               0 200 300 i
*> 3.1.10.0/24      10.1.12.2                               0 200 300 i
*> 3.1.11.0/24      0.0.0.0                  0         32768 i
*> 3.1.12.0/24      10.1.12.2                               0 200 300 i
*> 3.1.13.0/24      10.1.12.2                               0 200 300 i
*> 3.1.14.0/24      10.1.12.2                               0 200 300 i
*> 3.1.15.0/24      10.1.12.2                               0 200 300 i
```

Task 4

R2 should aggregate all the networks in 3.1.0.0 address space and advertise a single aggregate route. Ensure that the atomic-aggregate attribute is not attached to the aggregate route. This aggregation should be configured such that R1 in AS 100 is the only AS that receives the aggregate route. R3 in AS 300 should **not** receive the aggregate route. R1 should use R2 as the next hop to reach any of the specific routes within the aggregate; R1 should NOT use R2 if its network 3.1.11.0 /24 network is down. R3 does **not** need NLRI to network 3.1.11.0 /24 advertised by R1.

On R2:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

R2 should aggregate all the networks in 3.1.0.0 address space and advertise a single aggregate route that only aggregates the specific routes for subnets under the 3.0.0.0 network in its BGP table. Ensure that the atomic-aggregate attribute is not attached to the aggregate route.

```
R2(config)#router bgp 200
R2(config-router)#aggregate-address 3.1.0.0 255.255.240.0 summary-only as-set
```

This command aggregates networks 3.1.0.0 /24 – 3.1.15.0 /24 and only advertises the summary route and not all the specific routes; the "summary-only" argument accomplishes that. This aggregate route will not have the "atomic-aggregate" attribute attached because the "as-set" argument is used.

To verify the configuration:

On R2:

```
R2#Sh ip bgp | b Network
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	2.0.0.0	0.0.0.0	0		32768	i
s>	3.1.0.0/24	10.1.23.3	0		0	300 i
*>	3.1.0.0/20	0.0.0.0		100	32768	{300,100} i
s>	3.1.1.0/24	10.1.23.3	0		0	300 i
s>	3.1.2.0/24	10.1.23.3	0		0	300 i
s>	3.1.3.0/24	10.1.23.3	0		0	300 i
s>	3.1.4.0/24	10.1.23.3	0		0	300 i
s>	3.1.5.0/24	10.1.23.3	0		0	300 i
s>	3.1.6.0/24	10.1.23.3	0		0	300 i
s>	3.1.7.0/24	10.1.23.3	0		0	300 i
s>	3.1.8.0/24	10.1.23.3	0		0	300 i
s>	3.1.9.0/24	10.1.23.3	0		0	300 i
s>	3.1.10.0/24	10.1.23.3	0		0	300 i
s>	3.1.11.0/24	10.1.12.1	0		0	100 i
s>	3.1.12.0/24	10.1.23.3	0		0	300 i
s>	3.1.13.0/24	10.1.23.3	0		0	300 i
s>	3.1.14.0/24	10.1.23.3	0		0	300 i
s>	3.1.15.0/24	10.1.23.3	0		0	300 i

Note: The specific routes are suppressed (The letter "S" to the left of the ">" sign). The curly brackets in the path column identify the AS numbers that the specific routes originated from. Because the AS numbers are now included in the path column, neither R1 in AS 100 nor R3 in AS 300 will have the aggregate route in their BGP table. They can see their own AS number in the updates and because of the BGP's loop avoidance mechanism they will discard the aggregate.

On R1:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1#Sh ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.1.11.0/24      0.0.0.0                  0         32768 i
```

On R3:

```
R3#Sh ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.23.2                0             0 200 i
*> 3.1.0.0/24       0.0.0.0                  0         32768 i
*> 3.1.1.0/24       0.0.0.0                  0         32768 i
*> 3.1.2.0/24       0.0.0.0                  0         32768 i
*> 3.1.3.0/24       0.0.0.0                  0         32768 i
*> 3.1.4.0/24       0.0.0.0                  0         32768 i
*> 3.1.5.0/24       0.0.0.0                  0         32768 i
*> 3.1.6.0/24       0.0.0.0                  0         32768 i
*> 3.1.7.0/24       0.0.0.0                  0         32768 i
*> 3.1.8.0/24       0.0.0.0                  0         32768 i
*> 3.1.9.0/24       0.0.0.0                  0         32768 i
*> 3.1.10.0/24      0.0.0.0                  0         32768 i
*> 3.1.12.0/24      0.0.0.0                  0         32768 i
*> 3.1.13.0/24      0.0.0.0                  0         32768 i
*> 3.1.14.0/24      0.0.0.0                  0         32768 i
*> 3.1.15.0/24      0.0.0.0                  0         32768 i
```

Note: R1 and/or R3 can not see the aggregate route in their BGP table.

This aggregation should be configured such that R1 in AS 100 is the only AS that receives the aggregate route, R3 or future peer neighbors should not receive the aggregate route. R1 should use R2 as the next hop to reach any of the specific routes within the aggregate. This portion of the task can be configured using the "advertise-map" map, as follows:

```
R2(config)#ip as-path access-list 1 permit ^300$
```

The above command identifies AS number 300.

```
R2(config)#route-map TST permit 10
R2(config-route-map)#match as-path 1
```

The "route-map" command references the "as-path access-list 1".

```
R2(config)#router bgp 200
R2(config-router)#aggregate-address 3.1.0.0 255.255.240.0 as-set summary-
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
only advertise-map TST
```

The "advertise-map" command assigns the route-map "TST" to the "aggregate address" command.

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.1.0.0/20       10.1.12.2                0             0 200 300 i
*> 3.1.11.0/24      0.0.0.0                  0         32768 i
```

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.23.2                0             0 200 i
*> 3.1.0.0/24       0.0.0.0                  0         32768 i
*> 3.1.1.0/24       0.0.0.0                  0         32768 i
*> 3.1.2.0/24       0.0.0.0                  0         32768 i
*> 3.1.3.0/24       0.0.0.0                  0         32768 i
*> 3.1.4.0/24       0.0.0.0                  0         32768 i
*> 3.1.5.0/24       0.0.0.0                  0         32768 i
*> 3.1.6.0/24       0.0.0.0                  0         32768 i
*> 3.1.7.0/24       0.0.0.0                  0         32768 i
*> 3.1.8.0/24       0.0.0.0                  0         32768 i
*> 3.1.9.0/24       0.0.0.0                  0         32768 i
*> 3.1.10.0/24      0.0.0.0                  0         32768 i
*> 3.1.12.0/24      0.0.0.0                  0         32768 i
*> 3.1.13.0/24      0.0.0.0                  0         32768 i
*> 3.1.14.0/24      0.0.0.0                  0         32768 i
*> 3.1.15.0/24      0.0.0.0                  0         32768 i
```

Note: R3 gets the aggregate route but it rejects it because it sees its own AS number in the as-path list. R1 receives and processes the aggregate route because it does not see its own AS number in the as-path list advertised by R2.

The following shows all the routes that are advertised by R2 to its neighbor 10.1.23.3:

On R2:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2#Show ip bgp neighbors 10.1.23.3 advertised-routes

BGP table version is 64, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

     Network          Next Hop      Metric LocPrf Weight Path
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.1.0.0/20       0.0.0.0                              100   32768 300 i

Total number of prefixes 2
```

The output of the following display shows all the routes received and accepted by R3:

On R3:

```
R3#Show ip bgp neighbor 10.1.23.2 routes | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.23.2                0             0 200 i
```

R1 should not use R2 if its network 3.1.11.0 /24 network is down. This can be accomplished with a simple static route. If this static route is not configured, and R1's 3.1.11.0 /24 is down, R1 will use the longest match prefix in its routing table, which is the aggregate route.

On R1:

```
R1(config)#ip route 3.1.11.0 255.255.255.0 NULL0
```

To verify the configuration:

On R1:

```
R1#Show ip route | b Gate
Gateway of last resort is not set

B    2.0.0.0/8 [20/0] via 10.1.12.2, 00:32:00
     3.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
B       3.1.0.0/20 [20/0] via 10.1.12.2, 00:05:44
C       3.1.11.0/24 is directly connected, Loopback0
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.12.0 is directly connected, Serial0/0.12
```

Note: 3.1.11.0 /24 is directly connected, to test this condition we should shut down the interface and

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

check the routing table again, as follows:

On R1:

```
R1(config)#int lo0
R1(config-if)#shut

R1#Show ip route | b Gate

Gateway of last resort is not set

B    2.0.0.0/8 [20/0] via 10.1.12.2, 00:33:58
     3.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
B       3.1.0.0/20 [20/0] via 10.1.12.2, 00:07:42
S       3.1.11.0/24 is directly connected, Null0
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.12.0 is directly connected, Serial0/0.12
```

Note: If the interface is down all the traffic destined for the network is forwarded to the NULL0 interface, and therefore, it will not be forwarded to R2.

Let's enable the Lo0 interface of R1

On R1:

```
R1(config)#int lo0
R1(config-if)#NO shut
```

Task 5

Configure R2 such that a cost of 50 is assigned to the aggregate route.

On R2:

```
R2(config)#route-map COST permit 10
R2(config-route-map)#set metric 50
R2(config)#route-map COST permit 90

R2(config)#router bgp 200
R2(config-router)# aggregate-address 3.1.0.0 255.255.240.0 summary-only
advertise-map TST as-set attribute-map COST
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The attribute map identifies the name of the route-map used to set the attribute/s of the aggregate route.

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.1.0.0/20       10.1.12.2               50             0 200 300 i
*> 3.1.11.0/24      0.0.0.0                  0         32768 i
```

Note: R1 gets the aggregate route with a cost of 50.

Task 6

The policy for R1 requires that it should receive the aggregate route plus one of the specific subnets (3.1.3.0 /24). This policy should be configured and tested in three different ways using a suppress-map, and unsuppress-map.

To test the Suppress-map scenario #1:

On R2:

```
R2(config)#access-list 1 deny 3.1.3.0 0.0.0.255
R2(config)#access-list 1 permit any

R2(config)#route-map SUPP permit 10
R2(config-route-map)#match ip addr 1

R2(config)#router bgp 200
R2(config-router)# aggregate-address 3.1.0.0 255.255.240.0 summary-only
advertise-map TST as-set attribute-map COST suppress-map SUPP
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.1.0.0/20       10.1.12.2               50             0 200 300 i
*> 3.1.3.0/24       10.1.12.2                              0 200 300 i
*> 3.1.11.0/24      0.0.0.0                  0         32768 i
```

To test the suppress-map scenario #2:

```
R2(config)#No access-list 1

R2(config)#access-list 1 permit 3.1.3.0 0.0.0.255

R2(config)#router bgp 200
R2(config-router)#No aggregate-address 3.1.0.0 255.255.240.0

R2(config-router)#aggregate-address 3.1.0.0 255.255.240.0 summary-only
advertise-map TST as-set attribute-map COST

R2(config-router)#neighbor 10.1.12.1 unsuppress-map SUPP
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.1.0.0/20       10.1.12.2               50             0 200 300 i
*> 3.1.3.0/24       10.1.12.2                              0 200 300 i
*> 3.1.11.0/24      0.0.0.0                  0         32768 i
```

Task 7

Remove the configuration command/s from the previous step. The aggregate command should NOT be removed.

On R2:

```
R2(config)#No route-map SUPP
R2(config)#No access-list
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config)#router bgp 200
R2(config-router)#No neighbor 10.1.12.1 unsuppress-map SUPP
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Net

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.1.0.0/20       10.1.12.2               50             0 200 300 i
*> 3.1.3.0/24       10.1.12.2                              0 200 300 i
*> 3.1.11.0/24      0.0.0.0                  0         32768 i
```

Task 8

Configure R1 so it has the aggregate route plus the specific route that it wanted to have in its BGP table (3.1.3.0 /24). R1 should **not** advertise this subnet, configure a static route or use the redistribute command to accomplish this task. R2 should **not** be configured for this task.

To accomplish this task on R1, we can use the combination of an "exist-map" and an "inject –map". In the following configuration, the exist-map matches on the aggregate address and the router that advertised the aggregate address (the "route-source" command in the route-map called "EXIST"). The inject-map injects the IP address/es identified by the route-map called "INJECT", if the condition of the "exist-map" is true which means that the router that advertised the aggregate address and the aggregate address exist, then, inject whatever that is specified in the prefix-list that is referenced by the route-map INJECT.

To verify the BGP table of R1 before the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.1.0.0/20       10.1.12.2               50             0 200 300 i
*> 3.1.11.0/24      0.0.0.0                  0         32768 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To configure this task:

On R1:

The specific route, the aggregate and R2 are identified in the following prefix-lists:

```
R1(config)#ip prefix-list NET permit 3.1.3.0/24
R1(config)#ip prefix-list AGG permit 3.1.0.0/20
R1(config)#ip prefix-list R2 permit 10.1.12.2/32
```

The following route-map called "EXIST", tells the IOS that if the aggregate route exists and R2 is UP (which means that it also exists).

```
R1(config)#route-map EXIST permit 10
R1(config-route-map)#match ip addr prefix-list AGG
R1(config-route-map)#match ip route-source prefix-list R2
```

The following route-map called INJECT, matches on the specific route that must be injected in R1's BGP table:

```
R1(config)#route-map INJECT permit 10
R1(config-route-map)#set ip addr prefix-list NET
```

The "EXIST" and the "INJECT" route-maps are referenced under router bgp, instructing BGP that if the router that originated the aggregate route and the aggregate route itself are up, then, inject this specific route:

```
R1(config)#router bgp 100
R1(config-router)#bgp inject-map INJECT exist-map EXIST
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network
   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.1.0.0/20       10.1.12.2               50             0 200 300 i
*> 3.1.3.0/24       10.1.12.2                              0 ?
*> 3.1.11.0/24      0.0.0.0                  0         32768 i
```

Note: The subnet is in the BGP table of R1.

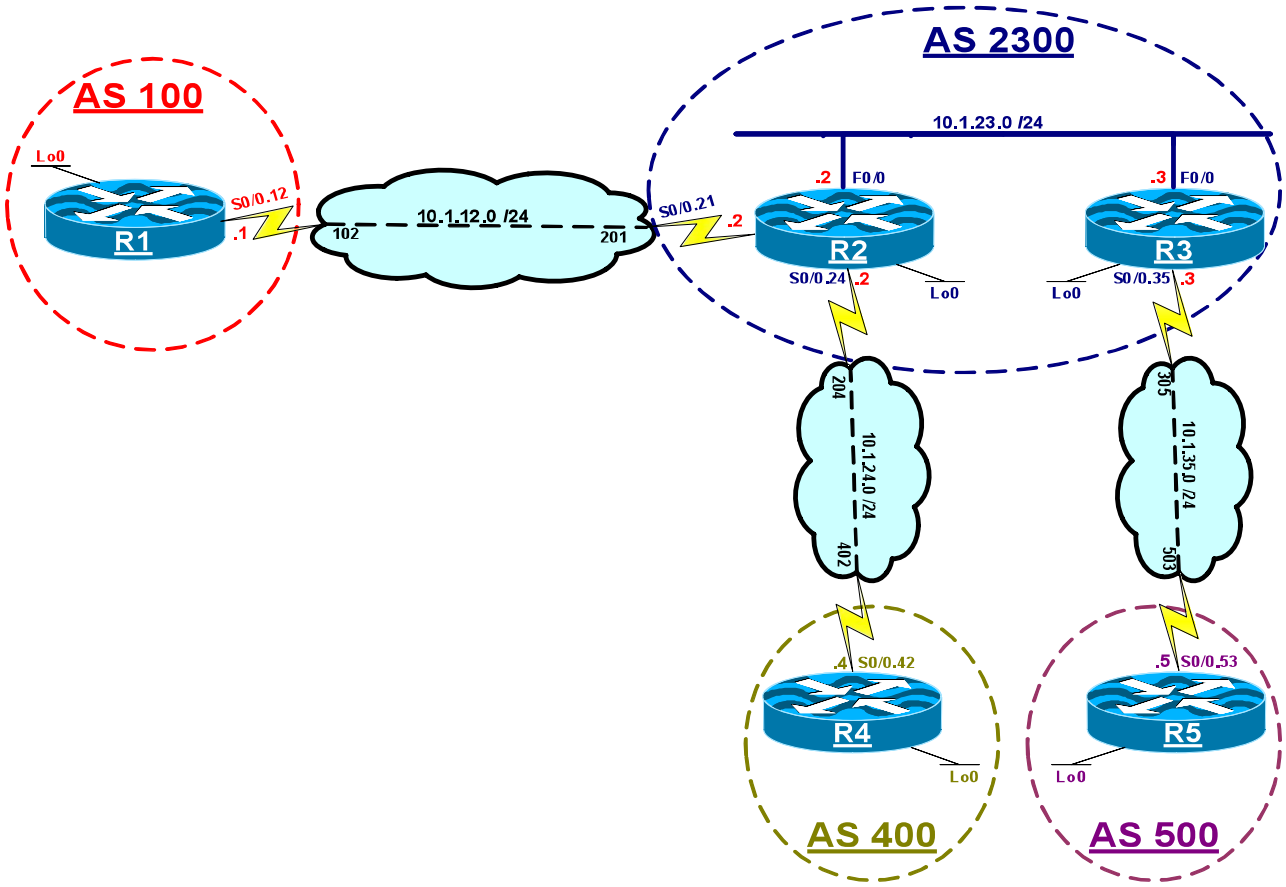
©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 9

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 6
The Community Attribute



Lab Setup:

- Configure the routers that are connected to the frame-relay clouds in a point-to-point manner.
- R2 and R3's F0/0 interface should be configured in VLAN 23.
- Use the following IP address chart for IP assignment.

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

AS & IP addressing chart:

Router	Interface	IP Address	AS Number
R1	Lo0 S0/0.12	1.1.1.1 /8 10.1.12.1 /24	**100**
R2	**Lo0** **S0/0.21** **S0/0.24** **F0/0 interface connection to R3**	**2.2.2.2 /8** **10.1.12.2 /24** **10.1.24.2 /24** **10.1.23.2 /24**	**2300**
R3	Lo0 F0/0 interface connection to R2 S0/0.35	3.3.3.3 /8 10.1.23.3 /24 10.1.35.3 /24	**2300**
R4	**Lo0** **S0/0.42**	**4.4.4.4 /8** **10.1.24.4 /24**	**400**
R5	Lo0 S0/0.53	5.5.5.5 /8 10.1.35.5 /24	**500**

Task 1

Configure EBGP peer session/s between the routers based on the above "AS & IP addressing chart". These routers should **only** advertise their loopback interface/s in BGP. These BGP routers should use their loopback0's IP address as their Router id. Ensure that every router has NLRI to every link in this lab using RIPv2 with auto summarization disabled.

On All Routers

```
Rx (config)#router rip
Rx(config-router)#No au
Rx(config-router)#ver 2
Rx(config-router)#network 10.0.0.0
```

On R1:

```
R1(config)#router bgp 100
R1(config-router)#bgp router-id 1.1.1.1
R1(config-router)#netw 1.0.0.0
R1(config-router)#No au
R1(config-router)#neighbor 10.1.12.2 remote-as 2300
```

On R2:

```
R2(config)#router bgp 2300
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config-router)#bgp router-id 2.2.2.2
R2(config-router)#netw 2.0.0.0
R2(config-router)#No au
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#neighbor 10.1.23.3 remote-as 2300
R2(config-router)#neighbor 10.1.24.4 remote-as 400
```

On R3:

```
R3(config)#router bgp 2300
R3(config-router)#bgp router-id 3.3.3.3
R3(config-router)#netw 3.0.0.0
R3(config-router)#No au
R3(config-router)#neighbor 10.1.23.2 remote-as 2300
R3(config-router)#neighbor 10.1.35.5 remote-as 500
```

On R4:

```
R4(config)#router bgp 400
R4(config-router)#bgp router-id 4.4.4.4
R4(config-router)#netw 4.0.0.0
R4(config-router)#No au
R4(config-router)#neighbor 10.1.24.2 remote-as 2300
```

On R5:

```
R5(config)#router bgp 500
R5(config-router)#bgp router-id 5.5.5.5
R5(config-router)#netw 5.0.0.0
R5(config-router)#No au
R5(config-router)#neighbor 10.1.35.3 remote-as 2300
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 2300 i
*> 3.0.0.0          10.1.12.2                              0 2300 i
*> 4.0.0.0          10.1.12.2                              0 2300 400 i
*> 5.0.0.0          10.1.12.2                              0 2300 500 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 2

Using the community attribute, configure R1 such that when it advertises network 1.0.0.0 /8 to R2 in AS 200, the network is not advertised to any of R2's IBGP or EBGP neighbor/s.

The community attribute is a numerical value that can be attached to a given prefix and advertised to a specific neighbor, once the neighbor receives the prefix, it will examine the community value and it will perform either filtering or use that value for route selection process.

By default, no community attribute/s is sent to any neighbor. To specify that a community attribute should be sent to a BGP neighbor, the "neighbor send-community" command is configured in the router config mode.

The well-known communities are as follows:

- **Internet – If assigned to a network/s, that network/s should be advertised to all neighbors.**
- **Local-as – If assigned to a network/s, that network/s should only be advertised within that AS.**
- **No-advertise – If assigned to a network/s. that network/s should not be advertised to any BGP neighbor.**
- **No-export – If assigned to a network/s, that network/s should not be advertised to an EBGP neighbor.**

On R1:

Note: Before configuring an access-list, always performs a "Show access-list" command to ensure that one does not exist.

```
R1(config)#access-list 1 permit 1.0.0.0 0.255.255.255
```

The access-list is used to identify the network; Prefix-lists can also be used for this purpose.

```
R1(config)#route-map TEST permit 10
R1(config-route-map)#match ip addr 1
R1(config-route-map)#set community no-advertise

R1(config-route-map)#route-map TEST permit 90
```

Note: The above route-map matches on the access-list and sets the community to "no-advertise", this well-known community attribute tells the receiving router not to advertise the prefix to any of its neighbor/s (IBGP or EBGP).

The "route-map TST permit 90" is the catchall route-map; it matches any network not matched with the match keyword in the "route-map TST permit 10".

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.12.2 send-community
R1(config-router)#neighbor 10.1.12.2 route-map TEST out
```

In the above commands, we are sending the community and assigning the route-map to a given neighbor.

The direction of the route-map specifies which router's decision should be influenced by this policy, if it should affect neighbor's decision, then, the direction of the route-map should be "out", but if the local router's decision should be influenced, then, the direction of the route-map should be "in".

To verify the configuration:

On R2:

```
R2#Sh ip bgp 1.0.0.0

BGP routing table entry for 1.0.0.0/8, version 7
Paths: (1 available, best #1, table Default-IP-Routing-Table, not
advertised to any peer)
Flag: 0x880
  Not advertised to any peer
  100
    10.1.12.1 from 10.1.12.1 (1.1.1.1)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Community: no-advertise
```

Note: The community attribute from R2's perspective. Since R2 does not advertise the network, R3 and the other EBGP and/or IBGP neighbor/s will not have any knowledge of this prefix.

On R3:

```
R3#Sh ip bgp 1.0.0.0

% Network not in table
```

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i2.0.0.0          10.1.23.2                0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          10.1.24.4                0    100      0 400 i
*> 5.0.0.0          10.1.35.5                0             0 500 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4:

```
R4#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.24.2                0             0 2300 i
*> 3.0.0.0          10.1.24.2                              0 2300 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*> 5.0.0.0          10.1.24.2                              0 2300 500 i
```

Task 3

Configure R5 such that when it advertises its network 5.0.0.0 to R3 in AS 2300, the routers in AS 2300 **do not** advertise that network to any of their EBGP peer/s. **Do not** configure filtering on R3 to accomplish this task.

On R5:

Note: Before configuring an access-list, always perform a "show access-list" command to ensure that one does not exist.

```
R5(config-router)#access-list 1 permit 5.0.0.0

R5(config)#route-map TST permit 10
R5(config-route-map)#match ip addr 1
R5(config-route-map)#set community no-export

R5(config-route-map)# route-map TST permit 90

R5(config)#router bgp 500
R5(config-router)#neighbor 10.1.35.3 send-community
R5(config-router)#neighbor 10.1.35.3 route-map TST out
```

This is another well-known community. In this case network 5.0.0.0 will only be advertised to the routers in AS 2300. The routers in AS 2300 will not advertise this network to any of their EBGP neighbor/s. Remember that by default routers will strip the community attibute, therefore, in this case, R3 should be configured to send community to R2, or else R2 will advertise that network to its EBGP peers.

On R3:

```
R3(config)#router bgp 2300
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-router)#neighbor 10.1.23.2 send-community
```

To verify the configuration:

On R3:

```
R3#Sh ip bgp 5.0.0.0

BGP routing table entry for 5.0.0.0/8, version 8
Paths: (1 available, best #1, table Default-IP-Routing-Table, not
advertised to EBGP peer)
Flag: 0x880
  Advertised to update-groups:
        1
  500
    10.1.35.5 from 10.1.35.5 (5.5.5.5)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Community: no-export
```

To test this configuration further, a point-to-point frame-relay connection and an EBGP peer session can be established between R3 and R4. R4 should not receive an update for network 5.0.0.0 from R3, but R4 will receive an update for network 5.0.0.0 from R2.

Task 4

Configure R3 in AS 2300 to advertise network 3.0.0.0 /8 to the routers in its own AS **only**.

On R3:

```
R3(config)#route-map TST permit 10
R3(config-route-map)#set community local-as

R3(config)#router bgp 2300
R3(config-router)#network 3.0.0.0 route-map TST
R3(config-router)#neighbor 10.1.23.2 send-community
```

To verify the configuration:

On R3:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3#Sh ip bgp 3.0.0.0

BGP routing table entry for 3.0.0.0/8, version 9
Paths: (1 available, best #1, table Default-IP-Routing-Table, not
advertised outside local AS)
  Advertised to update-groups:
        1
  Local
    0.0.0.0 from 0.0.0.0 (3.3.3.3)
      Origin IGP, metric 0, localpref 100, weight 32768, valid, sourced,
local, best
      Community: local-AS

R5#Sh ip bgp b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.35.3                              0 2300 i
*> 4.0.0.0          10.1.35.3                              0 2300 400 i
*> 5.0.0.0          0.0.0.0                  0         32768 i
```

Task 5

R1 is advertising network 1.0.0.0 which has an attached community attribute of "no-advertise" to R2 (task 2). Router R2 should be configured to advertise network 1.0.0.0 to all of its IBGP and EBGP peers. You should utilize a well-known community attribute to accomplish this task.

On R2:

Note: Before configuring R2, we should display the prefix in the BGP table:

```
R2#Show ip bgp 1.0.0.0

BGP routing table entry for 1.0.0.0/8, version 7
Paths: (1 available, best #1, table Default-IP-Routing-Table, not
advertised to any peer)
  Not advertised to any peer
  100
    10.1.12.1 from 10.1.12.1 (1.1.1.1)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Community: no-advertise
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

R2 can be configured to assign a well-known community of "internet" to this network; when the "internet" community is assigned to a network, that network will be advertised to all peers.

To configure:

On R2:

```
R2(config)#access-list 1 permit 1.0.0.0

R2(config)#route-map TST permit 10
R2(config-route-map)#match ip addr 1
R2(config-route-map)#set community internet

R2(config)#router bgp 2300
R2(config-router)#neighbor 10.1.12.1 route-map TST in
```

To verify the configuration:

On R2:

```
R2#Show ip bgp 1.0.0.0

BGP routing table entry for 1.0.0.0/8, version 10
Paths: (1 available, best #1, table Default-IP-Routing-Table)
Flag: 0x880
  Advertised to update-groups:
        1    2
  100
    10.1.12.1 from 10.1.12.1 (1.1.1.1)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Community: internet
```

On R4:

```
R4#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.24.2                              0 2300 100 i
*> 2.0.0.0          10.1.24.2                0             0 2300 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
```

On R3:

```
R3#Show ip bgp | b Network
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.12.1                0    100      0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          10.1.24.4                0    100      0 400 i
*> 5.0.0.0          10.1.35.5                0             0 500 i
```

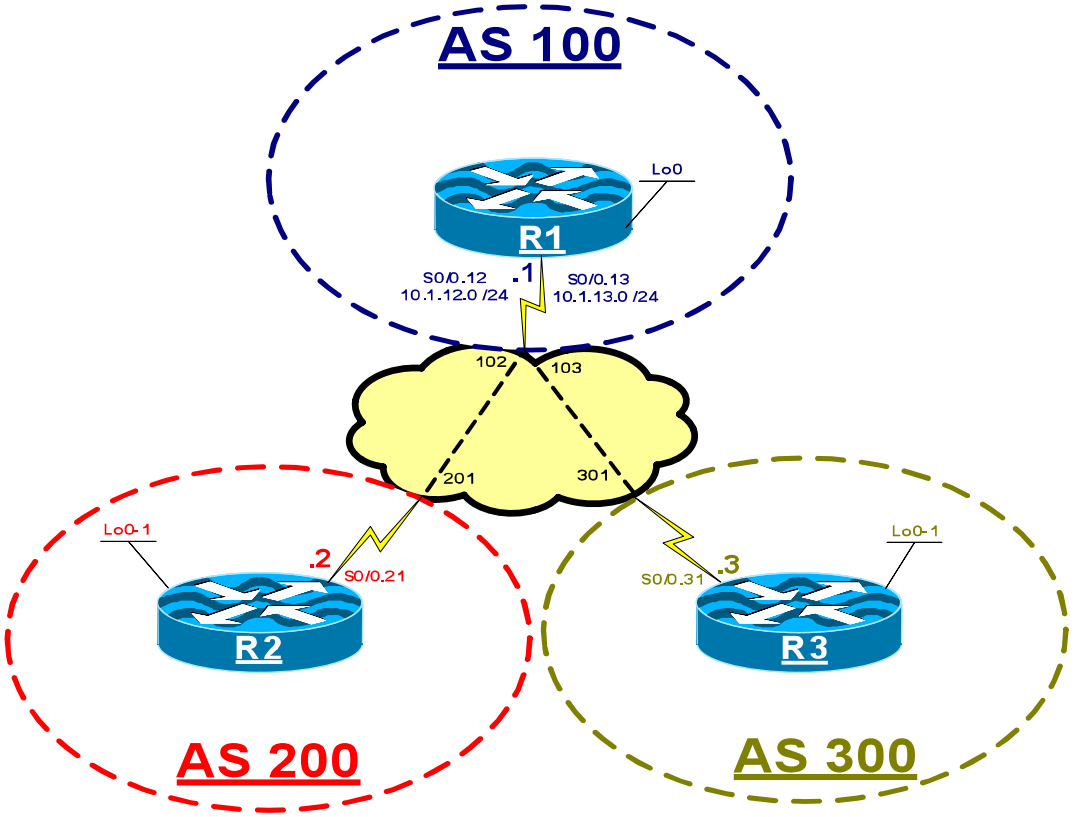
Task 6

Erase the startup config and reload the routers before proceeding to the next task.

Task 7

Configure the routers according to the diagram/chart below and use the IP addressing and AS numbering identified in the chart.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh



Lab Setup:

- Configure the routers that are connected to the frame-relay clouds in a point-to-point manner.
- R1 should have two point-to-point sub-interfaces, one connecting to R2 and the other connecting to R3.
- R2 and R3 should be configured with a single point-to-point connection to R1

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP addressing:

Router	Interface	IP Address	AS number
R1	Lo0 S0/0.12 S0/0.13	1.1.1.1 /8 10.1.12.1 /24 10.1.13.1 /24	100
R2	**Lo0** **Lo1** **S0/0.21**	**20.1.2.2 /24** **20.1.3.2 /24** **10.1.12.2 /24**	**200**
R3	Lo0 Lo1 S0/0.31	30.1.2.3 /24 30.1.3.3 /24 10.1.13.3 /24	300

On R1:

```
R1(config)#router bgp 100
R1(config-router)#netw 1.0.0.0
R1(config-router)#No au
R1(config-router)#neighbor 10.1.12.2 remote-as 200
R1(config-router)#neighbor 10.1.13.3 remote-as 300
```

On R2:

```
R2(config-if)#router bgp 200
R2(config-router)#No au
R2(config-router)#network 20.1.2.0 mask 255.255.255.0
R2(config-router)#network 20.1.3.0 mask 255.255.255.0
R2(config-router)#neighbor 10.1.12.1 remote-as 100
```

On R3:

```
R3(config-if)#router bgp 300
R3(config-router)#No au
R3(config-router)#netw 30.1.2.0 mask 255.255.255.0
R3(config-router)#netw 30.1.3.0 mask 255.255.255.0
3(config-router)#neighbor 10.1.13.1 remote-as 100
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 20.1.2.0/24      10.1.12.2                0             0 200 i
*> 20.1.3.0/24      10.1.12.2                0             0 200 i
*> 30.1.2.0/24      10.1.13.3                0             0 300 i
*> 30.1.3.0/24      10.1.13.3                0             0 300 i
```

On R2:

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.12.1                0             0 100 i
*> 20.1.2.0/24      0.0.0.0                  0         32768 i
*> 20.1.3.0/24      0.0.0.0                  0         32768 i
*> 30.1.2.0/24      10.1.12.1                              0 100 300 i
*> 30.1.3.0/24      10.1.12.1                              0 100 300 i
```

On R3:

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.13.1                0             0 100 i
*> 20.1.2.0/24      10.1.13.1                              0 100 200 i
*> 20.1.3.0/24      10.1.13.1                              0 100 200 i
*> 30.1.2.0/24      0.0.0.0                  0         32768 i
*> 30.1.3.0/24      0.0.0.0                  0         32768 i
```

Task 8

Ensure that R1 uses AS 200 to connect to networks in subnet 2 (20.1.**2**.0 /24 and 30.1.**2**.0 /24) and AS 300 to connect to networks in subnet 3 (20.1.**3**.0 /24 and 30.1.**3**.0 /24). You must use community tags in AS 200 and 300 and neighbor commands on R1 to accomplish this task.

On R2:

```
R2(config)#access-list 2 permit 20.1.2.0 0.0.0.255

R2(config)#access-list 3 permit 20.1.3.0 0.0.0.255

R2(config)#route-map TST permit 10
R2(config-route-map)#match ip addr 2
R2(config-route-map)#set community 2
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config)#route-map TST permit 20
R2(config-route-map)#match ip addr 3
R2(config-route-map)#set community 3
```

Note: The above command "set community" tags the route/s identified in the access-list.

```
R2(config)#route-map TST permit 90

R2(config)#router bgp 200
R2(config-router)#neighbor 10.1.12.1 send-community
R2(config-router)#neighbor 10.1.12.1 route-map TST out
```

On R3:

```
R3(config)#access-list 2 permit 30.1.2.0 0.0.0.255

R3(config)#access-list 3 permit 30.1.3.0 0.0.0.255

R3(config)#route-map TST permit 10
R3(config-route-map)#match ip addr 2
R3(config-route-map)#set community 2

R3(config)#route-map TST permit 20
R3(config-route-map)#match ip addr 3
R3(config-route-map)#set community 3

R3(config)#route-map TST permit 90

R3(config)#router bgp 300
R3(config-router)#neighbor 10.1.13.1 send-community
R3(config-router)#neighbor 10.1.13.1 route-map TST out
```

Note: R1 can display the routes via their assigned community tags:

On R1:

```
R1#Sh ip bgp 20.1.2.0/24

BGP routing table entry for 20.1.2.0/24, version 8
Paths: (1 available, best #1, table Default-IP-Routing-Table)
Flag: 0x880
  Advertised to update-groups:
        1
  200
    10.1.12.2 from 10.1.12.2 (20.1.3.2)
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
      Origin IGP, metric 0, localpref 100, valid, external, best
      Community: 2

R1#Show ip bgp community 2 | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 20.1.2.0/24      10.1.12.2                0             0 200 i
*> 30.1.2.0/24      10.1.13.3                0             0 300 i
```

Note: These are the routes that R2 and R3 tagged using community 2.

```
R1#Show ip bgp community 3 | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 20.1.3.0/24      10.1.12.2                0             0 200 i
*> 30.1.3.0/24      10.1.13.3                0             0 300 i
```

Note: These are the routes that R2 and R3 tagged using community 3.

On R1:

R1 is identifying the community tags using a community-list. This is like writing an access-list to identify a given route/s.

```
R1(config)#ip community-list standard TST2 permit 2

R1(config)#ip community-list standard TST3 permit 3
```

The communities are matched and the policy is assigned.

```
R1(config)#route-map TEST permit 10
R1(config-route-map)#match community TST2
R1(config-route-map)#set ip next-hop 10.1.12.2

R1(config)#route-map TEST permit 20
R1(config-route-map)#match community TST3
R1(config-route-map)#set ip next-hop 10.1.13.3

R1(config)#route-map TEST permit 90
```

The policy is applied inbound, since R1's decision must be affected:

```
R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.12.2 route-map TEST in
R1(config-router)#neighbor 10.1.13.3 route-map TEST in
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The policies are applied to the neighbors using the "neighbor route-map" commands.

Do not try to test reachability to the network through the newly assigned next hop IP addresses; the purpose of this lab is to understand the scope of the community attributes and its uses.

To verify this lab enter "show ip bgp" to see the next hop attribute.

```
R1#Show ip bgp | b Network
```

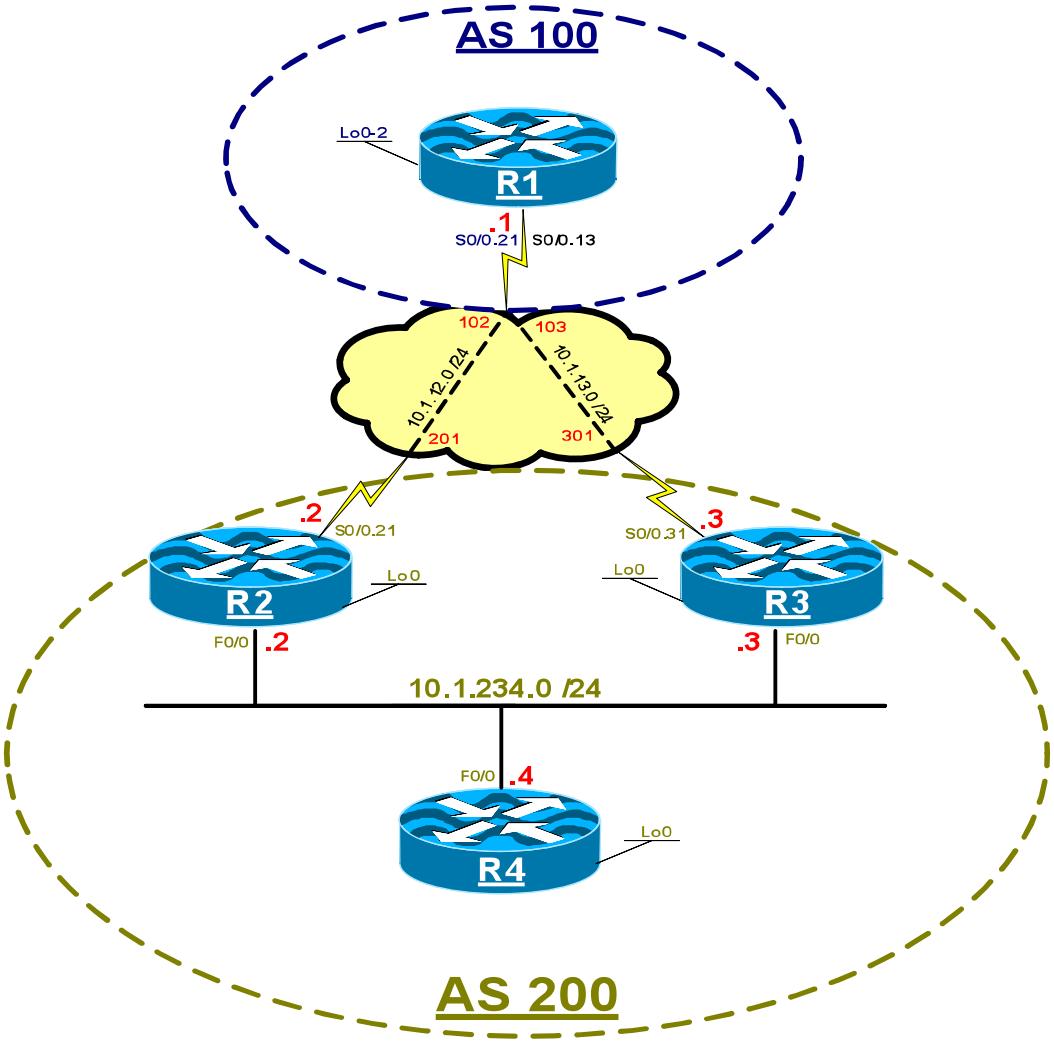
	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	1.0.0.0	0.0.0.0	0		32768	i
*>	20.1.2.0/24	10.1.12.2	0		0	200 i
*>	20.1.3.0/24	10.1.13.3	0		0	200 i
*>	30.1.2.0/24	10.1.12.2	0		0	300 i
*>	30.1.3.0/24	10.1.13.3	0		0	300 i

Task 9

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 7 – BGP Cost Community



Lab Setup:

- Configure R1 to be the hub and R2 and R3 to be the spokes, the frame-relay routers should be configured in a point-to-point manner.
- RIPv2 should be used to provide NLRI for the links. The loopback0 interfaces of R2, R3 and R4 should also be advertised in RIPv2 routing protocol.
- The F0/0 interface of R2, R3 and R4 should be configured in VLAN 234.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

- Use the following IP addressing chart for IP addressing assignment

IP addressing:

Router	Interface / IP address
R1	S0/0.12 = 10.1.12.1 /24 S0/0.13 = 10.1.13.1 /24 Loopback 0 = 1.1.1.1 /8 Loopback 1 = 100.1.1.1 /24 Loopback 2 = 200.1.1.1 /24
R2	**S0/0.21 = 10.1.12.2 /24** **F0/0 = 10.1.234.2 /24** **Loopback0 = 2.2.2.2 /8**
R3	S0/0.31 = 10.1.13.3 /24 F0/0 = 10.1.234.3 /24 Loopback0 = 3.3.3.3 /8
R4	**F0/0 = 10.1.234.4 /24** **Loopback0 = 4.4.4.4 /8**

Task 1

Configure R1 in AS 100; this router should establish an EBGP peer session with R2 and R3 in AS 200. R1 should advertise its Lo1 and Lo2 interfaces in BGP. All BGP routers should use their loopback 0 interface as their router-id.

On R1:

```
R1(config)#router bgp 100
R1(config-router)#bgp router-id 1.1.1.1
R1(config-router)#No auto-summary
R1(config-router)#neigh 10.1.12.2 remote-as 200
R1(config-router)#neigh 10.1.13.3 remote-as 200
R1(config-router)#network 100.1.1.0 mask 255.255.255.0
R1(config-router)#network 200.1.1.0
```

On R2:

```
R2(config)#router bgp 200
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Surinderjeet Singhh

Task 2

Configure an IBGP peer session between R2, R3 and R4; these routers should establish their peer session based on their loopback 0 interface. R4 should use its Lo0 as its BGP router ID. You should provide NLRI using RIPv2 with auto summarization disabled.

On R2:, R3 and R4

```
Rx(config)#router rip
Rx(config-router)#No au
Rx(config-router)#ver 2
Rx(config-router)#network 10.0.0.0
```

On R2:

```
R2(config)#router rip
R2(config-router)#netw 2.0.0.0

R2(config)#router bgp 200
R2(config-router)#neighbor 3.3.3.3 remote-as 200
R2(config-router)#neighbor 3.3.3.3 update-source lo0
R2(config-router)#neighbor 4.4.4.4 remote-as 200
R2(config-router)#neighbor 4.4.4.4 update-source lo0
```

On R3:

```
R3(config)#router rip
R3(config-router)#netw 3.0.0.0

R3(config)#router bgp 200
R3(config-router)#neighbor 2.2.2.2 remote-as 200
R3(config-router)#neighbor 2.2.2.2 update-source lo0
R3(config-router)#neighbor 4.4.4.4 remote-as 200
R3(config-router)#neighbor 4.4.4.4 update-source lo0
```

On R4:

```
R4(config)#router rip
R4(config-router)#netw 4.0.0.0

R4(config)#router bgp 200
R4(config-router)#No auto-summary
R4(config-router)#bgp router-id 4.4.4.4

R4(config-router)#neighbor 2.2.2.2 remote-as 200
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config-router)#neighbor 2.2.2.2 update-source lo0
R4(config-router)#neighbor 3.3.3.3 remote-as 200
R4(config-router)#neighbor 3.3.3.3 update-source lo0
```

To verify the configuration:

On R2:

```
R2#Show ip bgp sum | b Neigh

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
3.3.3.3         4   200       8       8        7    0    0 00:03:10        2
4.4.4.4         4   200       4       5        7    0    0 00:00:35        0
10.1.12.1       4   100      34      32        7    0    0 00:20:28        2
```

On R3:

```
R3#Show ip bgp sum | b Neigh

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
2.2.2.2         4   200       9       9        3    0    0 00:04:51        2
4.4.4.4         4   200       5       6        3    0    0 00:01:59        0
10.1.13.1       4   100      27      26        3    0    0 00:22:00        2
```

On R4:

```
R4#Show ip bgp sum | b Neigh

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
2.2.2.2         4   200       8       7        3    0    0 00:03:23        2
3.3.3.3         4   200       8       7        3    0    0 00:03:07        2
```

Task 3

Configure R2 and R3 such that the output of their "**show ip bgp**" command resembles the following:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2:

```
R2#show ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24     3.3.3.3                  0    100      0 100 i
*>                  10.1.12.1                0             0 100 i
* i200.1.1.0        3.3.3.3                  0    100      0 100 i
*>                  10.1.12.1                0             0 100 i
```

On R3:

```
R3#show ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24     2.2.2.2                  0    100      0 100 i
*>                  10.1.13.1                0             0 100 i
* i200.1.1.0        2.2.2.2                  0    100      0 100 i
*>                  10.1.13.1                0             0 100 i
```

On R4:

```
R4#show ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24     3.3.3.3                  0    100      0 100 i
*>i                 2.2.2.2                  0    100      0 100 i
* i200.1.1.0        3.3.3.3                  0    100      0 100 i
*>i                 2.2.2.2                  0    100      0 100 i
```

Let's check the BGP table of R2, R3 and R4:

On R2:

```
R2#Show ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24     10.1.13.1                0    100      0 100 i
*>                  10.1.12.1                0             0 100 i
* i200.1.1.0        10.1.13.1                0    100      0 100 i
*>                  10.1.12.1                0             0 100 i
```

On R3:

```
R3#Show ip bgp | b Network
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24     10.1.12.1                0    100      0 100 i
*>                  10.1.13.1                0             0 100 i
* i200.1.1.0        10.1.12.1                0    100      0 100 i
*>                  10.1.13.1                0             0 100 i
```

On R4:

```
R4#Show ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24     10.1.13.1                0    100      0 100 i
*>i                 10.1.12.1                0    100      0 100 i
* i200.1.1.0        10.1.13.1                0    100      0 100 i
*>i                 10.1.12.1                0    100      0 100 i
```

Let's configure R2 first:

On R2:

```
R2(config)#router bgp 200
R2(config-router)#neighbor 3.3.3.3 next-hop-self
R2(config-router)#neighbor 4.4.4.4 next-hop-self
```

On R3:

```
R3(config)#router bgp 200
R3(config-router)#neighbor 2.2.2.2 next-hop-self
R3(config-router)#neighbor 4.4.4.4 next-hop-self
```

To verify the configuration:

On R2:

```
R2#Show ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24     3.3.3.3                  0    100      0 100 i
*>                  10.1.12.1                0             0 100 i
* i200.1.1.0        3.3.3.3                  0    100      0 100 i
*>                  10.1.12.1                0             0 100 i
```

On R3:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24     2.2.2.2                  0    100      0 100 i
*>                  10.1.13.1                0             0 100 i
* i200.1.1.0        2.2.2.2                  0    100      0 100 i
*>                  10.1.13.1                0             0 100 i
```

On R4:

```
R4#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24     3.3.3.3                  0    100      0 100 i
*>i                 2.2.2.2                  0    100      0 100 i
* i200.1.1.0        3.3.3.3                  0    100      0 100 i
*>i                 2.2.2.2                  0    100      0 100 i
```

Task 4

Using an extended community attribute, configure R4 such that if a "**show ip bgp | b Network**" command is entered, it matches the output of the following:

On R4:

```
R4#show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i100.1.1.0/24     3.3.3.3                  0    100      0 100 i
* i                 2.2.2.2                  0    100      0 100 i
* i200.1.1.0        3.3.3.3                  0    100      0 100 i
*>i                 2.2.2.2                  0    100      0 100 i
```

In this scenario, the cost extended community attribute is used.

- **Cost is an extended community attribute**
- **It is a non-transitive extended community attribute that allows you to customize the local route preference, which can influence the best path selection process.**
- **This attribute is applied by configuring the "set ext community cost" command, using a route-map. This command is configured with a cost community id (0-255) and a cost value (0-4.29**

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Billion) with a default cost value of 2.145 billion. The lower value has more preference, but the lower cost community id value is used as the tiebreaker.

On R4:

The following identifies the prefix:

```
R4(config)#access-list 1 permit 100.1.1.0 0.0.0.255
```

A route-map matches the access-list and applies the extcommunity cost attribute with two numbers, the first number is the community id and the second number is the community value.

```
R4(config)#route-map TST permit 10
R4(config-route-map)#match ip address 1
R4(config-route-map)#Set extcommunity cost 1 1
R4(config-route-map)#route-map TST per 90
```

Lastly, it's applied by the neighbor command:

```
R4(config-route-map)#router bgp 200
R4(config-router)#neighbor 3.3.3.3 route-map TST in

R4#Clear ip bgp * in
```

To verify the configuration:

On R4:

```
R4#Show ip bgp | b Network
```

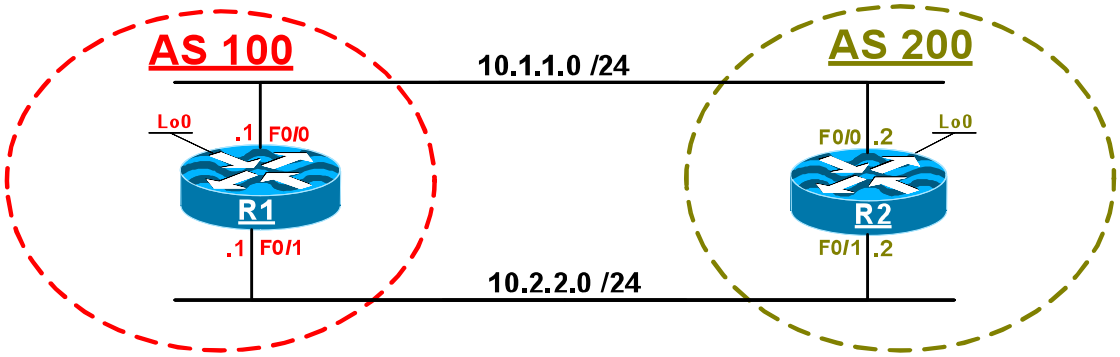
Network	Next Hop	Metric	LocPrf	Weight	Path
*>i100.1.1.0/24	3.3.3.3	0	100	0	100 i
* i	2.2.2.2	0	100	0	100 i
* i200.1.1.0	3.3.3.3	0	100	0	100 i
*>i	2.2.2.2	0	100	0	100 i

Task 5

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 8 – BGP & Load Balancing-I



Lab Setup:

- Configure F0/0 interface of R1 and R2 in VLAN 100
- Configure F0/1 interface of R1 and R2 in VLAN 200
- Use the IP addressing chart for IP address assignment

IP addressing:

Router	Interface / IP address
R1	F0/0 = 10.1.1.1 /24 F0/1 = 10.2.2.1 /24 Loopback 0 = 1.1.1.1 /8
R2	**F0/0 = 10.1.1.2 /24** **F0/1 = 10.2.2.2 /24** **Loopback0 = 2.2.2.2 /8**

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Task 1

Configure an EBGP peer session between R1 and R2; ensure that these routers perform load balancing using the two links. Use an IGP of your choice.

In this topology since the routers are directly connected, the load balancing can be performed if the EBGP peer session is established based on the loopback interfaces of these routers, up to 6 equal cost paths can be used. In this task, RIPv2 was chosen as the IGP.

On R1:

```
R1(config)#router bgp 100
R1(config-router)#No auto-summary
R1(config-router)#neighbor 2.2.2.2 remote-as 200
R1(config-router)#neighbor 2.2.2.2 ebgp-multihop 2
R1(config-router)#neighbor 2.2.2.2 update-source lo0

R1(config)#router rip
R1(config-router)#No auto-summary
R1(config-router)#ver 2
R1(config-router)#network 10.0.0.0
R1(config-router)#network 1.0.0.0
```

On R2:

```
R2(config)#router bgp 200
R2(config-router)#No auto-summary
R2(config-router)#neighbor 1.1.1.1 remote-as 100
R2(config-router)#neighbor 1.1.1.1 update-source lo0
R2(config-router)#neighbor 1.1.1.1 ebgp-multihop 2

R2(config)#router rip
R2(config-router)#no auto-summary
R2(config-router)#ver 2
R2(config-router)#network 10.0.0.0
R2(config-router)#network 2.0.0.0
```

To verify the configuration:

On R2:

```
R2#Sh ip bgp summary | b Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
1.1.1.1         4   100       3       3        1    0    0 00:00:46        0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2#Show ip route rip

R    1.0.0.0/8 [120/1] via 10.2.2.1, 00:00:05, FastEthernet0/1
               [120/1] via 10.1.1.1, 00:00:18, FastEthernet0/0
```

On R1:

```
R1#Sh ip bgp summary | b Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
2.2.2.2         4   200       7       7        1    0    0 00:02:46        0
```

On R1:

```
R1#Show ip route rip

R    2.0.0.0/8 [120/1] via 10.2.2.2, 00:00:21, FastEthernet0/1
               [120/1] via 10.1.1.2, 00:00:12, FastEthernet0/0
```

To test the configuration:

On R1:

```
R1#Traceroute 2.2.2.2

Type escape sequence to abort.
Tracing the route to 2.2.2.2

  1 10.1.1.2 0 msec
    10.2.2.2 0 msec *
```

On R2:

```
R2#Traceroute 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 10.2.2.1 4 msec
    10.1.1.1 0 msec *
```

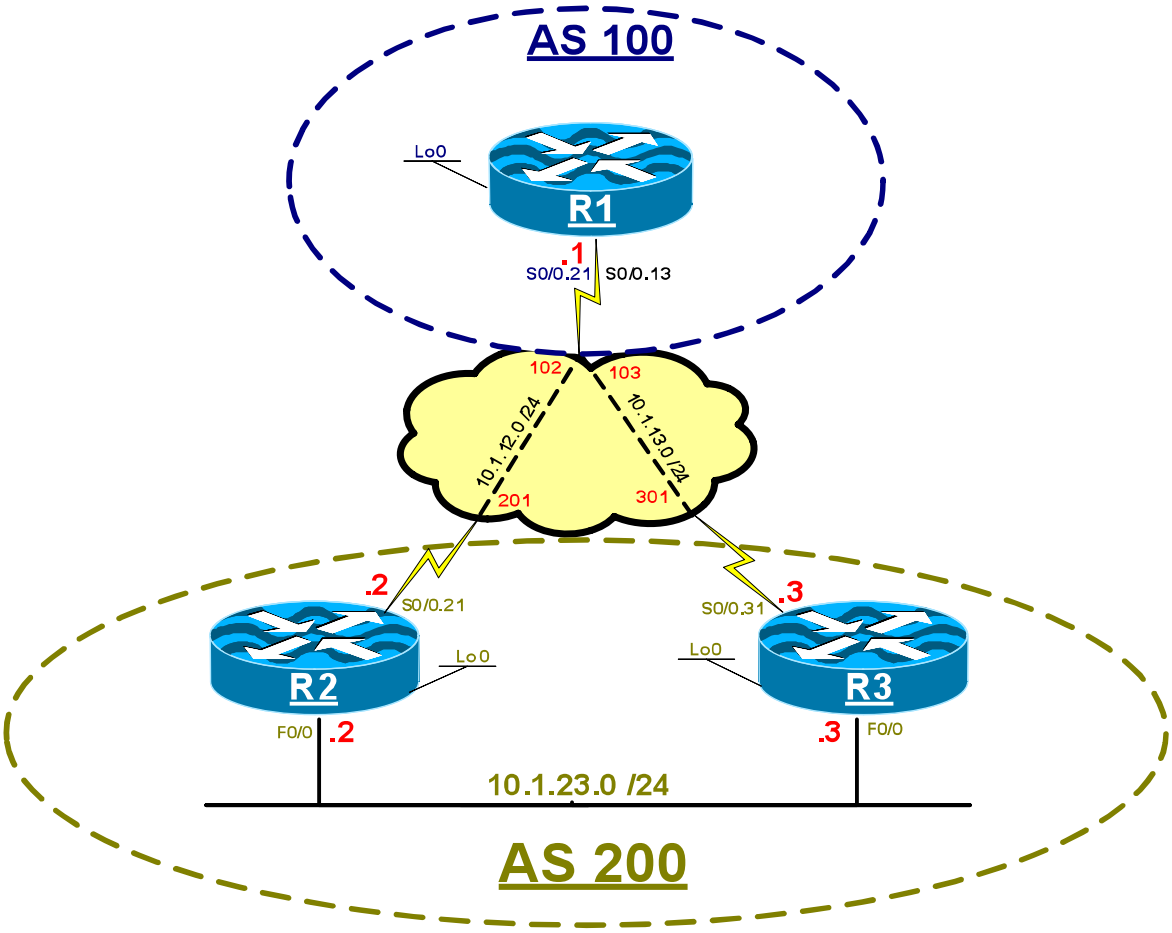

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 2

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 9 – BGP & Load Balancing-II



Lab Setup:

- Configure R1 as the hub and R2 and R3 as the spokes, all frame-relay links should be configured in a point-to-point manner.
- Configure F0/0 interface of R2 and R3 in VLAN 23.
- Use the IP addressing chart for IP addressing assignment.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP addressing:

Router	Interface / IP address
R1	S0/0.12 = 10.1.12.1 /24 S0/0.13 = 10.1.13.1 /24 Loopback 0 = 1.1.1.1 /8
R2	**S0/0.21 = 10.1.12.2 /24** **F0/0 = 10.1.23.2 /24** **Loopback0 = 2.2.2.2 /8**
R3	S0/0.31 = 10.1.13.3 /24 F0/0 = 10.1.23.3 /24 Loopback0 = 3.3.3.3 /8

Task 1

Configure R1 in AS 100 to establish EBGP peer sessions with R2 and R3 in AS 300. R2 and R3 should advertise netw 10.1.23.0 /24 in BGP.

On R1:

```
R1(config-router)#router bgp 100
R1(config-router)#no auto-summary
R1(config-router)#neighbor 10.1.12.2 remote-as 200
R1(config-router)#neighbor 10.1.13.3 remote-as 200
```

On R2:

```
R2(config)#router bgp 200
R2(config-router)#no auto-summary
R2(config-router)#Network 10.1.23.0 mask 255.255.255.0
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#neighbor 10.1.23.3 remote-as 200
```

On R3:

```
R3(config)#router bgp 200
R3(config-router)#no auto-summary
R3(config-router)#network 10.1.23.0 mask 255.255.255.0
R3(config-router)#neighbor 10.1.13.1 remote-as 100
R3(config-router)#neighbor 10.1.23.2 remote-as 200
```

To verify the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1:

```
R1#Sh ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*  10.1.23.0/24     10.1.13.3                0             0 200 i
*>                  10.1.12.2                0             0 200 i

R1#Show ip route | b Gateway
Gateway of last resort is not set

C    1.0.0.0/8 is directly connected, Loopback0
     10.0.0.0/24 is subnetted, 3 subnets
C       10.1.13.0 is directly connected, Serial0/0.13
C       10.1.12.0 is directly connected, Serial0/0.12
B       10.1.23.0 [20/0] via 10.1.12.2, 00:04:58
```

Task 2

Configure R1 such that it uses both neighbors (R2 and R3) to perform an equal cost load balancing.

BGP will only use one path to a given destination; therefore, it does not perform load balancing amongst multiple equal cost paths. The "maximum-paths" command can be configured to change this behavior. By default this command is set to 1:

On R1:

```
R1#Show ip protocol | inc Routing Protocol|Max

Routing Protocol is "bgp 100"
  Maximum path: 1
```

Let's set the maximum-paths to 2:

On R1:

```
R1(config)#router bgp 100
R1(config-router)#maximum-paths 2
```

To verify the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1:

```
R1#Show ip protocol | Routing Protocol|Max

Routing Protocol is "bgp 100"
  Maximum path: 2

R1#Show ip route | b Gateway
Gateway of last resort is not set

C    1.0.0.0/8 is directly connected, Loopback0
     10.0.0.0/24 is subnetted, 3 subnets
C       10.1.13.0 is directly connected, Serial0/0.13
C       10.1.12.0 is directly connected, Serial0/0.12
B       10.1.23.0 [20/0] via 10.1.13.3, 00:01:33
                  [20/0] via 10.1.12.2, 00:11:22
```

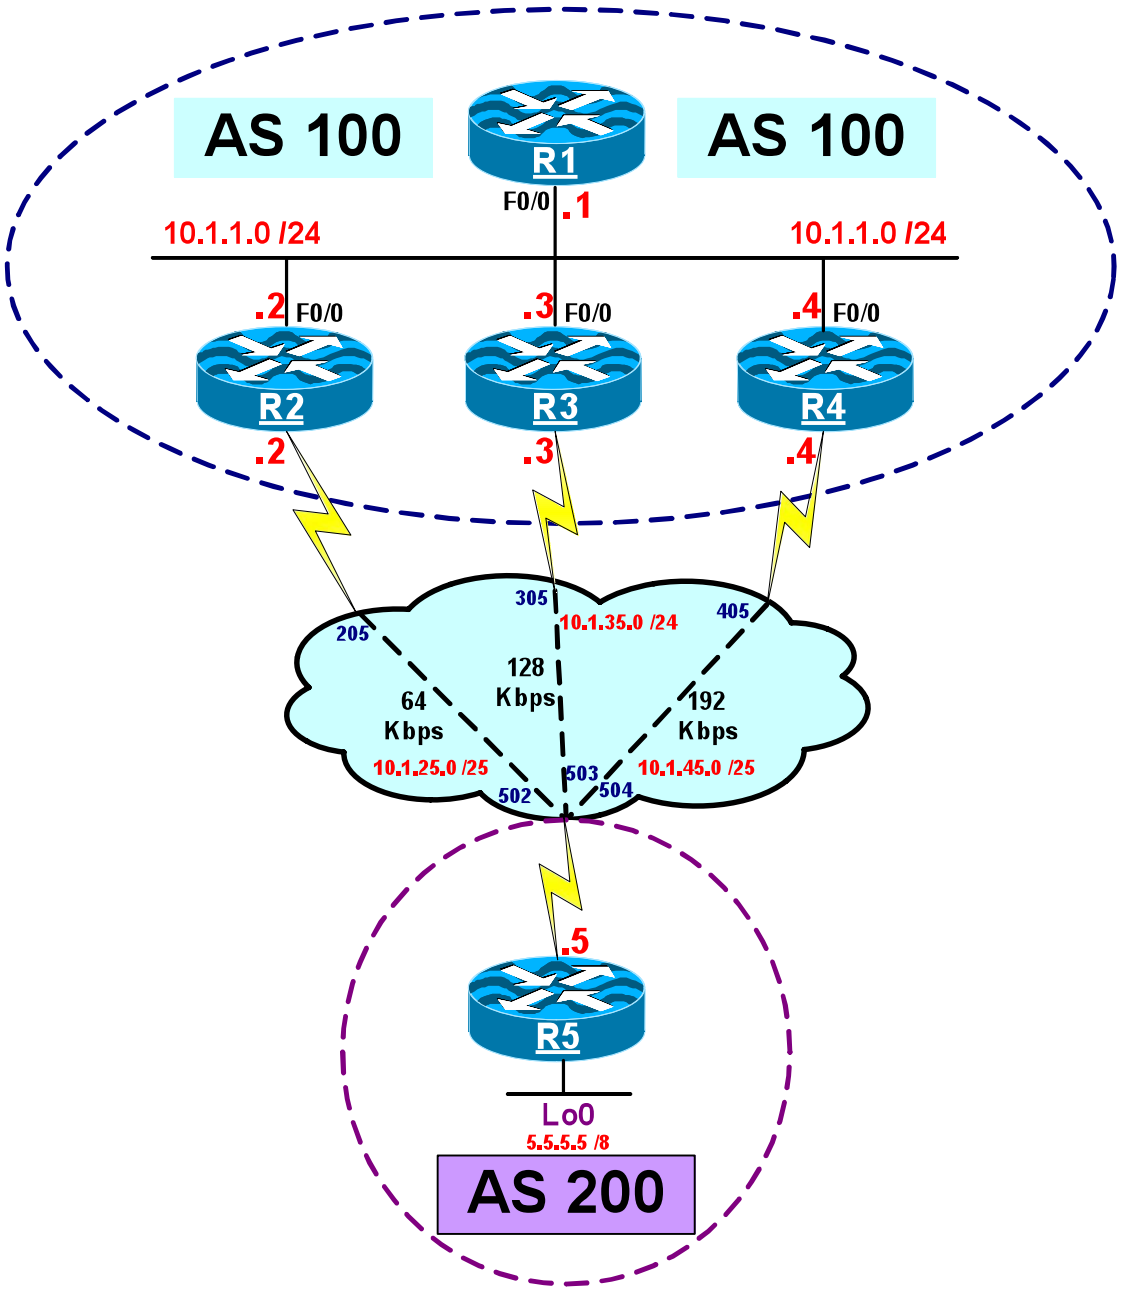
Note: R1 is performing equal cost load balancing across the two links.

Task 3

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 10 – BGP Unequal-Cost Load Balancing



©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Lab Setup:

- Configure F0/0 interface of R1, R2, R3 and R4 should be configured in VLAN 100.
- Configure the frame-relay connections in a point-to-point manner.
- Use the IP addressing chart below for IP address assignment.

IP addressing:

Router	Interface / IP address	VLAN
R1	F0/0 = 10.1.1.1 /24	100
R2	**F0/0 = 10.1.1.2 /24** **S0/0.25 = 10.1.25.2 /24**	**100**
R3	S0/0.35 = 10.1.35.3 /24 F0/0 = 10.1.1.3 /24	 100
R4	**F0/0 = 10.1.1.4 /24** **S0/0.45 = 10.1.45.4 /24**	**100**
R5	S0/0.54 = 10.1.45.5 /24 S0/0.53 = 10.1.35.5 /24 S0/0.52 = 10.1.25.5 /24 Lo0 = 5.5.5.5 /8	

Task 1

Configure peering according to the diagram.

On R1:

```
R1(config)#router bgp 100
R1(config-router)#No au
R1(config-router)#neighbor 10.1.1.2 remote 100
R1(config-router)#neighbor 10.1.1.3 remote 100
R1(config-router)#neighbor 10.1.1.4 remote 100
```

On R2:

```
R2(config)#router bgp 100
R2(config-router)#No au
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config-router)#neighbor 10.1.1.1 remote 100
R2(config-router)#neighbor 10.1.1.3 remote 100
R2(config-router)#neighbor 10.1.1.4 remote 100
R2(config-router)#neighbor 10.1.25.5 remote 200
```

On R3:

```
R3(config)#router bgp 100
R3(config-router)#No au
R3(config-router)#neighbor 10.1.1.1 remote 100
R3(config-router)#neighbor 10.1.1.2 remote 100
R3(config-router)#neighbor 10.1.1.4 remote 100
R3(config-router)#neighbor 10.1.35.5 remote 200
```

On R4:

```
R4(config)#router bgp 100
R4(config-router)#No au
R4(config-router)#neighbor 10.1.1.1 remote 100
R4(config-router)#neighbor 10.1.1.2 remote 100
R4(config-router)#neighbor 10.1.1.3 remote 100
R4(config-router)#neighbor 10.1.45.5 remote 200
```

On R5:

```
R5(config)#router bgp 200
R5(config-router)#No au
R5(config-router)#neighbor 10.1.25.2 remote 100
R5(config-router)#neighbor 10.1.35.3 remote 100
R5(config-router)#neighbor 10.1.45.4 remote 100
R5(config-router)#network 5.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
* i5.0.0.0         10.1.45.5                0    100      0 200 i
* i                10.1.35.5                0    100      0 200 i
* i                10.1.25.5                0    100      0 200 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 2

Configure the border routers to change the next hop IP address for external prefixes to their IP address.

On R2:, R3 and R4

```
R2(config)#router bgp 100
R2(config-router)#neighbor 10.1.1.1 next-hop-self
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
* i5.0.0.0         10.1.1.4                 0    100      0 200 i
* i                10.1.1.3                 0    100      0 200 i
*>i                10.1.1.2                 0    100      0 200 i
```

Since the next-hop IP address to reach network 5.0.0.0 /8 is reachable by the local router, the route is considered valid (The ">" sign to the left of the "i" in the network column. Whereas, in the previous task, the output of the show command revealed that the next-hop IP address was not reachable and therefore, the route was not valid. If the route is not valid from BGP's perspective, it will not be injected into the routing table.

Task 3

Configure the routers in AS 100 such that R1 distributes traffic proportionally over the external links to reach network 5.0.0.0 /8, the load balancing should be done based on the bandwidth of the links between the border routers of this AS and AS 200.

The unequal cost load balancing feature is used in conjunction with BGP multipath feature to advertise the exit link's bandwidth as an extended community to IBGP peers, this feature is used for links between directly connected EBGP neighbors and available in IOS release 12.2(2T) or better.
To configure this feature, the following steps should be performed:

1. **Enable the BGP dmzlink-bw feature:**
   This is accomplished by configuring the "BGP dmzlink-bw" router configuration command, this must be configured on the border router/s and the internal router/s

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

2. **Configure BGP to include the link bandwidth value of the external interface in extended community so they can be propagated to IBGP peers. This is accomplished through the "neighbor dmzlink-bw" router configuration mode command.**

Remember, for this feature to work, R1 must have an equal IGP cost and BGP attributes or else this feature will not work. Note the bandwidth of the links connecting the border routers R2, R3 and R4 to AS 200 is set to 64K, 128K and 192K respectively.

On R1:

```
R1#Show ip bgp 5.0.0.0

BGP routing table entry for 5.0.0.0/8, version 3
Paths: (3 available, best #3, table Default-IP-Routing-Table)
  Not advertised to any peer
  200
    10.1.1.4 from 10.1.1.4 (10.1.45.4)
      Origin IGP, metric 0, localpref 100, valid, internal
  200
    10.1.1.3 from 10.1.1.3 (10.1.35.3)
      Origin IGP, metric 0, localpref 100, valid, internal
  200
    10.1.1.2 from 10.1.1.2 (10.1.25.2)
      Origin IGP, metric 0, localpref 100, valid, internal, best

R1#Show ip route 5.0.0.0

Routing entry for 5.0.0.0/8
  Known via "bgp 100", distance 200, metric 0
  Tag 200, type internal
  Last update from 10.1.1.2 00:11:30 ago
  Routing Descriptor Blocks:
  * 10.1.1.2, from 10.1.1.2, 00:11:30 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 200
```

The BGP table identifies only one of the routes as the best, in this case, since all the attributes are equal and the "BGP Best path compare-routerid" command is not configured, the neighbor with the lowest IP address was chosen.

To configure the task:

On R1:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The following allows the local router to have three equal IBGP cost paths:

```
R1(config)#router bgp 100
R1(config-router)#maximum-path ibgp 3
```

On R2:, R3 and R4

```
Rx(config)#router bgp 100
Rx(config-router)#neighbor 10.1.1.1 send-community extended
Rx(config-router)#bgp dmzlink-bw
```

On R2:

```
R2(config)#router bgp 100
R2(config-router)#neighbor 10.1.25.5 dmzlink-bw
```

On R3:

```
R3(config)#router bgp 100
R3(config-router)#neighbor 10.1.35.5 dmzlink-bw
```

On R4:

```
R4(config)#router bgp 100
R4(config-router)#neighbor 10.1.45.5 dmzlink-bw
```

On R1:

```
R1#Clear ip bgp * in
```

To verify the configuration:

On R1:

```
R1#Show ip bgp 5.0.0.0

BGP routing table entry for 5.0.0.0/8, version 6
Paths: (3 available, best #3, table Default-IP-Routing-Table)
Multipath: iBGP
Flag: 0x800
  Not advertised to any peer
  200
    10.1.1.4 from 10.1.1.4 (10.1.45.4)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      DMZ-Link Bw 24 kbytes
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
  200
    10.1.1.3 from 10.1.1.3 (10.1.35.3)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      DMZ-Link Bw 16 kbytes
  200
    10.1.1.2 from 10.1.1.2 (10.1.25.2)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath, best
      DMZ-Link Bw 8 kbytes

R1#Show ip route 5.0.0.0

Routing entry for 5.0.0.0/8
  Known via "bgp 100", distance 200, metric 0
  Tag 200, type internal
  Last update from 10.1.1.3 00:07:18 ago
  Routing Descriptor Blocks:
    10.1.1.4, from 10.1.1.4, 00:07:18 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 200
    10.1.1.3, from 10.1.1.3, 00:07:18 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 200
  * 10.1.1.2, from 10.1.1.2, 00:21:13 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 200
```

Note: The traffic count is 1, in order to have the BGP table reflect on this counter, the "BGP dmzlink-bw" must be configured on R1.

On R1:

```
R1(config)#router bgp 100
R1(config-router)#bgp dmzlink-bw

R1#Clear ip bgp * in
```

To verify the configuration:

On R1:

```
R1#Show ip route 5.0.0.0

Routing entry for 5.0.0.0/8
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Known via "bgp 100", distance 200, metric 0
Tag 200, type internal
Last update from 10.1.1.2 00:00:43 ago
Routing Descriptor Blocks:
  10.1.1.4, from 10.1.1.4, 00:00:43 ago
    Route metric is 0, traffic share count is 2
    AS Hops 1
    Route tag 200
* 10.1.1.3, from 10.1.1.3, 00:00:43 ago
    Route metric is 0, traffic share count is 1
    AS Hops 1
    Route tag 200
  10.1.1.2, from 10.1.1.2, 00:00:43 ago
    Route metric is 0, traffic share count is 24
    AS Hops 1
    Route tag 200
```

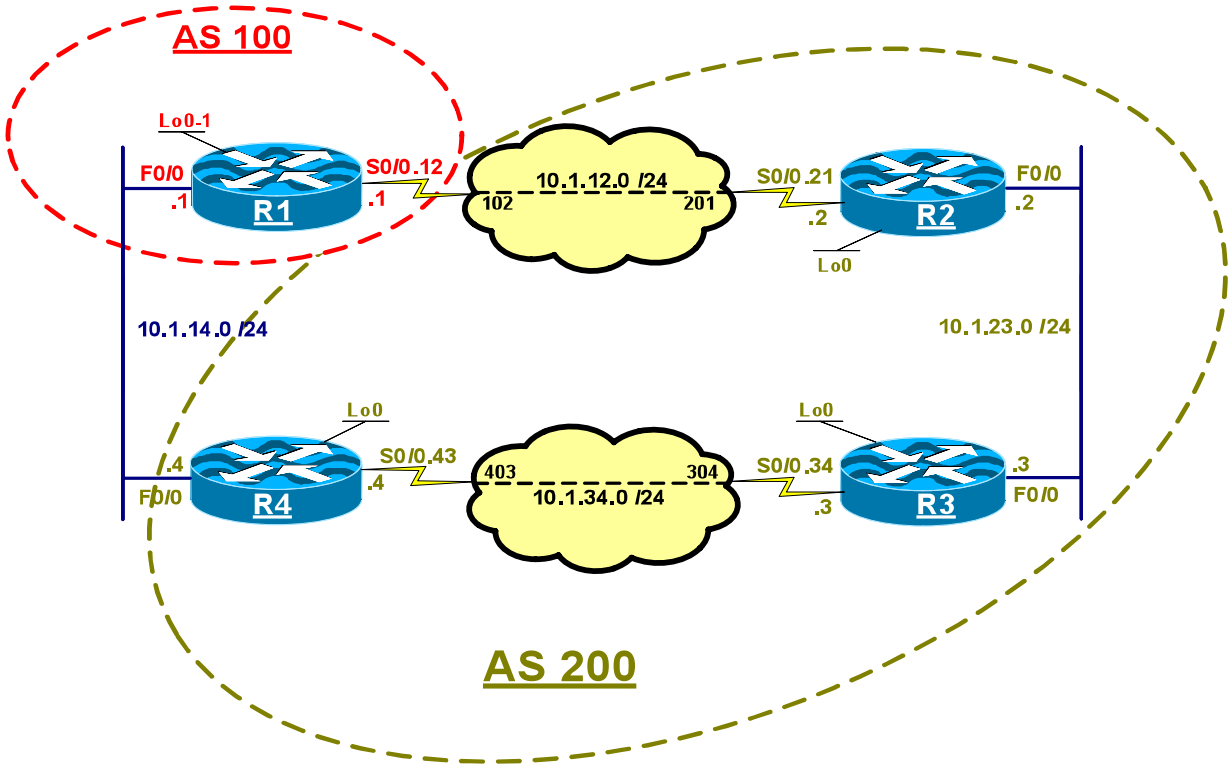
Note: The feature does not work unless it is enabled. In some IOS releases especially a "T" code, you may have to save the configuration and reload R1 to see the results.

Task 4

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 11 - Local-Preference -I



Lab Setup:

- Configure the routers that are connected to the frame-relay clouds in a point-to-point manner.
- R1 and R4's F0/0 interface should be configured in VLAN 14.
- R2 and R3's F0/0 interface should be configured in VLAN 23.
- Use the following IP addressing chart for IP address assignment.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP addressing:

Router	Interface	IP Address	AS number
R1	Lo0 Lo1 S0/0.12 F0/0	1.1.1.1 /8 11.1.1.1 /8 10.1.12.1 /24 10.1.14.1 /24	100
R2	**Lo0** **S0/0.21** **F0/0**	**2.2.2.2 /8** **10.1.12.2 /24** **10.1.23.2 /24**	**200**
R3	Lo0 F0/0 S0/0.34	3.3.3.3 /8 10.1.23.3 /24 10.1.34.3 /24	200
R4	**Lo0** **F0/0** **S0/0.43**	**4.4.4.4 /8** **10.1.14.4 /24** **10.1.34.4 /24**	**200**

Task 1

Configure routers R2, R3 and R4 in AS 200, these routers should have full mesh peer session between them. R2 and R4 should have an EBGP peer session to R1 in AS 100. BGP routers should **only** advertise their loopback interface/s in BGP. Provide NLRI for the links using RIPv2, disable automatic summarization.

On R1:

```
R1(config-if)#router bgp 100
R1(config-router)#No au
R1(config-router)#neighbor 10.1.14.4 remote-as 200
R1(config-router)#neighbor 10.1.12.2 remote-as 200
R1(config-router)#netw 11.0.0.0
R1(config-router)#netw 1.0.0.0

R1(config)#router rip
R1(config-router)#No au
R1(config-router)#ver 2
R1(config-router)#network 10.0.0.0
```

On R2:

```
R2(config)#router bgp 200
R2(config-router)#No au
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config-router)#netw 2.0.0.0
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#neighbor 10.1.23.3 remote-as 200
R2(config-router)#neighbor 10.1.34.4 remote-as 200

R2(config)#router rip
R2(config-router)#No au
R2(config-router)#ver 2
R2(config-router)#netw 10.0.0.0
```

On R3:

```
R3(config)#router bgp 200
R3(config-router)#No au
R3(config-router)#netw 3.0.0.0
R3(config-router)#neighbor 10.1.34.4 remote-as 200
R3(config-router)#neighbor 10.1.23.2 remote-as 200

R3(config)#router rip
R3(config-router)#No au
R3(config-router)#ver 2
R3(config-router)#netw 10.0.0.0
```

On R4:

```
R4(config)#router bgp 200
R4(config-router)#No au
R4(config-router)#netw 4.0.0.0
R4(config-router)#neighbor 10.1.34.3 remote-as 200
R4(config-router)#neighbor 10.1.23.2 remote-as 200
R4(config-router)#neighbor 10.1.14.1 remote-as 100

R4(config)#router rip
R4(config-router)#No au
R4(config-router)#ver 2
R4(config-router)#netw 10.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
*  2.0.0.0          10.1.14.4                              0 200 i
*>                  10.1.12.2                0             0 200 i
*  3.0.0.0          10.1.14.4                              0 200 i
*>                  10.1.12.2                              0 200 i
*  4.0.0.0          10.1.12.2                              0 200 i
*>                  10.1.14.4                0             0 200 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
```

Task 2

Ensure that the routers in AS 200 use R4 to reach network 1.0.0.0 /8 in AS 100. Local-Pref attribute must be used to accomplish this task.

Before this attribute is configured, the existing BGP table of the routers in AS 200 should be examined:

On R2:

```
R2#Show ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
* i1.0.0.0         10.1.14.1                0    100      0 100 i
*>                 10.1.12.1                0             0 100 i
*> 2.0.0.0         0.0.0.0                  0         32768 i
*>i3.0.0.0         10.1.23.3                0    100      0 i
*>i4.0.0.0         10.1.34.4                0    100      0 i
* i11.0.0.0        10.1.14.1                0    100      0 100 i
*>                 10.1.12.1                0             0 100 i
```

On R3:

```
R3#Show ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
* i1.0.0.0         10.1.14.1                0    100      0 100 i
*>i                10.1.12.1                0    100      0 100 i
*>i2.0.0.0         10.1.23.2                0    100      0 i
*> 3.0.0.0         0.0.0.0                  0         32768 i
*>i4.0.0.0         10.1.34.4                0    100      0 i
* i11.0.0.0        10.1.14.1                0    100      0 100 i
*>i                10.1.12.1                0    100      0 100 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4:

```
R4#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
* i1.0.0.0         10.1.12.1                0    100      0 100 i
*>                 10.1.14.1                0             0 100 i
*>i2.0.0.0         10.1.23.2                0    100      0 i
*>i3.0.0.0         10.1.34.3                0    100      0 i
*> 4.0.0.0         0.0.0.0                  0         32768 i
* i11.0.0.0        10.1.12.1                0    100      0 100 i
*>                 10.1.14.1                0             0 100 i
```

Note: R2 and R3 are taking the R2-R1 link to connect to network 1.0.0.0 /8, whereas, R4 is taking the R4-R1 link. To configure the task:

On R4:

```
R4(config)#access-list 1 permit 1.0.0.0 0.255.255.255

R4(config)#route-map TST permit 10
R4(config-route-map)#match ip addr 1
R4(config-route-map)#set Local-preference 101
R4(config)#route-map TST permit 90

R4(config-route-map)#router bgp 200
R4(config-router)#neighbor 10.1.14.1 route-map TST in

R4#Clear ip bgp * in
```

The local preference attribute is used to prefer an exit point from the local AS. Unlike the weight attribute, the local preference attribute is propagated throughout the local AS, hence an AS-Wide attribute. If there are multiple exit points from the local AS, the local preference attribute is used to select the exit point for a specific or all routes. Since the local preference attribute affects the routers within the AS, the route-map should be configured in the "in" direction. Remember that with local preference the higher value has better preference. When R4 and R2 in AS 200 receive an update for network 1.0.0.0 /8 from R1 in AS 100, they set the local-preference to 100 and share this information with all of their IBGP peer/s. In the above configuration, R4 has changed the local preference to 101, which is higher than the default value of 100, and since R4 will share this local preference value with all of its IBGP peers, they will all take the requested path to reach network 1.0.0.0.

To verify the configuration:

On R2:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.14.1                0    101      0 100 i
*                   10.1.12.1                0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          10.1.23.3                0    100      0 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
* i11.0.0.0         10.1.14.1                0    100      0 100 i
*>                  10.1.12.1                0             0 100 i
```

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.14.1                0    101      0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
* i11.0.0.0         10.1.14.1                0    100      0 100 i
*>i                 10.1.12.1                0    100      0 100 i
```

On R4:

```
R4#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.14.1                0    101      0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*>i3.0.0.0          10.1.34.3                0    100      0 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
* i11.0.0.0         10.1.12.1                0    100      0 100 i
*>                  10.1.14.1                0             0 100 i
```

Note: The routers in AS 200 take the R4-R1 link to connect to network 1.0.0.0/8.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 3

Ensure that the routers in AS 200 use R2 to reach network 11.0.0.0 /8. Local-Pref attribute must be used to accomplish this task.

Before this attribute is configured, the existing BGP table of the routers in AS 200 should be examined:

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.14.1                0    101      0 100 i
*                   10.1.12.1                0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          10.1.23.3                0    100      0 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
* i11.0.0.0         10.1.14.1                0    100      0 100 i
*>                  10.1.12.1                0             0 100 i
```

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.14.1                0    101      0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
* i11.0.0.0         10.1.14.1                0    100      0 100 i
*>i                 10.1.12.1                0    100      0 100 i
```

On R4:

```
R4#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.14.1                0    101      0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*>i3.0.0.0          10.1.34.3                0    100      0 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
* i11.0.0.0         10.1.12.1                0    100      0 100 i
*>                  10.1.14.1                0             0 100 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Note: R4 is taking the R4-R1 link to connect to network 11.0.0.0, whereas, the other routers in AS 200 are taking the R1-R2 link. Let's configure R2 for this task:

On R2:

```
R2(config)#access-list 11 permit 11.0.0.0 0.255.255.255

R2(config)#route-map TST permit 10
R2(config-route-map)#match ip addr 11
R2(config-route-map)#set Local-preference 101
R2(config)#route-map TST permit 90

R2(config-route-map)#router bgp 200
R2(config-router)#neighbor 10.1.12.1 route-map TST in

R2#Clear ip bgp * in
```

To verify the configuration:

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.14.1                0    101      0 100 i
*                   10.1.12.1                0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          10.1.23.3                0    100      0 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
*> 11.0.0.0         10.1.12.1                0    101      0 100 i
```

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.14.1                0    101      0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
*>i11.0.0.0         10.1.12.1                0    101      0 100 i
```

On R4:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.14.1                0    101      0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*>i3.0.0.0          10.1.34.3                0    100      0 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*>i11.0.0.0         10.1.12.1                0    101      0 100 i
*                   10.1.14.1                0             0 100 i
```

Note: The routers in AS 200 connect to network 11.0.0.0 /8 by going through R2-R1 link.

Let's test this configuration:

On R4:

```
R4#Traceroute 11.1.1.1

Type escape sequence to abort.
Tracing the route to 11.1.1.1

  1 10.1.14.1 0 msec *  0 msec
```

Why is R4 going directly to R1 for network 11.0.0.0 /8 when the BGP table uses 10.1.12.1?

Let's check the routing table of R4:

On R4:

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

B    1.0.0.0/8 [20/0] via 10.1.14.1, 00:08:59
B    2.0.0.0/8 [200/0] via 10.1.23.2, 00:08:59
B    3.0.0.0/8 [200/0] via 10.1.34.3, 00:08:59
C    4.0.0.0/8 is directly connected, Loopback0
     10.0.0.0/24 is subnetted, 4 subnets
C       10.1.14.0 is directly connected, FastEthernet0/0
R       10.1.12.0 [120/1] via 10.1.14.1, 00:00:13, FastEthernet0/0
R       10.1.23.0 [120/1] via 10.1.34.3, 00:00:06, Serial0/0.43
C       10.1.34.0 is directly connected, Serial0/0.43
B    11.0.0.0/8 [200/0] via 10.1.12.1, 00:02:49
```

Note: For the local router to reach 11.0.0.0 /24 network, it has to go to 10.1.12.1. This is because BGP border router (R2 in this case), did not change the "next-hop" attribute when it advertised this

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

network to its IBGP peers.

To fix this problem, the border routers (R2 and R4) must be configured to change the next-hop IP address to be an IP address from their local AS.

On R4:

```
R4(config)#router bgp 200
R4(config-router)#neigh 10.1.34.3 next-hop-self
R4(config-router)#neigh 10.1.23.2 next-hop-self
```

On R2:

```
R2(config)#router bgp 200
R2(config-router)#neigh 10.1.23.3 next-hop-self
R2(config-router)#neigh 10.1.34.4 next-hop-self
```

To verify the configuration:

On R4:

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

B    1.0.0.0/8 [20/0] via 10.1.14.1, 00:19:05
B    2.0.0.0/8 [200/0] via 10.1.23.2, 00:19:05
B    3.0.0.0/8 [200/0] via 10.1.34.3, 00:19:05
C    4.0.0.0/8 is directly connected, Loopback0
     10.0.0.0/24 is subnetted, 4 subnets
C       10.1.14.0 is directly connected, FastEthernet0/0
R       10.1.12.0 [120/1] via 10.1.14.1, 00:00:08, FastEthernet0/0
R       10.1.23.0 [120/1] via 10.1.34.3, 00:00:07, Serial0/0.43
C       10.1.34.0 is directly connected, Serial0/0.43
B    11.0.0.0/8 [200/0] via 10.1.23.2, 00:00:14
```

Note: The next-hop attribute is changed from 10.1.12.1 to 10.1.23.2. Let's try the traceroute again and verify the path to network 11.0.0.0 /24:

```
R4#Traceroute 11.1.1.1

Type escape sequence to abort.

Tracing the route to 11.1.1.1

  1 10.1.34.3 24 msec 24 msec 24 msec
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

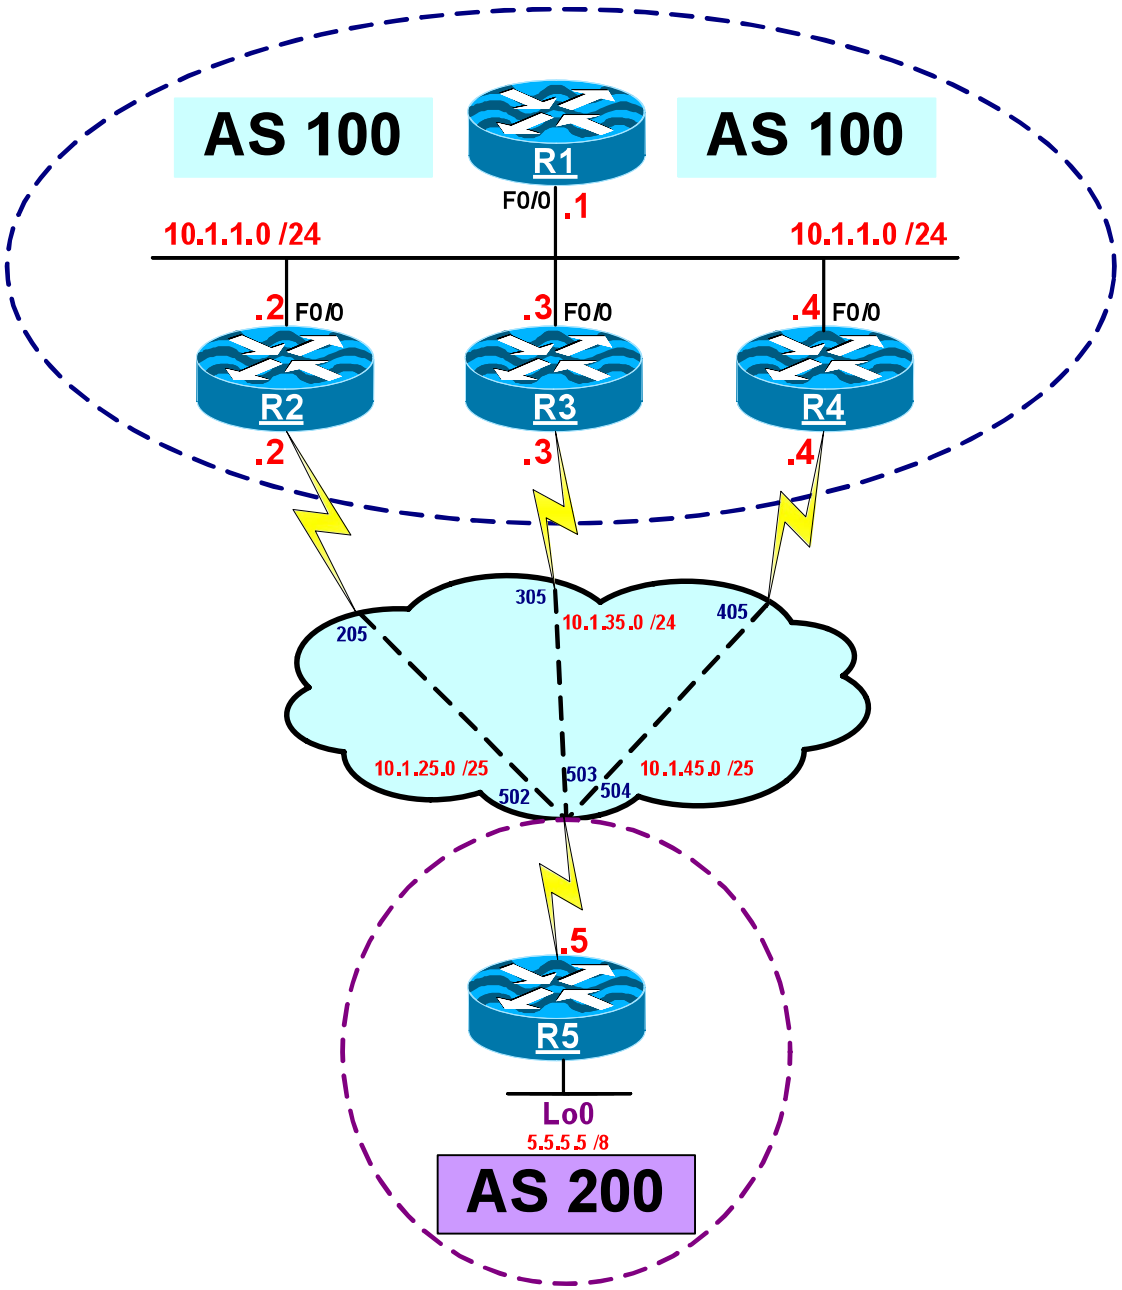
```
2 10.1.23.2 24 msec 24 msec 24 msec
3 10.1.12.1 16 msec *  16 msec
```

Task 4

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 12 – BGP Local-Preference - II



©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Lab Setup:

- Configure F0/0 interface of R1, R2, R3 and R4 should be configured in VLAN 100.
- Configure the frame-relay connections in a point-to-point manner.
- Use the IP addressing chart below for IP address assignment.

IP addressing:

Router	Interface / IP address	VLAN
R1	F0/0 = 10.1.1.1 /24	100
R2	**F0/0 = 10.1.1.2 /24** **S0/0.25 = 10.1.25.2 /24**	**100**
R3	S0/0.35 = 10.1.35.3 /24 F0/0 = 10.1.1.3 /24	 100
R4	**F0/0 = 10.1.1.4 /24** **S0/0.45 = 10.1.45.4 /24**	**100**
R5	S0/0.54 = 10.1.45.5 /24 S0/0.53 = 10.1.35.5 /24 S0/0.52 = 10.1.25.5 /24 Lo0 = 5.5.5.5 /8	

Task 1

Configure peering according to the diagram.

On R1:

```
R1(config)#router bgp 100
R1(config-router)#No au
R1(config-router)#neighbor 10.1.1.2 remote 100
R1(config-router)#neighbor 10.1.1.3 remote 100
R1(config-router)#neighbor 10.1.1.4 remote 100
```

On R2:

```
R2(config)#router bgp 100
R2(config-router)#No au
R2(config-router)#neighbor 10.1.1.1 remote 100
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config-router)#neighbor 10.1.1.3 remote 100
R2(config-router)#neighbor 10.1.1.4 remote 100
R2(config-router)#neighbor 10.1.25.5 remote 200
```

On R3:

```
R3(config)#router bgp 100
R3(config-router)#No au
R3(config-router)#neighbor 10.1.1.1 remote 100
R3(config-router)#neighbor 10.1.1.2 remote 100
R3(config-router)#neighbor 10.1.1.4 remote 100
R3(config-router)#neighbor 10.1.35.5 remote 200
```

On R4:

```
R4(config)#router bgp 100
R4(config-router)#No au
R4(config-router)#neighbor 10.1.1.1 remote 100
R4(config-router)#neighbor 10.1.1.2 remote 100
R4(config-router)#neighbor 10.1.1.3 remote 100
R4(config-router)#neighbor 10.1.45.5 remote 200
```

On R5:

```
R5(config)#router bgp 200
R5(config-router)#No au
R5(config-router)#neighbor 10.1.25.2 remote 100
R5(config-router)#neighbor 10.1.35.3 remote 100
R5(config-router)#neighbor 10.1.45.4 remote 100
R5(config-router)#network 5.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
* i5.0.0.0         10.1.45.5                0    100      0 200 i
* i                10.1.35.5                0    100      0 200 i
* i                10.1.25.5                0    100      0 200 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 2

Configure the border routers to change the next-hop IP address to an internal (the link within their AS) IP address.

On R2:, R3 and R4

```
Rx(config)#router bgp 100
Rx(config-router)#neighbor 10.1.1.1 next-hop-self
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
* i5.0.0.0         10.1.1.4                 0    100      0 200 i
* i                10.1.1.3                 0    100      0 200 i
*>i               10.1.1.2                 0    100      0 200 i
```

Task 3

Configure R2, R3 and R4 such that R1 takes R4 as the primary and R3 as the backup, if R4 and R3 are both down, then, R1 should take R2 to reach network 5.0.0.0 /8. You must use local preference and minimum number of commands to accomplish this task.

On R2:

```
R2(config)#router bgp 100
R2(config-router)#bgp default local-preference 200

R2#Clear ip bgp * in
```

On R3:

```
R3(config)#router bgp 100
R3(config-router)#bgp default local-preference 300

R3#Clear ip bgp * in
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4:

```
R4(config-router)#bgp default local-preference 400

R4#Clear ip bgp * in
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i5.0.0.0          10.1.1.4                 0    400      0 200 i
* i                 10.1.1.3                 0    300      0 200 i
* i                 10.1.1.2                 0    200      0 200 i

R1#Show ip bgp 5.0.0.0

BGP routing table entry for 5.0.0.0/8, version 6
Paths: (3 available, best #1, table Default-IP-Routing-Table)
Flag: 0x840
  Not advertised to any peer
  200
    10.1.1.4 from 10.1.1.4 (10.1.45.4)
      Origin IGP, metric 0, localpref 400, valid, internal, best
  200
    10.1.1.3 from 10.1.1.3 (10.1.35.3)
      Origin IGP, metric 0, localpref 300, valid, internal
  200
    10.1.1.2 from 10.1.1.2 (10.1.25.2)
      Origin IGP, metric 0, localpref 200, valid, internal
```

To test the configuration:

On R1:

```
R1#Traceroute 5.5.5.5

Type escape sequence to abort.
Tracing the route to 5.5.5.5

  1 10.1.1.4 0 msec 0 msec 0 msec
  2 10.1.45.5 24 msec *  56 msec
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 4

Remove the configuration from the previous task and re-configure the same task using another method. **Do not** use neighbor 10.1.25.5, 10.1.35.5 or 10.1.45.5. You should use local preference to accomplish this task.

On R2:

```
R2(config)#router bgp 100
R2(config-router)#no bgp default local-preference 200

R2#Clear ip bgp * in
```

On R3:

```
R3(config)#router bgp 100
R3(config-router)#no bgp default local-preference 300

R3#Clear ip bgp * in
```

On R4:

```
R4(config)#router bgp 100
R4(config-router)#no bgp default local-preference 400

R4#Clear ip bgp * in
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
* i5.0.0.0         10.1.1.4                 0    100      0 200 i
* i                10.1.1.3                 0    100      0 200 i
*>i               10.1.1.2                 0    100      0 200 i
```

To configure the task:

On R2:

```
R2(config)#route-map TST permit 10
R2(config-route-map)#set local-preference 200
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config-route-map)#route-map TST permit 90

R2(config)#router bgp 100
R2(config-router)#neighbor 10.1.1.1 route-map TST OUT

R2#Clear ip bgp * in
```

On R3:

```
R3(config)#route-map TST permit 10
R3(config-route-map)#set local-preference 300
R3(config-route-map)#route-map TST permit 90

R3(config)#router bgp 100
R3(config-router)#neighbor 10.1.1.1 route-map TST OUT

R3#Clear ip bgp * in
```

On R4:

```
R4(config)#route-map TST permit 10
R4(config-route-map)#set local-preference 400
R4(config-route-map)#route-map TST permit 90

R4(config)#router bgp 100
R4(config-router)#neighbor 10.1.1.1 route-map TST OUT

R4#Clear ip bgp * in
```

To test the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i5.0.0.0          10.1.1.4                 0    400      0 200 i
* i                 10.1.1.3                 0    300      0 200 i
* i                 10.1.1.2                 0    200      0 200 i

R1#Traceroute 5.5.5.5

Type escape sequence to abort.

Tracing the route to 5.5.5.5
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

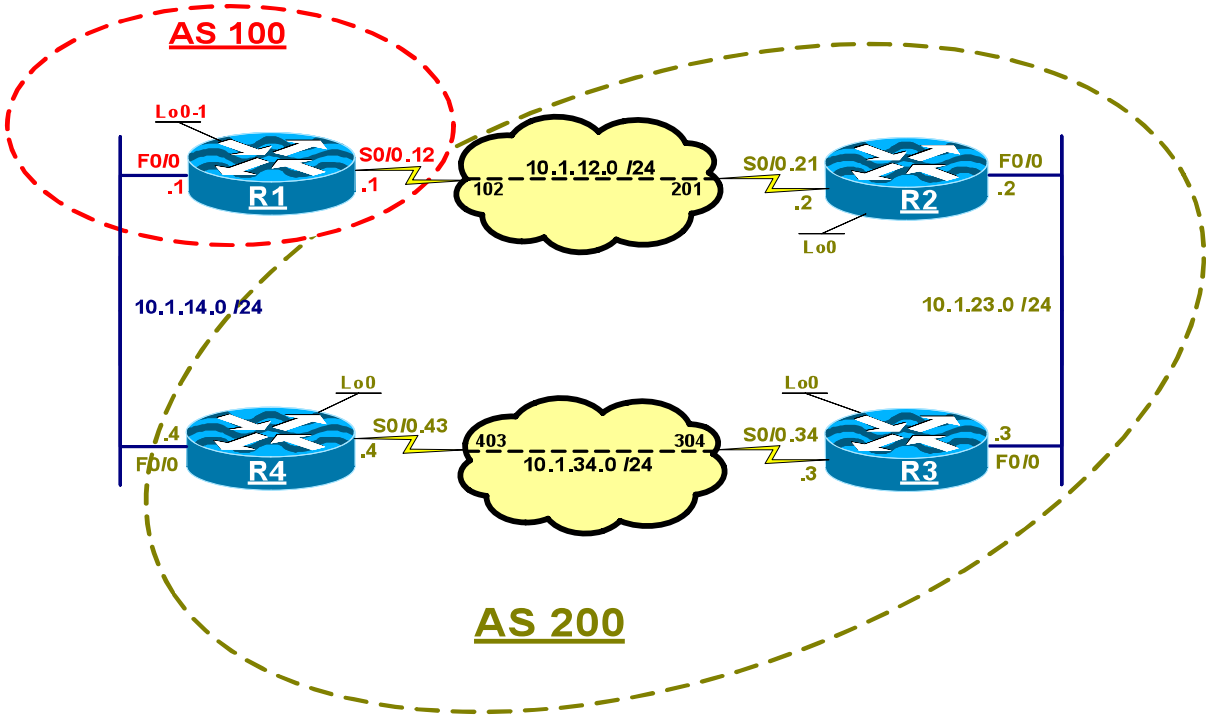
```
1 10.1.1.4 0 msec 0 msec 0 msec
2 10.1.45.5 24 msec *  24 msec
```

Task 5

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 13 - The AS-Path Attribute



Lab Setup:

- Configure the routers that are connected to the frame-relay clouds in a point-to-point manner.
- R1 and R4's F0/0 interface should be configured in VLAN 14.
- R2 and R3's F0/0 interface should be configured in VLAN 23.
- Use the following IP addressing chart for IP address assignment.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP addressing:

Router	Interface	IP Address	AS number
R1	Lo0 Lo1 S0/0.12 F0/0	1.1.1.1 /8 11.1.1.1 /8 10.1.12.1 /24 10.1.14.1 /24	100
R2	**Lo0** **S0/0.21** **F0/0**	**2.2.2.2 /8** **10.1.12.2 /24** **10.1.23.2 /24**	**200**
R3	Lo0 F0/0 S0/0.34	3.3.3.3 /8 10.1.23.3 /24 10.1.34.3 /24	200
R4	**Lo0** **F0/0** **S0/0.43**	**4.4.4.4 /8** **10.1.14.4 /24** **10.1.34.4 /24**	**200**

Task 1

Configure R2, R3 and R4 in AS 200, these routers should have full mesh peer session between them. R2 and router R4 should have an EBGP peer session to R1 in AS 100. These routers should **only** advertise their loopback interface/s in BGP. Provide NLRI for the links using RIPv2, disable automatic summarization.

On R1:

```
R1(config-if)#router bgp 100
R1(config-router)#No au
R1(config-router)#neighbor 10.1.14.4 remote-as 200
R1(config-router)#neighbor 10.1.12.2 remote-as 200
R1(config-router)#netw 11.0.0.0
R1(config-router)#netw 1.0.0.0

R1(config)#router rip
R1(config-router)#No au
R1(config-router)#ver 2
R1(config-router)#network 10.0.0.0
```

On R2:

```
R2(config)#router bgp 200
R2(config-router)#No au
R2(config-router)#netw 2.0.0.0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#neighbor 10.1.23.3 remote-as 200
R2(config-router)#neighbor 10.1.23.3 next-hop-self
R2(config-router)#neighbor 10.1.34.4 next-hop-self
R2(config-router)#neighbor 10.1.34.4 remote-as 200

R2(config)#router rip
R2(config-router)#No au
R2(config-router)#ver 2
R2(config-router)#netw 10.0.0.0
```

On R3:

```
R3(config)#router bgp 200
R3(config-router)#No au
R3(config-router)#netw 3.0.0.0
R3(config-router)#neighbor 10.1.34.4 remote-as 200
R3(config-router)#neighbor 10.1.23.2 remote-as 200

R3(config)#router rip
R3(config-router)#No au
R3(config-router)#ver 2
R3(config-router)#netw 10.0.0.0
```

On R4:

```
R4(config)#router bgp 200
R4(config-router)#No au
R4(config-router)#netw 4.0.0.0
R4(config-router)#neighbor 10.1.34.3 remote-as 200
R4(config-router)#neighbor 10.1.34.3 next-hop-self
R4(config-router)#neighbor 10.1.23.2 remote-as 200
R4(config-router)#neighbor 10.1.23.2 next-hop-self
R4(config-router)#neighbor 10.1.14.1 remote-as 100

R4(config)#router rip
R4(config-router)#No au
R4(config-router)#ver 2
R4(config-router)#netw 10.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*  2.0.0.0          10.1.14.4                              0 200 i
*>                  10.1.12.2                0             0 200 i
*  3.0.0.0          10.1.12.2                              0 200 i
*>                  10.1.14.4                              0 200 i
*  4.0.0.0          10.1.12.2                              0 200 i
*>                  10.1.14.4                0             0 200 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
```

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
* i1.0.0.0          10.1.14.1                0    100      0 100 i
*>                  10.1.12.1                0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          10.1.23.3                0    100      0 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
* i11.0.0.0         10.1.14.1                0    100      0 100 i
*>                  10.1.12.1                0             0 100 i
```

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
* i1.0.0.0          10.1.14.1                0    100      0 100 i
*>i                 10.1.12.1                0    100      0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
* i11.0.0.0         10.1.14.1                0    100      0 100 i
*>i                 10.1.12.1                0    100      0 100 i
```

On R4:

```
R4#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.14.1                0             0 100 i
* i                 10.1.12.1                0    100      0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*>i3.0.0.0          10.1.34.3                0    100      0 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
*> 11.0.0.0          10.1.14.1                0             0 100 i
* i                  10.1.12.1                0    100      0 100 i
```

Task 2

Configure R1 in AS 100 such that routers in AS 200 use the link through R4-R1 to reach its network 1.0.0.0 /8. Use the AS-Path attribute to accomplish this task.

To configure R1 so the routers in AS 200 take the R4-R1 link to reach network 1.0.0.0 /8:

On R1:

```
R1(config)#access-list 1 permit 1.0.0.0 0.255.255.255

R1(config)#route-map TST1 permit 10
R1(config-route-map)#match ip addr 1
R1(config-route-map)#set as-path prepend 100
R1(config)#route-map TST1 permit 90

R1(config-route-map)#router bgp 100
R1(config-router)#neighbor 10.1.12.2 route-map TST1 out

R1#Clear ip bgp *
```

To verify the configuration:

On R2:

```
R2#Sh ip bgp | b Network
   Network          Next Hop            Metric LocPrf Weight Path
*  1.0.0.0          10.1.12.1                0             0 100 100 i
*>i                 10.1.34.4                0    100      0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          10.1.23.3                0    100      0 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
* i11.0.0.0         10.1.34.4                0    100      0 100 i
*>                  10.1.12.1                0             0 100 i
```

On R3:

```
R3#Sh ip bgp | b Network
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.34.4                0    100      0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
* i11.0.0.0         10.1.34.4                0    100      0 100 i
*>i                 10.1.23.2                0    100      0 100 i
```

On R4:

```
R4#Sh ip bgp | b Network
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.14.1                0             0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*>i3.0.0.0          10.1.34.3                0    100      0 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
* i11.0.0.0         10.1.23.2                0    100      0 100 i
*>                  10.1.14.1                0             0 100 i
```

Note: The AS-path attribute is used to influence the degree of preference in another AS. R2, R3 and R4 will go through R4 to reach network 1.0.0.0 /8.

Task 3

Configure R1 in AS 100 such that the routers in AS 200 use the link through R2-R1 to reach network 11.0.0.0 /8. Use the AS-Path attribute to accomplish this task.

To configure R1 so the routers in AS 200 take the R4-R1 link to reach network 11.0.0.0 /8:

On R1:

```
R1(config)#access-list 11 permit 11.0.0.0 0.255.255.255

R1(config)#route-map TST11 permit 10
R1(config-route-map)#match ip addr 11
R1(config-route-map)#set as-path prepend 100
R1(config)#route-map TST11 permit 20

R1(config-route-map)#router bgp 100
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config-router)#neighbor 10.1.14.4 route-map TST11 out

R1#Clear ip bgp * out
```

To verify the configuration:

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.12.1                0             0 100 i
* i                 10.1.34.4                0    100      0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          10.1.23.3                0    100      0 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
*> 11.0.0.0         10.1.12.1                0             0 100 i
```

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.23.2                0    100      0 100 i
* i                 10.1.34.4                0    100      0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
*>i11.0.0.0         10.1.23.2                0    100      0 100 i
```

On R4:

```
R4#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
* i1.0.0.0          10.1.23.2                0    100      0 100 i
*>                  10.1.14.1                0             0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*>i3.0.0.0          10.1.34.3                0    100      0 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*>i11.0.0.0         10.1.23.2                0    100      0 100 i
*                   10.1.14.1                0             0 100 100 i
```

Note: All the routers will take the R2-R1 link to connect to network 11.0.0.0/8.

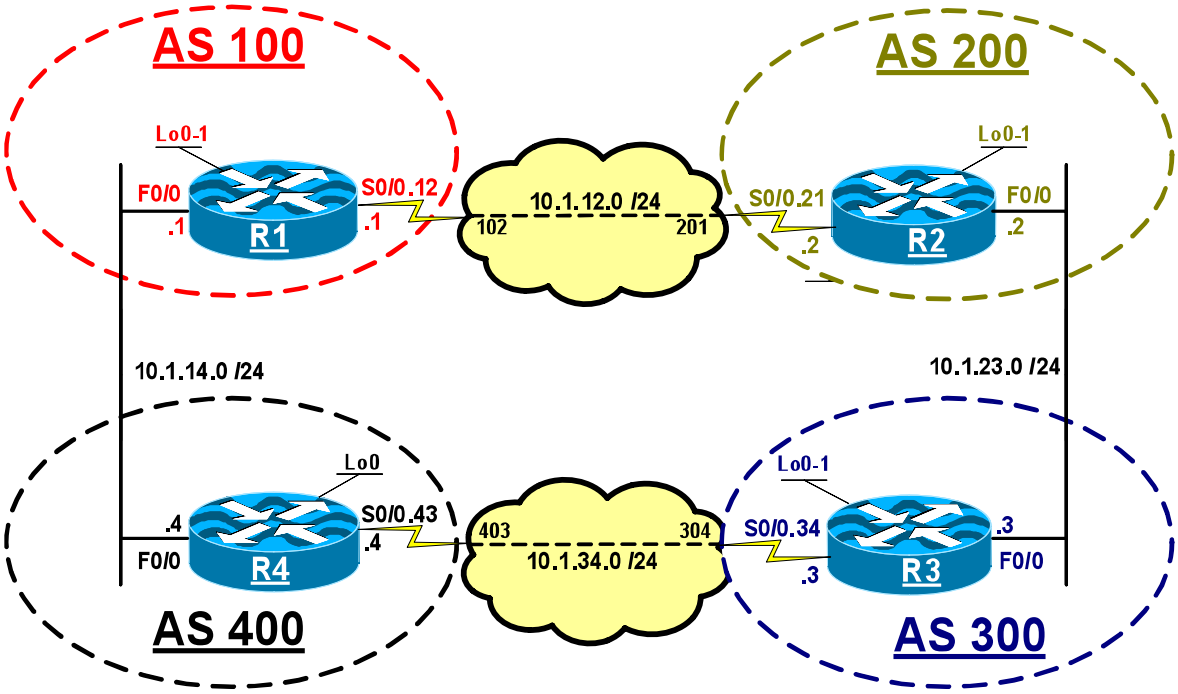
©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 4

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 14 - The Weight Attribute



Lab Setup:

- Configure the routers that are connected to the frame-relay clouds in a point-to-point manner.
- R1 and R4's F0/0 interface should be configured in VLAN 14.
- R2 and R3's F0/0 interface should be configured in VLAN 23.
- Use the following IP addressing chart for IP address assignment.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP addressing:

Router	Interface	IP Address	AS number
R1	Lo0 Lo1 S0/0.12 F0/0	1.1.1.1 /8 11.1.1.1 /8 10.1.12.1 /24 10.1.14.1 /24	100
R2	**Lo0** **Lo1** **S0/0.21** **F0/0**	**2.2.2.2 /8** **22.2.2.2 /8** **10.1.12.2 /24** **10.1.23.2 /24**	**200**
R3	Lo0 Lo1 F0/0 S0/0.34	3.3.3.3 /8 33.3.3.3 /8 10.1.23.3 /24 10.1.34.3 /24	300
R4	**Lo0** **F0/0** **S0/0.43**	**4.4.4.4 /8** **10.1.14.4 /24** **10.1.34.4 /24**	**400**

Task 1

Configure R1 in AS 100 to establish an EBGP peer session with R2 in AS 200 and R4 in AS 400. Router R2 should establish EBGP peer sessions with R1 in AS 100 and R3 in AS 300.Router R3 should establish EBGP peer sessions with R2 in AS 200 and R4 in AS 400.Router R4 should establish EBGP peer sessions with R1 in AS 100 and R3 in AS 300.The BGP routers should **only** advertise their loopback/s in BGP. Provide NLRI for the links using RIPv2 with auto summarization disabled.

On R1:

```
R1(config-if)#router bgp 100
R1(config-router)#No au
R1(config-router)#neighbor 10.1.14.4 remote-as 400
R1(config-router)#neighbor 10.1.12.2 remote-as 200
R1(config-router)#netw 11.0.0.0
R1(config-router)#netw 1.0.0.0

R1(config)#router rip
R1(config-router)#No au
R1(config-router)#ver 2
R1(config-router)#network 10.0.0.0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2:

```
R2(config)#router bgp 200
R2(config-router)#No au
R2(config-router)#netw 2.0.0.0
R2(config-router)#netw 22.0.0.0
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#neighbor 10.1.23.3 remote-as 300

R2(config)#router rip
R2(config-router)#No au
R2(config-router)#ver 2
R2(config-router)#netw 10.0.0.0
```

On R3:

```
R3(config)#router bgp 300
R3(config-router)#No au
R3(config-router)#netw 3.0.0.0
R3(config-router)#netw 33.0.0.0
R3(config-router)#neighbor 10.1.34.4 remote-as 400
R3(config-router)#neighbor 10.1.23.2 remote-as 200

R3(config)#router rip
R3(config-router)#No au
R3(config-router)#ver 2
R3(config-router)#netw 10.0.0.0
```

On R4:

```
R4(config)#router bgp 400
R4(config-router)#No au
R4(config-router)#netw 4.0.0.0
R4(config-router)#neighbor 10.1.34.3 remote-as 300
R4(config-router)#neighbor 10.1.14.1 remote-as 100

R4(config)#router rip
R4(config-router)#No au
R4(config-router)#ver 2
R4(config-router)#netw 10.0.0.0
```

To verify the configuration:

On R4:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4#Sh ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*  1.0.0.0          10.1.34.3                              0 300 200 100 i
*>                  10.1.14.1                0             0 100 i
*  2.0.0.0          10.1.34.3                              0 300 200 i
*>                  10.1.14.1                              0 100 200 i
*  3.0.0.0          10.1.14.1                              0 100 200 300 i
*>                  10.1.34.3                0             0 300 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*  11.0.0.0         10.1.34.3                              0 300 200 100 i
*>                  10.1.14.1                0             0 100 i
*  22.0.0.0         10.1.34.3                              0 300 200 i
*>                  10.1.14.1                              0 100 200 i
*  33.0.0.0         10.1.14.1                              0 100 200 300 i
*>                  10.1.34.3                0             0 300 i
```

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*  2.0.0.0          10.1.14.4                              0 400 300 200 i
*>                  10.1.12.2                0             0 200 i
*> 3.0.0.0          10.1.14.4                              0 400 300 i
*                   10.1.12.2                              0 200 300 i
*> 4.0.0.0          10.1.14.4                0             0 400 i
*                   10.1.12.2                              0 200 300 400 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
*  22.0.0.0         10.1.14.4                              0 400 300 200 i
*>                  10.1.12.2                0             0 200 i
*> 33.0.0.0         10.1.14.4                              0 400 300 i
*                   10.1.12.2                              0 200 300 i
```

Task 2

Configure R1 in AS 100 to use AS 200 to reach all the prefixes within this topology; you must use the "weight" attribute to accomplish this task.

The weight attribute is a cisco-defined attribute that is local to the router. This attribute is not advertised to any BGP neighbor/s. If there is more than one route to a given destination, the weight attribute can decide which path is better. The higher the value the better the preference.

On R1:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.12.2 weight 40000

R1#Clear ip bgp *
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0         40000 200 i
*                   10.1.14.4                              0 400 300 200 i
*> 3.0.0.0          10.1.12.2                         40000 200 300 i
*                   10.1.14.4                              0 400 300 i
*> 4.0.0.0          10.1.12.2                         40000 200 300 400 i
*                   10.1.14.4                0             0 400 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
*> 22.0.0.0         10.1.12.2                0         40000 200 i
*                   10.1.14.4                              0 400 300 200 i
*> 33.0.0.0         10.1.12.2                         40000 200 300 i
*                   10.1.14.4                              0 400 300 i
```

Note: R1 has to traverse through R2 to reach all the networks within this topology.

Task 3

The policy of AS 100 is changed to the following:

- R1 in AS 100 should use R4 in AS 400 to reach networks 33.0.0.0 /8 and network 4.0.0.0 /8.
- Network 3.0.0.0 /8 and existing and future prefixes from AS 200 should have a weight attribute of 54000 through R2.

Let's remove the "weight" command from the previous task:

On R1:

```
R1(config)#router bgp 100
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config-router)#No neigh 10.1.12.2 weight 40000

R1#Clear ip bgp *
```

Let's configure the task:

The following access-list identifies networks 33.0.0.0 /8 and 4.0.0.0 /8:

```
R1(config)#access-list 1 permit 33.0.0.0 0.255.255.255
R1(config)#access-list 1 permit 4.0.0.0 0.255.255.255
```

The following access-list identifies networks 3.0.0.0 /8:

```
R1(config)#access-list 2 permit 3.0.0.0 0.255.255.255
```

The following as-path access-list identifies existing and future prefixes in AS 200

```
R1(config)#ip as-path access-list 1 permit ^200$
```

The following route-map (TST) assigns a weight attribute of 45000 to the networks identified in the access-list 1.

```
R1(config)#route-map TST permit 10
R1(config-route-map)#match ip addr 1
R1(config-route-map)#set weight 45000
R1(config)#route-map TST permit 90
```

The following route-map (TEST) assigns a weight attribute of 54000 to the networks identified in the as-path access-list 1 and the prefix identified in access-list 2.

```
R1(config-route-map)#route-map TEST per 10
R1(config-route-map)#match as-path 1
R1(config-route-map)#set weight 54000

R1(config-route-map)#route-map TEST per 20
R1(config-route-map)#match ip addr 2
R1(config-route-map)#set weight 54000
R1(config-route-map)#route-map TEST per 90

R1(config-route-map)#router bgp 100
R1(config-router)#neighbor 10.1.14.4 route-map TST in
R1(config-router)#neighbor 10.1.12.2 route-map TEST in

R1#Clear ip bgp * in
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The above commands assign the attributes to neighbors R2 and R4.

To verify the configuration:

On R1:

```
R1#Sh ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0         54000 200 i
*                   10.1.14.4                              0 400 300 200 i
*> 3.0.0.0          10.1.12.2                         54000 200 300 i
*                   10.1.14.4                              0 400 300 i
*  4.0.0.0          10.1.12.2                              0 200 300 400 i
*>                  10.1.14.4                0         45000 400 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
*> 22.0.0.0         10.1.12.2                0         54000 200 i
*                   10.1.14.4                              0 400 300 200 i
*  33.0.0.0         10.1.12.2                              0 200 300 i
*>                  10.1.14.4                         45000 400 300 i
```

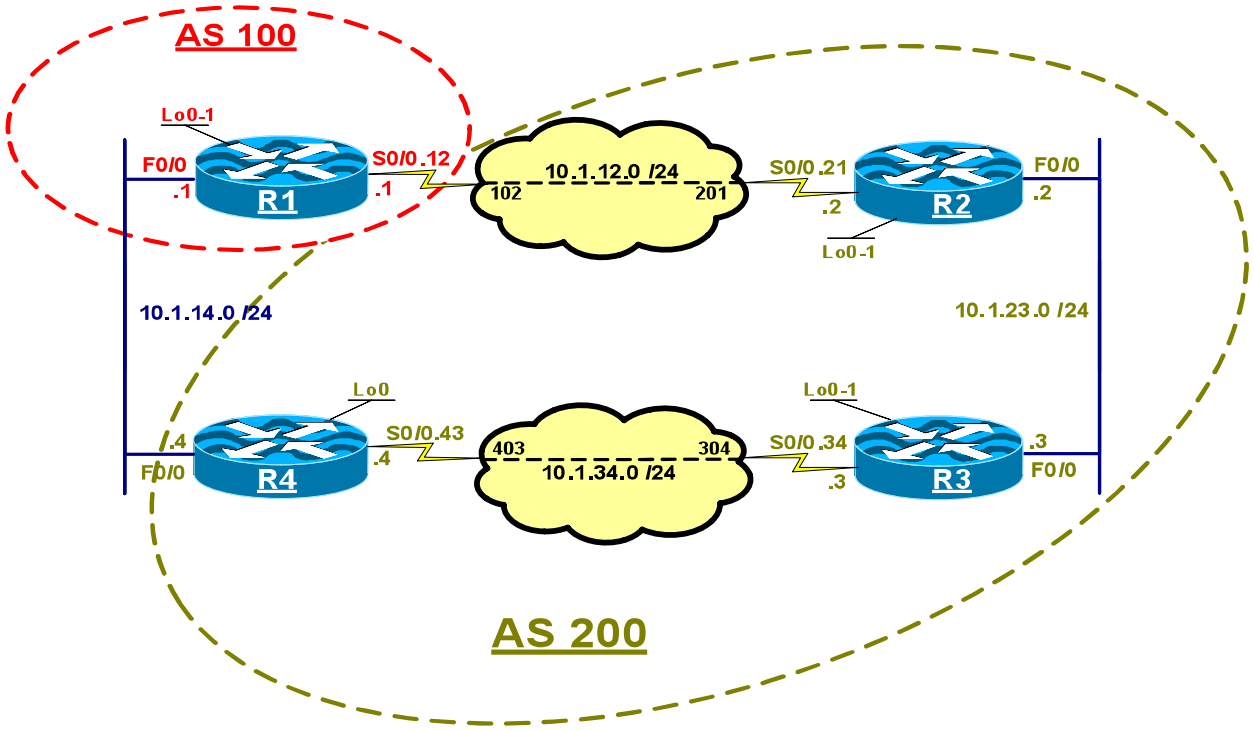
Note: Networks from AS 200 are assigned a weight attribute of 54000, and networks 4.0.0.0 and 33.0.0.0 /8 are assigned a weight attribute of 45000.

Task 4

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 15
Multi Exit Discriminator Attribute



Lab Setup:

- Configure the routers that are connected to the frame-relay clouds in a point-to-point manner.
- R1 and R4's F0/0 interface should be configured in VLAN 14.
- R2 and R3's F0/0 interface should be configured in VLAN 23.
- Use the following IP addressing chart for IP address assignment.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP addressing:

Router	Interface	IP Address	AS number
R1	Lo0 Lo1 S0/0.12 F0/0	1.1.1.1 /8 11.1.1.1 /8 10.1.12.1 /24 10.1.14.1 /24	100
R2	**Lo0** **Lo1** **S0/0.21** **F0/0**	**2.2.2.2 /8** **22.2.2.2 /8** **10.1.12.2 /24** **10.1.23.2 /24**	**200**
R3	**Lo0** **Lo1** **F0/0** **S0/0.34**	**3.3.3.3 /8** **33.3.3.3 /8** **10.1.23.3 /24** **10.1.34.3 /24**	**200**
R4	**Lo0** **F0/0** **S0/0.43**	**4.4.4.4 /8** **10.1.14.4 /24** **10.1.34.4 /24**	**200**

Task 1

Configure routers R2, R3 and R4 in AS 200, these routers should have full mesh peer session between them. Routers R2 and router R4 should have EBGP peer session to R1 in AS 100. BGP routers should **only** advertise their loopback interface/s in BGP. Provide NLRI for the links using RIPv2, disable automatic summarization.

On R1:

```
R1(config-if)#router bgp 100
R1(config-router)#No au
R1(config-router)#neighbor 10.1.14.4 remote-as 200
R1(config-router)#neighbor 10.1.12.2 remote-as 200
R1(config-router)#netw 11.0.0.0
R1(config-router)#netw 1.0.0.0

R1(config)#router rip
R1(config-router)#No au
R1(config-router)#ver 2
R1(config-router)#network 10.0.0.0
```

On R2:

```
R2(config)#router bgp 200
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config-router)#No au
R2(config-router)#netw 2.0.0.0
R2(config-router)#netw 22.0.0.0
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#neighbor 10.1.23.3 remote-as 200
R2(config-router)#neighbor 10.1.23.3 next-hop-self
R2(config-router)#neighbor 10.1.34.4 remote-as 200
R2(config-router)#neighbor 10.1.34.4 next-hop-self

R2(config)#router rip
R2(config-router)#No au
R2(config-router)#ver 2
R2(config-router)#netw 10.0.0.0
```

On R3:

```
R3(config)#router bgp 200
R3(config-router)#No au
R3(config-router)#netw 3.0.0.0
R3(config-router)#netw 33.0.0.0
R3(config-router)#neighbor 10.1.34.4 remote-as 200
R3(config-router)#neighbor 10.1.23.2 remote-as 200

R3(config)#router rip
R3(config-router)#No au
R3(config-router)#ver 2
R3(config-router)#netw 10.0.0.0
```

On R4:

```
R4(config)#router bgp 200
R4(config-router)#No au
R4(config-router)#netw 4.0.0.0
R4(config-router)#neighbor 10.1.34.3 remote-as 200
R4(config-router)#neighbor 10.1.34.3 next-hop-self
R4(config-router)#neighbor 10.1.23.2 remote-as 200
R4(config-router)#neighbor 10.1.23.2 next-hop-self
R4(config-router)#neighbor 10.1.14.1 remote-as 100

R4(config)#router rip
R4(config-router)#No au
R4(config-router)#ver 2
R4(config-router)#netw 10.0.0.0
```

To verify the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1:

```
R1#Sh ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*  2.0.0.0          10.1.14.4                              0 200 i
*>                  10.1.12.2                0             0 200 i
*  3.0.0.0          10.1.14.4                              0 200 i
*>                  10.1.12.2                              0 200 i
*  4.0.0.0          10.1.12.2                              0 200 i
*>                  10.1.14.4                0             0 200 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
*  22.0.0.0         10.1.14.4                              0 200 i
*>                  10.1.12.2                0             0 200 i
*  33.0.0.0         10.1.14.4                              0 200 i
*>                  10.1.12.2                              0 200 i
```

On R2:

```
R2#Sh ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
* i1.0.0.0          10.1.34.4                0    100      0 100 i
*>                  10.1.12.1                0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          10.1.23.3                0    100      0 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
* i11.0.0.0         10.1.34.4                0    100      0 100 i
*>                  10.1.12.1                0             0 100 i
*> 22.0.0.0         0.0.0.0                  0         32768 i
*>i33.0.0.0         10.1.23.3                0    100      0 i
```

On R3:

```
R3#Sh ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.34.4                0    100      0 100 i
* i                 10.1.23.2                0    100      0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
*>i11.0.0.0         10.1.34.4                0    100      0 100 i
* i                 10.1.23.2                0    100      0 100 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
*>i22.0.0.0          10.1.23.2                0    100      0 i
*> 33.0.0.0          0.0.0.0                  0         32768 i
```

On R4:

```
R4#Sh ip bgp | b Network
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.14.1                0             0 100 i
* i                 10.1.23.2                0    100      0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*>i3.0.0.0          10.1.34.3                0    100      0 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*> 11.0.0.0         10.1.14.1                0             0 100 i
* i                 10.1.23.2                0    100      0 100 i
*>i22.0.0.0         10.1.23.2                0    100      0 i
*>i33.0.0.0         10.1.34.3                0    100      0 i
```

Task 2

Configure AS 200 such that R1 in AS 100 takes R4 to reach any prefix advertised in AS 200. Manipulate MED to accomplish this task.

MED is used as a suggestion to an external AS regarding the preferred route into the AS that is advertising the metric. The reason suggestion is used here is because the AS that is receiving MED attribute could use another attribute such as "weight" that will override the MED attribute. As far as MED is concerned, the lower value has better preference.

On R2:

```
R2(config)#route-map TST permit 10
R2(config-route-map)#set metric 100

R2(config)#router bgp 200
R2(config-router)#neighbor 10.1.12.1 route-map TST out

R2#Clear ip bgp * Out
```

To verify the configuration:

On R1:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1#Sh ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.14.4                              0 200 i
*                   10.1.12.2              100             0 200 i
*> 3.0.0.0          10.1.14.4                              0 200 i
*                   10.1.12.2              100             0 200 i
*  4.0.0.0          10.1.12.2              100             0 200 i
*>                  10.1.14.4                0             0 200 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
*> 22.0.0.0         10.1.14.4                              0 200 i
*                   10.1.12.2              100             0 200 i
*> 33.0.0.0         10.1.14.4                              0 200 i
*                   10.1.12.2              100             0 200 i

R1#Traceroute 2.2.2.2

Type escape sequence to abort.
Tracing the route to 2.2.2.2

  1 10.1.14.4 0 msec 4 msec 0 msec
  2 10.1.34.3 24 msec 24 msec 24 msec
  3 10.1.23.2 24 msec *  24 msec
```

Note: R1 receives a MED value of 100 for all prefixes from R2, since R1 does not receive a MED value from R4, R1 considers the MED coming from R4 as zero. When R1 compares the two MED values, it takes the path with a lower MED value, which is R4.

Task 3

Remove the configuration command from the previous task before proceeding to the next task.

On R2:

```
R2(config)#No route-map TST

R2(config)#router bgp 200
R2(config-router)#No neighbor 10.1.12.1 route-map TST out
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 4

Configure the appropriate router/s in AS 200 such that AS 100 goes through R4 to reach network 33.0.0.0 /24 and R2 to reach network 3.0.0.0 /8. Utilize MED to accomplish this task.

The following output shows the existing BGP table of R1 before manipulating the MED attribute:

On R1:

```
R1#Sh ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.14.4                              0 200 i
*                   10.1.12.2                0             0 200 i
*> 3.0.0.0          10.1.14.4                              0 200 i
*                   10.1.12.2                              0 200 i
*  4.0.0.0          10.1.12.2                              0 200 i
*>                  10.1.14.4                0             0 200 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
*> 22.0.0.0         10.1.14.4                              0 200 i
*                   10.1.12.2                0             0 200 i
*> 33.0.0.0         10.1.14.4                              0 200 i
*                   10.1.12.2                              0 200 i
```

Note: Both networks are reachable via R4

On R2:

```
R2(config)#access-list 3 permit 3.0.0.0
R2(config)#access-list 33 permit 33.0.0.0
```

Note: It is very easy to remember that access-list 3 is referencing network 3.0.0.0 and access-list 33 is referencing network 33.0.0.0. If possible, you should choose an access-list name or number that matches the network.

```
R2(config)#route-map TST permit 10

R2(config-route-map)#match ip addr 3
R2(config-route-map)#set metric 50

R2(config)#route-map TST permit 20
R2(config-route-map)#match ip addr 33
R2(config-route-map)#set metric 100
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config)#route-map TST permit 90

R2(config)#router bgp 200
R2(config-router)#neighbor 10.1.12.1 route-map TST out

R2#Clear ip bgp * in
R2#Clear ip bgp * out
```

On R4:

```
R4(config)#access-list 3 permit 3.0.0.0
R4(config)#access-list 33 permit 33.0.0.0

R4(config)#route-map TST permit 10
R4(config-route-map)#match ip addr 3
R4(config-route-map)#set metric 100

R4(config)#route-map TST permit 20
R4(config-route-map)#match ip addr 33
R4(config-route-map)#set metric 50

R4(config)#route-map TST permit 90

R4(config)#router bgp 200
R4(config-router)#neighbor 10.1.14.1 route-map TST out

R4#Clear ip bgp * in
R4#Clear ip bgp * out
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.14.4                              0 200 i
*                   10.1.12.2                0             0 200 i
*  3.0.0.0          10.1.14.4              100             0 200 i
*>                  10.1.12.2               50             0 200 i
*  4.0.0.0          10.1.12.2                              0 200 i
*>                  10.1.14.4                0             0 200 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
*> 22.0.0.0         10.1.14.4                              0 200 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
*                  10.1.12.2                0             0 200 i
*> 33.0.0.0        10.1.14.4               50             0 200 i
*                  10.1.12.2              100             0 200 i
```

Note: Network 3.0.0.0 is reachable via R2 and network 33.0.0.0 is reachable via R4

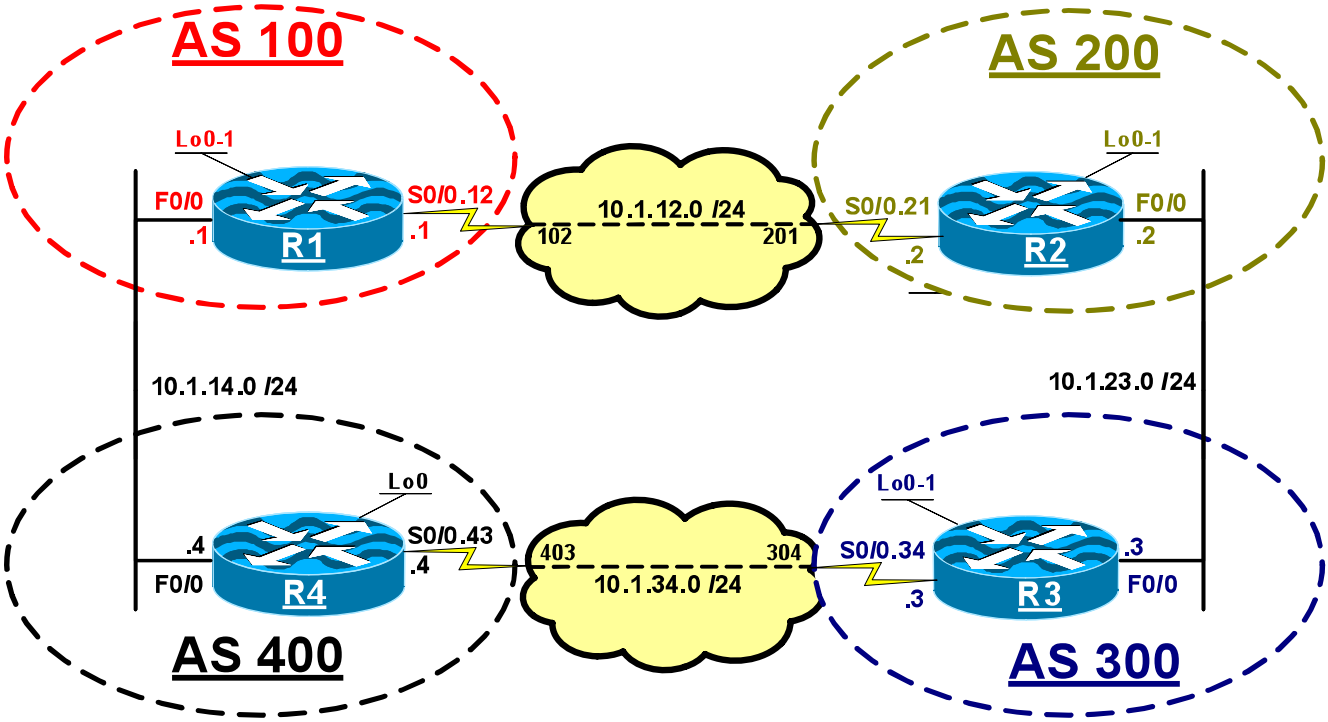
Task 5

Remove BGP routing protocol and any previous BGP related command/s from the previous tasks and reconfigure them based on the chart below. These routers should **only** advertise their loopback interface/s in BGP. The BGP peering should be established as follows:

R1 should establish EBGP peer sessions with R2 and R4 in AS 200 and 400 respectively.
R2 should establish EBGP peer sessions with R1 and R3 in AS 100 and 300 respectively.
R3 should establish EBGP peer sessions with R2 and R4 in AS 200 and 400 respectively.
R4 should establish EBGP peer sessions with R3 and R1 in AS 300 and 100 respectively.

Router	Interface	IP Address	AS number
R1	Lo0 Lo1 S0/0.12 F0/0	1.1.1.1 /8 11.1.1.1 /8 10.1.12.1 /24 10.1.14.1 /24	100
R2	**Lo0** **Lo1** **S0/0.21** **F0/0**	**2.2.2.2 /8** **22.2.2.0 /8** **10.1.12.2 /24** **10.1.23.2 /24**	**200**
R3	Lo0 Lo1 F0/0 S0/0.34	3.3.3.3 /8 33.3.3.0 /8 10.1.23.3 /24 10.1.34.3 /24	300
R4	**Lo0** **F0/0** **S0/0.43**	**4.4.4.4 /8** **10.1.14.4 /24** **10.1.34.4 /24**	**400**

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh



On R1:

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

R1(config)#router bgp 100
R1(config-router)#No au
R1(config-router)#netw 1.0.0.0
R1(config-router)#netw 11.0.0.0
R1(config-router)#neighbor 10.1.12.2 remote-as 200
R1(config-router)#neighbor 10.1.14.4 remote-as 400
```

On R2:

```
R2(config)#No router bgp 200
R2(config)#No access-list 3
R2(config)#No access-list 33
R2(config)#No route-map TST

R2(config)#router bgp 200
R2(config-router)#netw 2.0.0.0
```


©2012 Micronics Networking & Training Inc. All rights reserved

```
R2(config-router)#netw 22.0.0.0
R2(config-router)#No au
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#neighbor 10.1.23.3 remote-as 300
```

On R3:

```
R3(config)#No router bgp 200

R3(config)#router bgp 300
R3(config-router)#No au
R3(config-router)#netw 3.0.0.0
R3(config-router)#netw 33.0.0.0
R3(config-router)#neighbor 10.1.23.2 remote-as 200
R3(config-router)#neighbor 10.1.34.4 remote-as 400
```

On R4:

```
R4(config)#No router bgp 200
R4(config)#No access-list 3
R4(config)#No access-list 33
R4(config)#No route-map TST

R4(config)#router bgp 400
R4(config-router)#No au
R4(config-router)#netw 4.0.0.0
R4(config-router)#neighbor 10.1.14.1 remote-as 100
R4(config-router)#neighbor 10.1.34.3 remote-as 300
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*  2.0.0.0          10.1.14.4                              0 400 300 200 i
*>                  10.1.12.2                0             0 200 i
*  3.0.0.0          10.1.14.4                              0 400 300 i
*>                  10.1.12.2                              0 200 300 i
*> 4.0.0.0          10.1.14.4                0             0 400 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
*  22.0.0.0         10.1.14.4                              0 400 300 200 i
*>                  10.1.12.2                0             0 200 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
*   33.0.0.0          10.1.14.4                              0 400 300 i
*>                    10.1.12.2                              0 200 300 i
```

Task 6

Configure R4 in AS 400 to pass a MED value of 100 to R1 in AS 100 and configure R2 in AS 200 to pass a MED of 120 to R1 in AS 100.

On R2:

```
R2(config)#route-map TST permit 10
R2(config-route-map)#set metric 120

R2(config)#router bgp 200
R2(config-router)#neighbor 10.1.12.1 route-map TST out

R2#Clear ip bgp * out
```

On R4:

```
R4(config)#route-map TST permit 10
R4(config-route-map)#set metric 100

R4(config)#router bgp 400
R4(config-router)#neighbor 10.1.14.1 route-map TST out

R4#Clear ip bgp * out
```

To verify the configuration:

On R1:

```
R1#Clear ip bgp *

R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*  2.0.0.0          10.1.14.4              100             0 400 300 200 i
*>                  10.1.12.2              120             0 200 i
*> 3.0.0.0          10.1.14.4              100             0 400 300 i
*                   10.1.12.2              120             0 200 300 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
*> 4.0.0.0          10.1.14.4                100          0 400 i
*> 11.0.0.0         0.0.0.0                  0          32768 i
*  22.0.0.0         10.1.14.4                100          0 400 300 200 i
*>                  10.1.12.2                120          0 200 i
*> 33.0.0.0         10.1.14.4                100          0 400 300 i
*                   10.1.12.2                120          0 200 300 i
```

Task 7

Ensure that R1 in AS 100 always takes R4 to reach all other prefixes advertised in this topology. You should configure R1 to accomplish this task.

Note: R1 goes through R2 to reach network 22.0.0.0 /8 even though the MED value is lower through R4. The reason this happens is that the AS-PATH is longer through R4. The following command ignores the AS-PATH attribute when the local router performs best path calculation.

On R1:

```
R1(config)#router bgp 100
R1(config-router)#bgp always-compare-med
R1(config-router)#bgp bestpath as-path ignore

R1#Clear ip bgp *
```

Note: The tab key will not work when entering the "bgp bestpath as-path ignore" command. This command is a hidden one.

The MED as stated in RFC 1771, is an optional non-transitive attribute. The value of this attribute may be used by the BGP best path selection process to discriminate among multiple exit points to a neighboring AS. The lower the value the better the path. The MED value comparison is done only among paths from the same AS. The "bgp always-compare-med" command is used to change this behavior by enforcing MED comparison between all paths, regardless of the AS from which the paths are received.

Note: The second command "bgp bestpath as-path ignore" is also needed as part of the solution in accomplishing this task, this command is needed because as-path and the origin attributes are looked at before the MED attribute. Since the origin attribute is identical for all routes and BGP routing protocol is told to ignore the as-path attribute in best path calculation, BGP will look at the MED value next.

To verify the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1:

```
R1#Sh ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*  2.0.0.0          10.1.12.2              120             0 200 i
*>                  10.1.14.4              100             0 400 300 200 i
*  3.0.0.0          10.1.12.2              120             0 200 300 i
*>                  10.1.14.4              100             0 400 300 i
*  4.0.0.0          10.1.12.2              120             0 200 300 400 i
*>                  10.1.14.4              100             0 400 i
*> 11.0.0.0          0.0.0.0                  0         32768 i
*  22.0.0.0         10.1.12.2              120             0 200 i
*>                  10.1.14.4              100             0 400 300 200 i
*  33.0.0.0         10.1.12.2              120             0 200 300 i
*>                  10.1.14.4              100             0 400 300 i
```

Note: R1 takes R4 to reach all the networks in this topology

Task 8

Remove the BGP configuration performed in Task 6. Configure R4 in AS 400 to pass a MED value of 100 to R1 in AS 100; R2 in AS 200 should **not** pass any MED values to R1. Ensure that R1 in AS 100 takes R4 to reach any prefix advertised in this topology except the ones originated by R2. **Do not** use any global config command/s as part of the solution to this task.

To remove the configuration performed in Task 6 and verify the configuration:

On R2:

```
R2(config)#router bgp 200
R2(config-router)#No neighbor 10.1.12.1 route-map TST out

R2(config)#No route-map TST
```

On R4:

```
R4(config)#router bgp 400
R4(config-router)#No neighbor 10.1.14.1 route-map TST out

R4(config)#No route-map TST
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the configuration:

On R1:

```
R1#Sh ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
*                   10.1.14.4                              0 400 300 200 i
*> 3.0.0.0          10.1.12.2                              0 200 300 i
*                   10.1.14.4                              0 400 300 i
*> 4.0.0.0          10.1.12.2                              0 200 300 400 i
*                   10.1.14.4                0             0 400 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
*> 22.0.0.0         10.1.12.2                0             0 200 i
*                   10.1.14.4                              0 400 300 200 i
*> 33.0.0.0         10.1.12.2                              0 200 300 i
*                   10.1.14.4                              0 400 300 i
```

To configure R4 to pass a MED value of 100 to R1:

On R4:

```
R4(config)#route-map TST per 10
R4(config-route-map)#set metric 100

R4(config-route-map)#router bgp 400
R4(config-router)#neigh 10.1.14.1 route-map TST out
```

On R1:

```
R1#Sh ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
*                   10.1.14.4              100             0 400 300 200 i
*> 3.0.0.0          10.1.12.2                              0 200 300 i
*                   10.1.14.4              100             0 400 300 i
*> 4.0.0.0          10.1.12.2                              0 200 300 400 i
*                   10.1.14.4              100             0 400 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
*> 22.0.0.0         10.1.12.2                0             0 200 i
*                   10.1.14.4              100             0 400 300 200 i
*> 33.0.0.0         10.1.12.2                              0 200 300 i
*                   10.1.14.4              100             0 400 300 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Note: R2 is not giving a MED value to R1. R1 will always take the path with the lowest MED. R4 is giving a MED value of 100 to R1. Therefore, R1 will take the path through R2 with a MED value of zero.

Note: The following command tells BGP to set the path with the missing MED value to the worst possible value.

On R1:

```
R1(config)#router bgp 100
R1(config-router)#bgp bestpath med missing-as-worst
```

To verify the configuration:

On R1:

```
R1#Sh ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
*                   10.1.14.4              100             0 400 300 200 i
*> 3.0.0.0          10.1.12.2       4294967295             0 200 300 i
*                   10.1.14.4              100             0 400 300 i
*> 4.0.0.0          10.1.12.2       4294967295             0 200 300 400 i
*                   10.1.14.4              100             0 400 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
*> 22.0.0.0         10.1.12.2                0             0 200 i
*                   10.1.14.4              100             0 400 300 200 i
*> 33.0.0.0         10.1.12.2       4294967295             0 200 300 i
*                   10.1.14.4              100             0 400 300 i
```

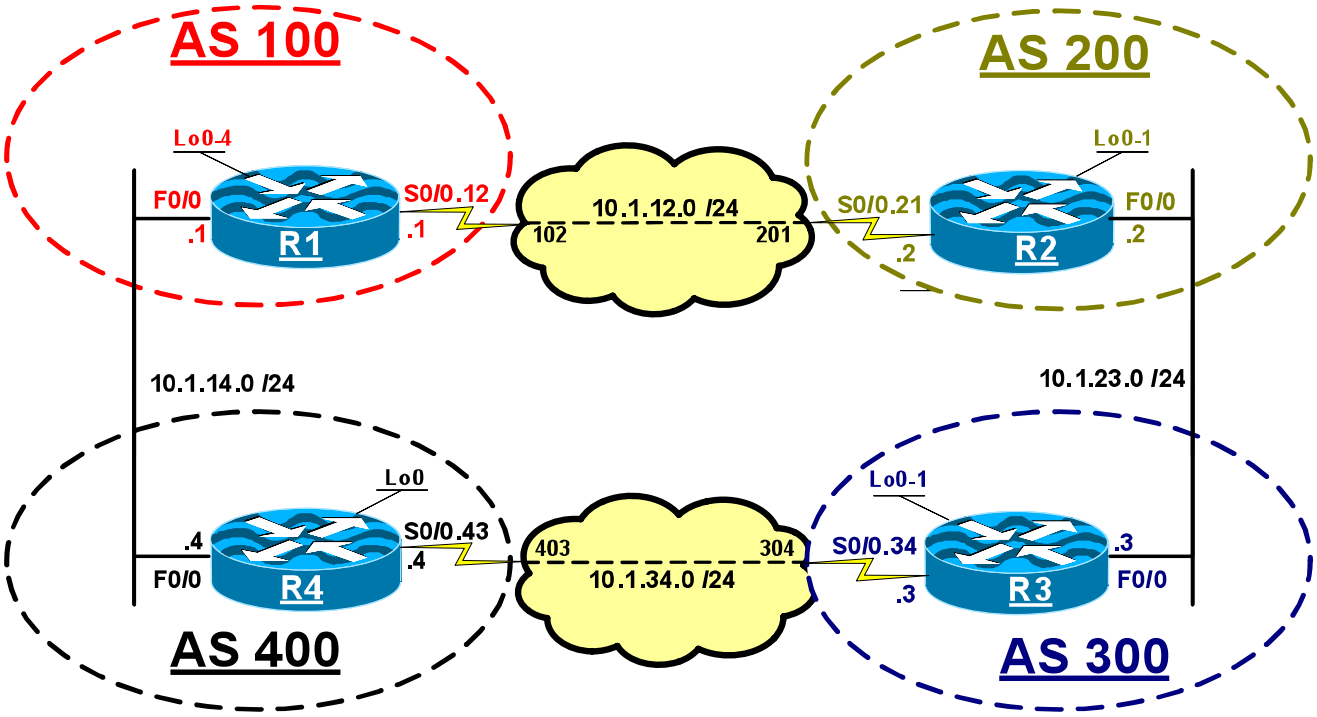
Note: R1 takes the path through R4 to get to all the routes except the network/s advertised by R2. The reason that R1 goes through R2 to reach the networks that are originated by R2 is because R2 is giving a MED value of 0 for the networks that it originated. When R1 compares 0 to 100, it will take the route with the lowest MED.

Task 9

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 16 -
Filtering Using Access-lists and Prefix-lists



Lab Setup:

- Configure the routers that are connected to the frame-relay clouds in a point-to-point manner.
- R1 and R4's F0/0 interface should be configured in VLAN 14.
- R2 and R3's F0/0 interface should be configured in VLAN 23.
- Use the following IP addressing chart for IP address assignment.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP addressing:

Router	Interface	IP Address	AS number
R1	Lo0 Lo1 Lo2 Lo3 Lo4 S0/0.12 F0/0	1.1.0.1 /24 1.1.1.1 /24 1.1.2.1 /24 1.1.3.1 /24 1.1.4.1 /24 10.1.12.1 /24 10.1.14.1 /24	100
R2	**Lo0** **Lo1** **S0/0.21** **F0/0**	**2.2.2.2 /8** **22.2.2.2 /8** **10.1.12.2 /24** **10.1.23.2 /24**	**200**
R3	Lo0 Lo1 F0/0 S0/0.34	3.3.3.3 /8 33.3.3.3 /8 10.1.23.3 /24 10.1.34.3 /24	300
R4	**Lo0** **F0/0** **S0/0.43**	**4.4.4.4 /8** **10.1.14.4 /24** **10.1.34.4 /24**	**400**

Task 1

Configure the routers according to the above IP addressing chart; these routers should **only** advertise their loopback interface/s in BGP, the peering between the routers should be established as follows:
R1 should establish EBGP peer sessions with R2 and R4 in AS 200 and 400 respectively.
R2 should establish EBGP peer sessions with R1 and R3 in AS 100 and 300 respectively.
R3 should establish EBGP peer sessions with R2 and R4 in AS 200 and 400 respectively.
R4 should establish EBGP peer sessions with R3 and R1 in AS 300 and 100 respectively.
Provide NLRI for the links using RIPv2, disable automatic summarization.

On R1:

```
R1(config)#router bgp 100
R1(config-router)#No au
R1(config-router)#netw 1.1.0.0 mask 255.255.255.0
R1(config-router)#netw 1.1.1.0 mask 255.255.255.0
R1(config-router)#netw 1.1.2.0 mask 255.255.255.0
R1(config-router)#netw 1.1.3.0 mask 255.255.255.0
R1(config-router)#netw 1.1.4.0 mask 255.255.255.0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config-router)#neighbor 10.1.12.2 remote-as 200
R1(config-router)#neighbor 10.1.14.4 remote-as 400
```

On R2:

```
R2(config)#router bgp 200
R2(config-router)#No au
R2(config-router)#netw 2.0.0.0
R2(config-router)#netw 22.0.0.0
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#neighbor 10.1.23.3 remote-as 300
```

On R3:

```
R3(config)#router bgp 300
R3(config-router)#No au
R3(config-router)#netw 3.0.0.0
R3(config-router)#netw 33.0.0.0
R3(config-router)#neighbor 10.1.23.2 remote-as 200
R3(config-router)#neighbor 10.1.34.4 remote-as 400
```

On R4:

```
R4(config)#router bgp 400
R4(config-router)#No au
R4(config-router)#netw 4.0.0.0
R4(config-router)#neighbor 10.1.14.1 remote-as 100
R4(config-router)#neighbor 10.1.34.3 remote-as 300
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.0.0/24       0.0.0.0                  0         32768 i
*> 1.1.1.0/24       0.0.0.0                  0         32768 i
*> 1.1.2.0/24       0.0.0.0                  0         32768 i
*> 1.1.3.0/24       0.0.0.0                  0         32768 i
*> 1.1.4.0/24       0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
*  3.0.0.0          10.1.14.4                              0 400 300 i
*>                  10.1.12.2                              0 200 300 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
   Network          Next Hop            Metric LocPrf Weight Path
*  4.0.0.0          10.1.12.2                              0 200 300 400 i
*>                  10.1.14.4                0             0 400 i
*> 22.0.0.0         10.1.12.2                0             0 200 i
*  33.0.0.0         10.1.14.4                              0 400 300 i
*>                  10.1.12.2                              0 200 300 i
```

On R2:

```
R2#Sh ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.0.0/24       10.1.12.1                0             0 100 i
*> 1.1.1.0/24       10.1.12.1                0             0 100 i
*> 1.1.2.0/24       10.1.12.1                0             0 100 i
*> 1.1.3.0/24       10.1.12.1                0             0 100 i
*> 1.1.4.0/24       10.1.12.1                0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.0.0.0          10.1.23.3                0             0 300 i
*  4.0.0.0          10.1.12.1                              0 100 400 i
*>                  10.1.23.3                              0 300 400 i
*> 22.0.0.0         0.0.0.0                  0         32768 i
*> 33.0.0.0         10.1.23.3                0             0 300 i
```

On R3:

```
R3#Sh ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*  1.1.0.0/24       10.1.34.4                              0 400 100 i
*>                  10.1.23.2                              0 200 100 i
*  1.1.1.0/24       10.1.34.4                              0 400 100 i
*>                  10.1.23.2                              0 200 100 i
*  1.1.2.0/24       10.1.34.4                              0 400 100 i
*>                  10.1.23.2                              0 200 100 i
*  1.1.3.0/24       10.1.34.4                              0 400 100 i
*>                  10.1.23.2                              0 200 100 i
*  1.1.4.0/24       10.1.34.4                              0 400 100 i
*>                  10.1.23.2                              0 200 100 i
*  2.0.0.0          10.1.34.4                              0 400 100 200 i
*>                  10.1.23.2                0             0 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*> 4.0.0.0          10.1.34.4                0             0 400 i
*  22.0.0.0         10.1.34.4                              0 400 100 200 i
*>                  10.1.23.2                0             0 200 i
*> 33.0.0.0         0.0.0.0                  0         32768 i
```

On R4:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4#Sh ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.0.0/24       10.1.14.1                0             0 100 i
*                   10.1.34.3                              0 300 200 100 i
*> 1.1.1.0/24       10.1.14.1                0             0 100 i
*                   10.1.34.3                              0 300 200 100 i
*> 1.1.2.0/24       10.1.14.1                0             0 100 i
*                   10.1.34.3                              0 300 200 100 i
*> 1.1.3.0/24       10.1.14.1                0             0 100 i
*                   10.1.34.3                              0 300 200 100 i
*> 1.1.4.0/24       10.1.14.1                0             0 100 i
*                   10.1.34.3                              0 300 200 100 i
*> 2.0.0.0          10.1.14.1                              0 100 200 i
*                   10.1.34.3                              0 300 200 i
*  3.0.0.0          10.1.14.1                              0 100 200 300 i
*>                  10.1.34.3                0             0 300 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*> 22.0.0.0         10.1.14.1                              0 100 200 i
*                   10.1.34.3                              0 300 200 i
*  33.0.0.0         10.1.14.1                              0 100 200 300 i
*>                  10.1.34.3                0             0 300 i
```

Task 2

Configure R2 to block network 1.1.4.0 /24 from getting into its routing and BGP tables. Use distribute-list and access-list to accomplish this task.

On R2:

```
R2(config)#access-list 4 deny 1.1.4.0 0.0.0.255
R2(config)#access-list 4 permit any

R2(config)#router bgp 200
R2(config-router)#neighbor 10.1.12.1 distribute-list 4 in
R2(config-router)#neighbor 10.1.23.3 distribute-list 4 in
```

Note: The tricky part was to understand the topology, if the topology in not understood, silly mistakes can occur which can cost points. The tricky part of this task is to block the prefix from both neighbors.

To verify the configuration:

On R2:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2#Clear ip bgp * in

R2#Sh ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.0.0/24       10.1.12.1                0             0 100 i
*> 1.1.1.0/24       10.1.12.1                0             0 100 i
*> 1.1.2.0/24       10.1.12.1                0             0 100 i
*> 1.1.3.0/24       10.1.12.1                0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.0.0.0          10.1.23.3                0             0 300 i
*  4.0.0.0          10.1.12.1                              0 100 400 i
*>                  10.1.23.3                              0 300 400 i
*> 22.0.0.0         0.0.0.0                  0         32768 i
*> 33.0.0.0         10.1.23.3                0             0 300 i
```

Note: The network is filtered.

Task 3

Remove the configuration command from previous task, and accomplish the same task using prefix-list and distribute-list.

On R2:

```
R2(config)#No access-list 4

R2(config)#router bgp 200
R2(config-router)#No neighbor 10.1.12.1 distribute-list 4 in
R2(config-router)#No neighbor 10.1.23.3 distribute-list 4 in

R2(config)#ip prefix-list TST seq 5 deny 1.1.4.0/24
R2(config)#ip prefix-list TST seq 10 permit 0.0.0.0/0 le 32

R2(config)#router bgp 200
R2(config-router)#neighbor 10.1.12.1 prefix-list TST in
R2(config-router)#neighbor 10.1.23.3 prefix-list TST in

R2#Clear ip bgp * in
```

Note: There are many ways to accomplish a given task, understanding and remembering the different ways can be the key to success.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the configuration:

On R2:

```
R2#Sh ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.0.0/24       10.1.12.1                0             0 100 i
*                   10.1.23.3                              0 300 400 100 i
*> 1.1.1.0/24       10.1.12.1                0             0 100 i
*                   10.1.23.3                              0 300 400 100 i
*> 1.1.2.0/24       10.1.12.1                0             0 100 i
*                   10.1.23.3                              0 300 400 100 i
*> 1.1.3.0/24       10.1.12.1                0             0 100 i
*                   10.1.23.3                              0 300 400 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*  3.0.0.0          10.1.12.1                              0 100 400 300 i
*>                  10.1.23.3                0             0 300 i
*  4.0.0.0          10.1.12.1                              0 100 400 i
*>                  10.1.23.3                              0 300 400 i
*> 22.0.0.0         0.0.0.0                  0         32768 i
*  33.0.0.0         10.1.12.1                              0 100 400 300 i
*>                  10.1.23.3                0             0 300 i
```

Note: The network is filtered.

Task 4

Configure R3 in AS 300 to block network 22.0.0.0 /8 from entering its routing and BGP table. DO NOT use distribute-list or prefix-list. A route-map and an access-list should be used to accomplish this task.

On R3:

```
R3(config)#access-list 22 deny 22.0.0.0
R3(config)#access-list 22 permit any

R3(config)#route-map TST permit 10
R3(config-route-map)#match ip addr 22

R3(config)#router bgp 300
R3(config-router)#neighbor 10.1.23.2 route-map TST in
R3(config-router)#neighbor 10.1.34.4 route-map TST in
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the configuration:

On R3:

```
R3#Sh ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*  1.1.0.0/24       10.1.23.2                              0 200 100 i
*>                  10.1.34.4                              0 400 100 i
*  1.1.1.0/24       10.1.23.2                              0 200 100 i
*>                  10.1.34.4                              0 400 100 i
*  1.1.2.0/24       10.1.23.2                              0 200 100 i
*>                  10.1.34.4                              0 400 100 i
*  1.1.3.0/24       10.1.23.2                              0 200 100 i
*>                  10.1.34.4                              0 400 100 i
*> 1.1.4.0/24       10.1.34.4                              0 400 100 i
*  2.0.0.0          10.1.34.4                              0 400 100 200 i
*>                  10.1.23.2                0             0 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*> 4.0.0.0          10.1.34.4                0             0 400 i
*> 33.0.0.0         0.0.0.0                  0         32768 i
```

Task 5

Remove the Access-list and the route-map configuration from the previous task. Use minimum number of lines in the access-list to filter network 22.0.0.0 /8 from the BGP and IP routing table of R3. You should use an access-list and a route-map to accomplish this task.

On R3:

```
R3(config)#No access-list 22
R3(config)#No route-map TST

R3(config)#access-list 22 permit 22.0.0.0

R3(config)#route-map TST deny 10
R3(config-route-map)#match ip addr 22
R3(config)#route-map TST permit 90
```

Remember that the route-map TST is already configured in the last step.

To verify the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3:

```
R3#Clear ip bgp * in

R3#Sh ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*  1.1.0.0/24       10.1.34.4                              0 400 100 i
*>                  10.1.23.2                              0 200 100 i
*  1.1.1.0/24       10.1.34.4                              0 400 100 i
*>                  10.1.23.2                              0 200 100 i
*  1.1.2.0/24       10.1.34.4                              0 400 100 i
*>                  10.1.23.2                              0 200 100 i
*  1.1.3.0/24       10.1.34.4                              0 400 100 i
*>                  10.1.23.2                              0 200 100 i
*> 1.1.4.0/24       10.1.34.4                              0 400 100 i
*  2.0.0.0          10.1.34.4                              0 400 100 200 i
*>                  10.1.23.2                0             0 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*> 4.0.0.0          10.1.34.4                0             0 400 i
*                   10.1.23.2                              0 200 100 400 i
*> 33.0.0.0         0.0.0.0                  0         32768 i
```

When we are asked to configure an access-list with minimum number of lines, we should always see if the task can be accomplished using a single statement in the access-list.

Task 6

R4 should NOT see networks 1.1.0.0 /24 – 1.1.4.0 /24 in its BGP and/or IP routing table, this filtering should be configured on R1 and R3. Use prefix-list/s to accomplish this task.

On R1: and R3

```
Rx(config)#ip prefix-list TST seq 5 deny 1.1.0.0/24
Rx(config)#ip prefix-list TST seq 10 deny 1.1.1.0/24
Rx(config)#ip prefix-list TST seq 15 deny 1.1.2.0/24
Rx(config)#ip prefix-list TST seq 20 deny 1.1.3.0/24
Rx(config)#ip prefix-list TST seq 25 deny 1.1.4.0/24
Rx(config)#ip prefix-list TST seq 30 permit 0.0.0.0/0 le 32
```

On R1:

```
R1(config)#router bgp 100
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config-router)#neigh 10.1.14.4 prefix-list TST Out

R1#Clear ip bgp * out
```

On R3:

```
R3(config)#router bgp 100
R3(config-router)#neigh 10.1.34.4 prefix-list TST Out

R3#Clear ip bgp * out
```

To verify the configuration:

On R4:

```
R4#Sh ip bgp | b Network

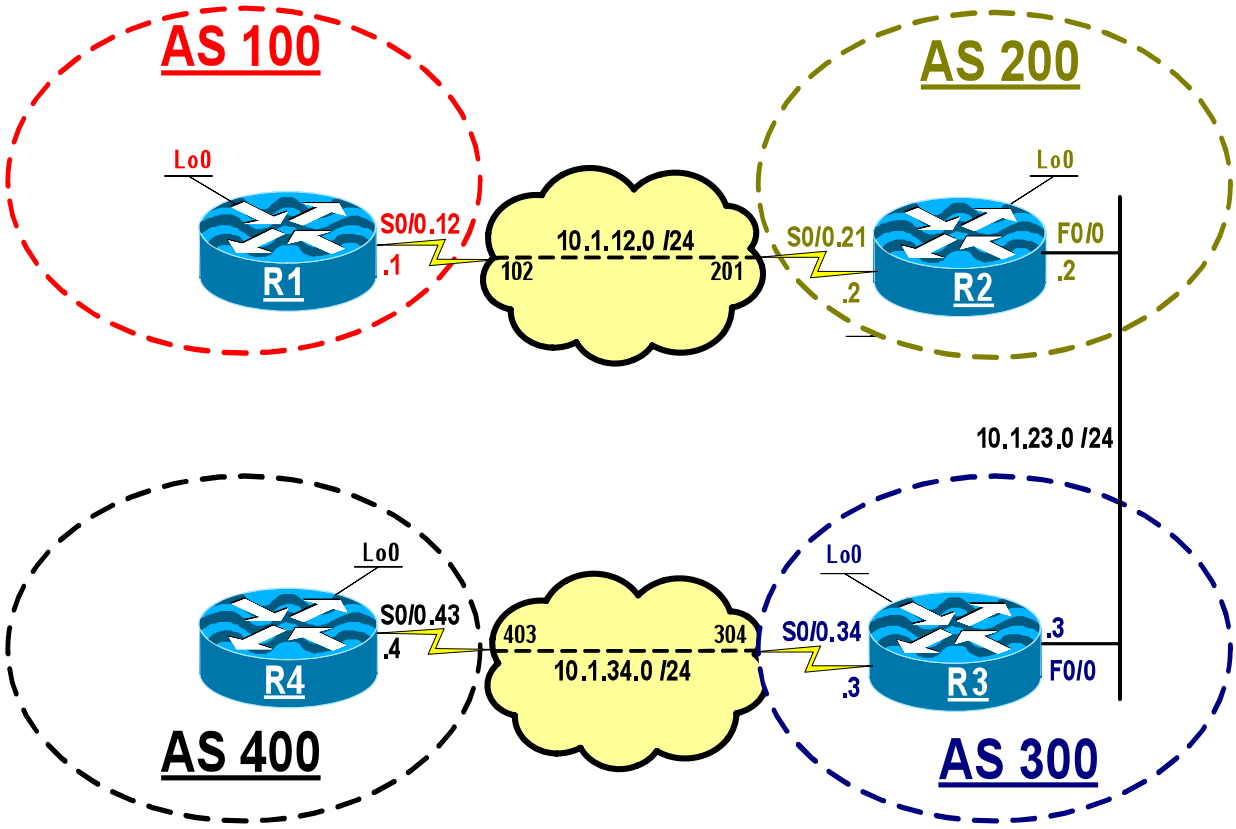
   Network          Next Hop            Metric LocPrf Weight Path
*  2.0.0.0          10.1.14.1                              0 100 200 i
*>                  10.1.34.3                              0 300 200 i
*> 3.0.0.0          10.1.34.3                0             0 300 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*> 22.0.0.0         10.1.14.1                              0 100 200 i
*> 33.0.0.0         10.1.34.3                0             0 300 i
```

Task 7

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 17 - Regular Expressions



Lab Setup:

- Configure the routers that are connected to the frame-relay clouds in a point-to-point manner.
- R2 and R3's F0/0 interface should be configured in VLAN 23.
- Use the following IP addressing chart for IP address assignment.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP addressing:

Router	Interface	IP Address	AS number
R1	Lo0 S0/0.12	1.1.1.1 /8 10.1.12.1 /24	100
R2	**Lo0** **S0/0.21** **F0/0**	**2.2.2.2 /8** **10.1.12.2 /24** **10.1.23.2 /24**	**200**
R3	Lo0 F0/0 S0/0.34	3.3.3.3 /8 10.1.23.3 /24 10.1.34.3 /24	300
R4	**Lo0** **S0/0.43**	**4.4.4.4 /8** **10.1.34.4 /24**	**400**

Task 1

Configure the routers according to the above IP addressing chart; these routers should **on,ly** advertise their loopback0 interface in BGP, the peering between the routers should be established as follows:
R1 should establish EBGP peer sessions with R2 in AS 200.
R2 should establish EBGP peer sessions with R1 and R3 in AS 100 and 300 respectively.
R3 should establish EBGP peer sessions with R2 and R4 in AS 200 and 400 respectively.
R4 should establish EBGP peer sessions with R3 in AS 300.
Provide NLRI for the links using RIPv2, disable automatic summarization.

On R1:

```
R1(config)#router bgp 100
R1(config-router)#No au
R1(config-router)#netw 1.0.0.0
R1(config-router)#neighbor 10.1.12.2 remote-as 200
```

On R2:

```
R2(config)#router bgp 200
R2(config-router)#netw 2.0.0.0
R2(config-router)#No au
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#neighbor 10.1.23.3 remote-as 300
```

On R3:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config)#router bgp 300
R3(config-router)#No au
R3(config-router)#netw 3.0.0.0
R3(config-router)#neighbor 10.1.23.2 remote-as 200
R3(config-router)#neighbor 10.1.34.4 remote-as 400
```

On R4:

```
R4(config)#router bgp 400
R4(config-router)#No au
R4(config-router)#netw 4.0.0.0
R4(config-router)#neighbor 10.1.34.3 remote-as 300
```

On All Routers:

```
Rx(config)#router rip
Rx(config-router)#No au
Rx(config-router)#ver 2
Rx(config-router)#Network 10.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.0.0.0          10.1.12.2                              0 200 300 i
*> 4.0.0.0          10.1.12.2                              0 200 300 400 i
```

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.12.1                0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.0.0.0          10.1.23.3                0             0 300 i
*> 4.0.0.0          10.1.23.3                              0 300 400 i
```

On R3:

```
R3#Show ip bgp | b Network
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.23.2                              0 200 100 i
*> 2.0.0.0          10.1.23.2                0             0 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*> 4.0.0.0          10.1.34.4                0             0 400 i
```

On R4:

```
R4#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.34.3                              0 300 200 100 i
*> 2.0.0.0          10.1.34.3                              0 300 200 i
*> 3.0.0.0          10.1.34.3                0             0 300 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
```

Task 2

Configure R1 such that it blocks all the prefixes that originated in AS 300.

On R1:

```
R1(config)#ip as-path access-list 1 deny _300$
R1(config)#ip as-path access-list 1 permit .*

R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.12.2 filter-list 1 in

R1#Clear ip bgp * in
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 4.0.0.0          10.1.12.2                              0 200 300 400 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 3

Remove the configuration command/s from the previous task before proceeding to the next task.

On R1:

```
R1(config)#No ip as-path access-list 1

R1(config)#router bgp 100
R1(config-router)#No neighbor 10.1.12.2 filter-list 1 in

R1#Clear ip bgp * in
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.0.0.0          10.1.12.2                              0 200 300 i
*> 4.0.0.0          10.1.12.2                              0 200 300 400 i
```

Task 4

Configure R1 such that it blocks all the prefixes that traversed through AS 300.

On R1:

```
R1(config)#ip as-path access-list 1 deny _300_
R1(config)#ip as-path access-list 1 permit .*

R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.12.2 filter-list 1 in

R1#Clear ip bgp * in
```

To verify the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
```

Task 5

Remove the configuration command from the previous task before proceeding to the next task.

On R1:

```
R1(config)#No ip as-path access-list 1

R1(config)#router bgp 100
R1(config-router)#No neighbor 10.1.12.2 filter-list 1 in

R1#Clear ip bgp * in
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.0.0.0          10.1.12.2                              0 200 300 i
*> 4.0.0.0          10.1.12.2                              0 200 300 400 i
```

Task 6

Configure R3 such that it does not advertise the prefixes that originated in its own AS to any of its neighbors.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3:

```
R3(config)#ip as-path access-list 1 deny ^$
R3(config)#ip as-path access-list 1 permit .*

R3(config)#router bgp 300
R3(config-router)#neighbor 10.1.23.2 filter-list 1 out
R3(config-router)#neighbor 10.1.34.4 filter-list 1 out

R3#Clear ip bgp * out
```

To verify the configuration:

On R4:

```
R4#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.34.3                              0 300 200 100 i
*> 2.0.0.0          10.1.34.3                              0 300 200 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
```

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.12.1                0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 4.0.0.0          10.1.23.3                              0 300 400 i
```

Task 7

Remove the configuration command from the previous task before proceeding to the next task.

On R3:

```
R3(config)#No ip as-path access-list 1 deny ^$

R3(config)#router bgp 300
R3(config-router)#No neighbor 10.1.23.2 filter-list 1 out
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-router)#No neighbor 10.1.34.4 filter-list 1 out

R3#Clear ip bgp * out
```

To verify the configuration:

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.12.1                0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.0.0.0          10.1.23.3                0             0 300 i
*> 4.0.0.0          10.1.23.3                              0 300 400 i
```

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.23.2                              0 200 100 i
*> 2.0.0.0          10.1.23.2                0             0 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*> 4.0.0.0          10.1.34.4                0             0 400 i
```

On R4:

```
R4#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.34.3                              0 300 200 100 i
*> 2.0.0.0          10.1.34.3                              0 300 200 i
*> 3.0.0.0          10.1.34.3                0             0 300 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
```

Task 8

Configure R3 such that it blocks the entire network from its neighboring AS 200.

On R3:

```
R3(config)#ip as-path access-list 1 deny ^200$
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config)#ip as-path access-list 1 permit .*

R3(config)#router bgp 300
R3(config-router)#neighbor 10.1.23.2 filter-list 1 in

R3#Clear ip bgp * in
```

To verify the configuration:

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.23.2                              0 200 100 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*> 4.0.0.0          10.1.34.4                0             0 400 i
```

Task 9

Remove the configuration command from the previous task before proceeding to the next task.

On R3:

```
R3(config)#No ip as-path access-list 1

R3(config)#router bgp 300
R3(config-router)#No neighbor 10.1.23.2 filter-list 1 in

R3#Clear ip bgp * in
```

To verify the configuration:

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.23.2                              0 200 100 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.23.2                0             0 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*> 4.0.0.0          10.1.34.4                0             0 400 i
```

Task 10

Configure R3 such that it blocks all the prefixes from its directly connected neighbors.

On R3:

```
R3(config)#ip as-path access-list 1 deny ^[0-9]+$
R3(config)#ip as-path access-list 1 permit .*

R3(config)#router bgp 300
R3(config-router)#neighbor 10.1.23.2 filter-list 1 in
R3(config-router)#neighbor 10.1.34.4 filter-list 1 in

R3#Clear ip bgp * in
```

To verify the configuration:

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.23.2                             0 200 100 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
```

Task 11

Remove the configuration command from the previous task before proceeding to the next task.

On R3:

```
R3(config)#No ip as-path access-list 1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config)#router bgp 300
R3(config-router)#No neighbor 10.1.23.2 filter-list 1 in
R3(config-router)#No neighbor 10.1.34.4 filter-list 1 in

R3#Clear ip bgp * in
```

To verify the configuration:

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.23.2                              0 200 100 i
*> 2.0.0.0          10.1.23.2                0             0 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*> 4.0.0.0          10.1.34.4                0             0 400 i
```

Task 12

Configure R1 such that it blocks all the prefixes that originated in AS 300 and traversed through AS 200.

On R1:

```
R1(config)#ip as-path access-list 1 deny _200_300$
R1(config)#ip as-path access-list 1 permit .*

R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.12.2 filter-list 1 in

R1#Clear ip bgp * in
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
*> 2.0.0.0          10.1.12.2                0              0 200 i
*> 4.0.0.0          10.1.12.2                                0 200 300 400 i
```

Task 13

Remove the configuration command from the previous task before proceeding to the next task.

On R1:

```
R1(config)#No ip as-path access-list 1

R1(config)#router bgp 100
R1(config-router)#No neighbor 10.1.12.2 filter-list 1 in

R1#Clear ip bgp * in
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.0.0.0          10.1.12.2                              0 200 300 i
*> 4.0.0.0          10.1.12.2                              0 200 300 400 i
```

Task 14

Your company has decided to use more complex regular expressions in the future, configure the routers such that they do not use recursive algorithm when processing regular expressions.

On R1:

```
R1(config)#router bgp 100
R1(config-router)#bgp regexp deterministic
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2:

```
R2(config)#router bgp 200
R2(config-router)#bgp regexp deterministic
```

On R3:

```
R3(config)#router bgp 300
R3(config-router)#bgp regexp deterministic
```

On R2:

```
R4(config)#router bgp 400
R4(config-router)#bgp regexp deterministic
```

Task 15

Configure R1 in AS 100 to **only** allow prefixes from its existing and future directly connected ASes, these ASes should be allowed to prepend.

Before the "As-path access-list" is written and applied, a "show ip bgp regexp" command should be issued, if the desired output is displayed, then, the "as-path access-list" should be written and applied.

Note: In the following regular expression the "(_\1)" section can be thought of the memory button in a calculator, the expression before it "^([0-9]+)" is what you are putting in the memory location 1, and the "*" specifies zero or more of the expression that is in the memory location 1.

On R1:

```
R1#Sh ip bgp regexp ^([0-9]+)(_\1)*$

BGP table version is 19, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.12.2                0             0 200 i

R1(config)#ip as-path access-list 1 permit ^([0-9]+)(_\1)*$
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the output of the "IP as-path access-list 1" command:

```
R1#Show ip bgp filter-list 1

BGP table version is 19, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.12.2                0             0 200 i
```

Therefore there are two ways to test the regular expression before applying it to a given neighbor/s:

1. **Show ip bgp regexp**
2. **Show ip bgp filter-list**

The next step is to apply it to neighbor/s:

```
R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.12.2 filter-list 1 in
```

The following uses the refresh messages so the changes can apply to the existing and new prefixes.

```
R1#Clear ip bgp * in
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
```

Task 16

Remove the "IP as-path access-list" and the "neighbor filter-list" commands from the previous step.

On R1:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config)#No ip as-path access-list 1

R1(config)#router bgp 100
R1(config-router)#No neighbor 10.1.12.2 filter-list 1 in

R1#Clear ip bgp * in
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.0.0.0          10.1.12.2                              0 200 300 i
*> 4.0.0.0          10.1.12.2                              0 200 300 400 i
```

Task 17

Configure R1 to prepend its own AS number once in addition to the default behavior.

On R1:

```
R1(config)#route-map TST permit 10
R1(config-route-map)#set as-path prepend 100

R1(config-route-map)#router bgp 100
R1(config-router)#neighbor 10.1.12.2 route-map TST out

R1#Clear ip bgp * out
```

To verify the configuration:

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.12.1                0             0 100 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
*> 3.0.0.0          10.1.23.3                0              0 300 i
*> 4.0.0.0          10.1.23.3                               0 300 400 i
```

Task 18

Configure R2 such that it allows AS-Path prepend from AS 100 **only** if they have prepended their own AS number and **not** another AS number.

On R2:

```
R2(config)#ip as-path access-list 1 permit ^([0-9]+)(_\1)*$

R2(config)#router bgp 200
R2(config-router)#neighbor 10.1.12.1 filter-list 1 in

R2#Clear ip bgp * in
```

To verify the configuration:

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.12.1                0             0 100 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.0.0.0          10.1.23.3                0             0 300 i
*> 4.0.0.0          10.1.23.3                              0 300 400 i
```

To test the configuration:

Change the configuration of R1 to prepend different AS numbers

On R1:

```
R1(config)#No route-map TST


R1(config)#route-map TST permit 10
R1(config-route-map)#set as-path prepend 100 600 700 800
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1#Clear ip bgp * out
```

On R2:

```
R2#Show ip bgp | b Network

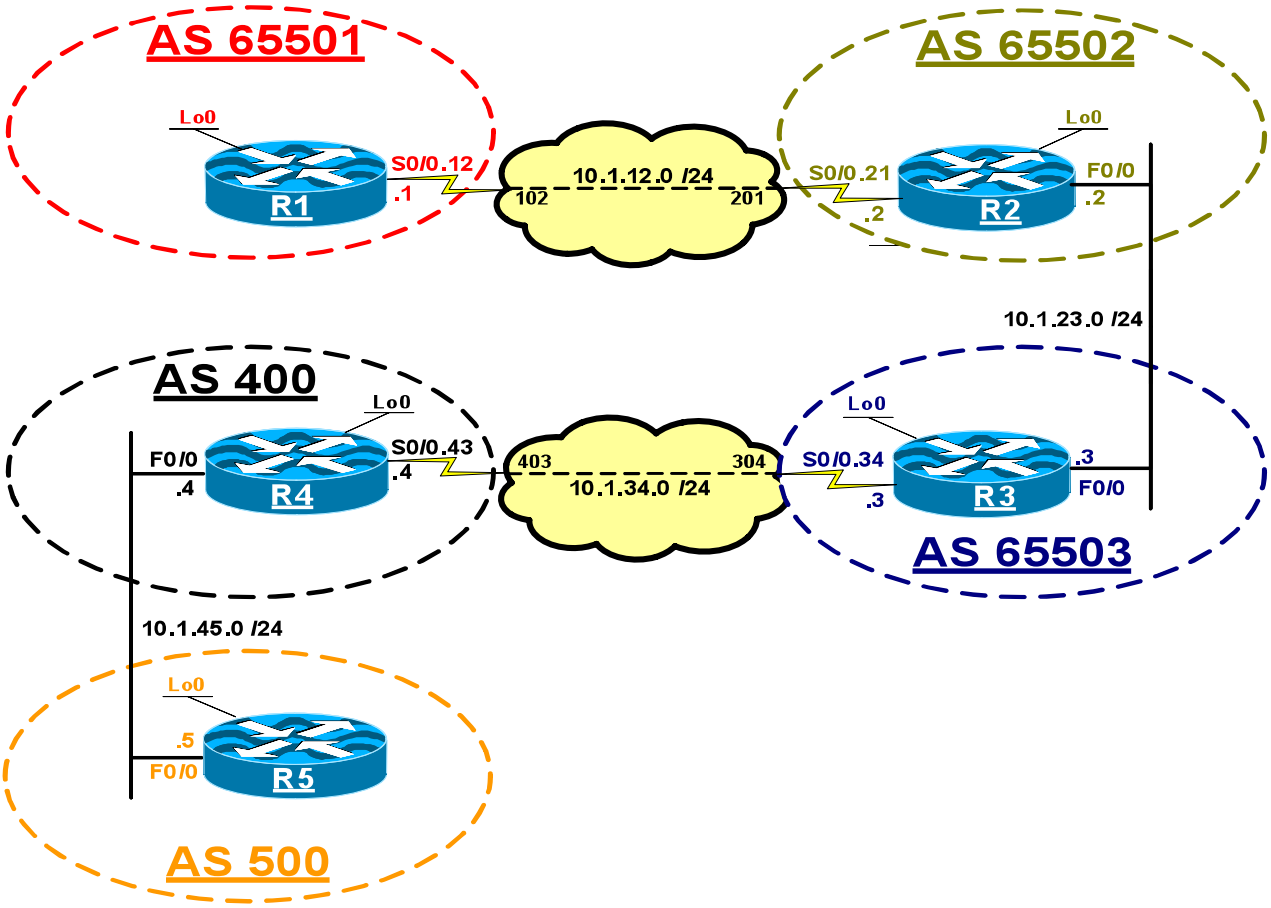
   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.0.0.0          10.1.23.3                0             0 300 i
*> 4.0.0.0          10.1.23.3                              0 300 400 i
```

Task 19

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 18 - Advanced BGP Configurations



Lab Setup:

- Configure the routers that are connected to the frame-relay clouds in a point-to-point manner.
- R2 and R3's F0/0 interface should be configured in VLAN 23.
- R4 and R5's F0/0 interface should be configured in VLAN 45.
- Use the following IP addressing chart for IP address assignment.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP addressing:

Router	Interface	IP Address	AS number
R1	Lo0 S0/0.12	1.1.1.1 /8 10.1.12.1 /24	65501
R2	**Lo0** **S0/0.21** **F0/0**	**2.2.2.2 /8** **10.1.12.2 /24** **10.1.23.2 /24**	**65502**
R3	Lo0 S0/0.34 F0/0	3.3.3.3 /8 10.1.34.3 /24 10.1.23.3 /24	65503
R4	**Lo0** **S0/0.43** **F0/0**	**4.4.4.4 /8** **10.1.34.4 /24** **10.1.45.4 /24**	**400**
R5	Lo0 F0/0	5.5.5.5 /8 10.1.45.5 /24	500

Task 1

Configure BGP peering on the routers as follows:

R1 in AS 65501, should establish an EBGP peer session with R2 in AS 65502.
R2 in AS 65502, should establish EBGP peer sessions with R1 and R3 in AS 65501 and 65503 respectively.
R3 in AS 65503, should establish EBGP peer sessions with R2 and R4 in AS 65502 and AS 400 respectively.
R4 in AS 400, should establish EBGP peer sessions with R3 and R5 in AS 65503 and AS 500 respectively.
R5 in AS 500 should establish an EBGP peer session with R4 in AS 400.
Provide NLRI to the links that connect the routers using RIPv2.
These routers should advertise their loopback interface/s in their assigned AS.

On R1:

```
R1(config)#router bgp 65501
R1(config-router)#No au
R1(config-router)#netw 1.0.0.0
R1(config-router)#neighbor 10.1.12.2 remote-as 65502

R1(config)#router rip
R1(config-router)#No au
R1(config-router)#ver 2
R1(config-router)#netw 10.0.0.0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2:

```
R2(config)#router bgp 65502
R2(config-router)#No au
R2(config-router)#netw 2.0.0.0
R2(config-router)#neighbor 10.1.12.1 remote-as 65501
R2(config-router)#neighbor 10.1.23.3 remote-as 65503

R2(config-router)#router rip
R2(config-router)#ver 2
R2(config-router)#No au
R2(config-router)#netw 10.0.0.0
```

On R3:

```
R3(config)#router bgp 65503
R3(config-router)#No au
R3(config-router)#netw 3.0.0.0
R3(config-router)#neighbor 10.1.23.2 remote-as 65502
R3(config-router)#neighbor 10.1.34.4 remote-as 400

R3(config-router)#router rip
R3(config-router)#No au
R3(config-router)#ver 2
R3(config-router)#netw 10.0.0.0
```

On R4:

```
R4(config)#router bgp 400
R4(config-router)#No au
R4(config-router)#netw 4.0.0.0
R4(config-router)#neighbor 10.1.34.3 remote-as 65503
R4(config-router)#neighbor 10.1.45.5 remote-as 500

R4(config-router)#router rip
R4(config-router)#ver 2
R4(config-router)#No au
R4(config-router)#netw 10.0.0.0
```

On R5:

```
R5(config-if)#router bgp 500
R5(config-router)#No au
R5(config-router)#netw 5.0.0.0
R5(config-router)#neighbor 10.1.45.4 remote-as 400
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R5(config-router)#router rip
R5(config-router)#No au
R5(config-router)#ver 2
R5(config-router)#netw 10.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 65502 i
*> 3.0.0.0          10.1.12.2                              0 65502 65503 i
*> 4.0.0.0          10.1.12.2                              0 65502 65503 400 i
*> 5.0.0.0          10.1.12.2                              0 65502 65503 400 500 i
```

Task 2

Configure R2 such that if any of its links to any of its external peers go down, the routes from that particular neighbor are **not** removed from the BGP table. The routes should only be removed if a given neighbor is down for longer than 5 minutes.

On R2:

```
R2(config-if)#router bgp 65502
R2(config-router)#No bgp fast-external-fallover
R2(config-router)#neighbor 10.1.12.1 timers 60 300
R2(config-router)#neighbor 10.1.23.3 timers 60 300
```

On R1:

```
R1(config-if)#router bgp 65501
R1(config-router)#neighbor 10.1.12.2 timers 60 300
```

On R3:

```
R3(config-if)#router bgp 65503
R3(config-router)#neighbor 10.1.23.2 timers 60 300
```

This feature only supports the directly connected peers. If BGP fast external-fallover is disabled, the

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

BGP routing process will wait until the configured hold timer expires before the peer session is reset. If the other neighbors are not configured with the timers, the lowest value (the default = 180 seconds) will be negotiated.

Task 3

Configure R5 in AS 500 so it sets the hello and hold timer values to double its default value for its neighboring router R4.

On R5:

```
R5(config)#router bgp 500
R5(config-router)#neighbor 10.1.45.4 timer 120 360
```

Task 4

Configure R4 such that it removes the private AS numbers when it advertises prefixes to R5.

Before configuring R4, we should display the BGP table of R5.

On R5:

```
R5#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.45.4                              0 400 65503 65502 65501 i
*> 2.0.0.0          10.1.45.4                              0 400 65503 65502 i
*> 3.0.0.0          10.1.45.4                              0 400 65503 i
*> 4.0.0.0          10.1.45.4                0             0 400 i
*> 5.0.0.0          0.0.0.0                  0         32768 i
```

On R4:

```
R4(config)#router bgp 400
R4(config-router)#neighbor 10.1.45.5 remove-private-as

R4#Clear ip bgp * out
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the configuration:

On R5:

```
R5#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.45.4                              0 400 i
*> 2.0.0.0          10.1.45.4                              0 400 i
*> 3.0.0.0          10.1.45.4                              0 400 i
*> 4.0.0.0          10.1.45.4                0             0 400 i
*> 5.0.0.0          0.0.0.0                  0         32768 i
```

Task 5

Configure R1 such that it sets the minimum time between sending BGP advertisement for its EBGP neighbors to 60 seconds.

On R1:

```
R1(config)#router bgp 65501
R1(config-router)#neighbor 10.1.12.2 advertisement-interval 60
```

The default values are as follows:

- **External peers – 30 seconds.**
- **Internal peers – 5 seconds.**

This command sets the minimum interval between the sending of BGP routing updates for internal and external neighbors.

Task 6

Configure the following loopback1 interfaces On R4: and R5 and advertise them in RIPv2 only.
Ensure that R4 and R5 establish their EBGP peer session based on their Lo1 interface. **Do not** use "**neighbor ebgp multihop**" command to establish this peer session.

Configure R4's - Lo1 = 44.4.4.4 /8 and R5's – Lo1 = 55.5.5.5 /8

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4:

```
R4(config)#int lo1
R4(config-if)#ip addr 44.4.4.4 255.0.0.0

R4(config-router)#router rip
R4(config-router)#netw 44.0.0.0
```

On R5:

```
R5(config)#int lo1
R5(config-if)#ip addr 55.5.5.5 255.0.0.0

R5(config-router)#router rip
R5(config-router)#netw 55.0.0.0
```

On R4:

```
R4(config-if)#router bgp 400
R4(config-router)#No neighbor 10.1.45.5 remote-as 500
R4(config-router)#neighbor 55.5.5.5 remote-as 500
R4(config-router)#neighbor 55.5.5.5 update-source lo1
R4(config-router)#neighbor 55.5.5.5 disable-connected-check
```

On R5:

```
R5(config-if)#router bgp 500
R5(config-router)#No neighbor 10.1.45.4 remote-as 400
R5(config-router)#neighbor 44.4.4.4 remote-as 400
R5(config-router)#neighbor 44.4.4.4 update-source lo1
R5(config-router)#neighbor 44.4.4.4 disable-connected-check
```

The default behavior of BGP:

A BGP speaker will verify the connection of a single-hop EBGP peering session to determine if the EBGP peer is directly connected to the same network segment. If the peer is not directly connected to the same network segment, the connection verification will fail and it will prevent the peering session from being established.

The "neighbor disable-connected-check" command will disable the connection verification process for EBGP peering session that are reachable by a single hop but are configured on a loopback interface or configured with a non-directly connected IP address.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 7

For security purposes, configure R1 and R2 such that they only accept IP packets with a TTL count in the IP header that is equal to or greater than 253. If the TTL count of the IP packet is 252 or less, these routers should ignore the packet/s.

On R1:

```
R1(config)#router bgp 65501
R1(config-router)#neighbor 10.1.12.2 ttl-security hops 2
```

On R2:

```
R2(config)#router bgp 65502
R2(config-router)#neighbor 10.1.12.1 ttl-security hops 2
```

This feature is enabled in 12.3(7)T and it provides a lightweight security for the BGP routers against CPU utilization attacks. These types of attacks flood the network with IP packets that contain forged source and destination IP addresses in the packet headers. This configuration accepts only IP packets with a TTL count that is equal to or greater than the locally configured value.

Task 8

Remove the configuration from task 6 and ensure that the two routers establish a peer session using their directly connected interface's IP address. R4 in AS 400 should establish a peer session with R5 in AS 555, and R5 in AS 500 should establish a peer session with R4 in AS 400.

On R4:

```
R4(config)#No int lo1
```

On R5:

```
R5(config)#No int lo1
```

On R4:

```
R4(config-if)#router bgp 400
R4(config-router)#neighbor 10.1.45.5 remote-as 500
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config-router)#No neighbor 55.5.5.5 remote-as 500
R4(config-router)#No neighbor 55.5.5.5 update-source lo1
R4(config-router)#No neighbor 55.5.5.5 disable-connected-check

R4(config-router)#router rip
R4(config-router)#No netw 44.0.0.0
```

On R5:

```
R5(config-if)#router bgp 500
R5(config-router)#neighbor 10.1.45.4 remote-as 400
R5(config-router)#No neighbor 44.4.4.4 remote-as 400
R5(config-router)#No neighbor 44.4.4.4 update-source lo1
R5(config-router)#No neighbor 44.4.4.4 disable-connected-check

R5(config-router)#router rip
R5(config-router)#No netw 55.0.0.0
```

On R4:

```
R4(config-if)#router bgp 400
R4(config-router)#No neighbor 10.1.45.5 remote-as 500
R4(config-router)#neighbor 10.1.45.5 remote-as 555
```

On R5:

```
R5(config-if)#router bgp 500
R5(config-router)#neighbor 10.1.45.4 local-as 555
```

Typically used for AS migration and should be removed when the migration is complete. The "local-as" command prepends the AS number specified in the command to the as-path. With this command, the local router appears to be in another AS.

Task 9

Configure R3 such that it limits the number of AS-path segments that are permitted in inbound routes to 20.

On R3:

```
R3(config)#router bgp 65503
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-router)#bgp maxas-limit 20
```

This command discards routes that have a number of AS-Path segments that exceed the specified value. The range is 1 – 2000. The default value in BGP is 75. This command was introduced in IOS release 12.2.

Task 10

Configure R3 to inject a default route to neighbor 10.1.34.4 ONLY if there is a route to 2.0.0.0 /8 in R3's routing table.

On R3:

```
R3(config)#access-list 2 permit 2.0.0.0

R3(config)#route-map TEST permit 10
R3(config-route-map)#match ip addr 2

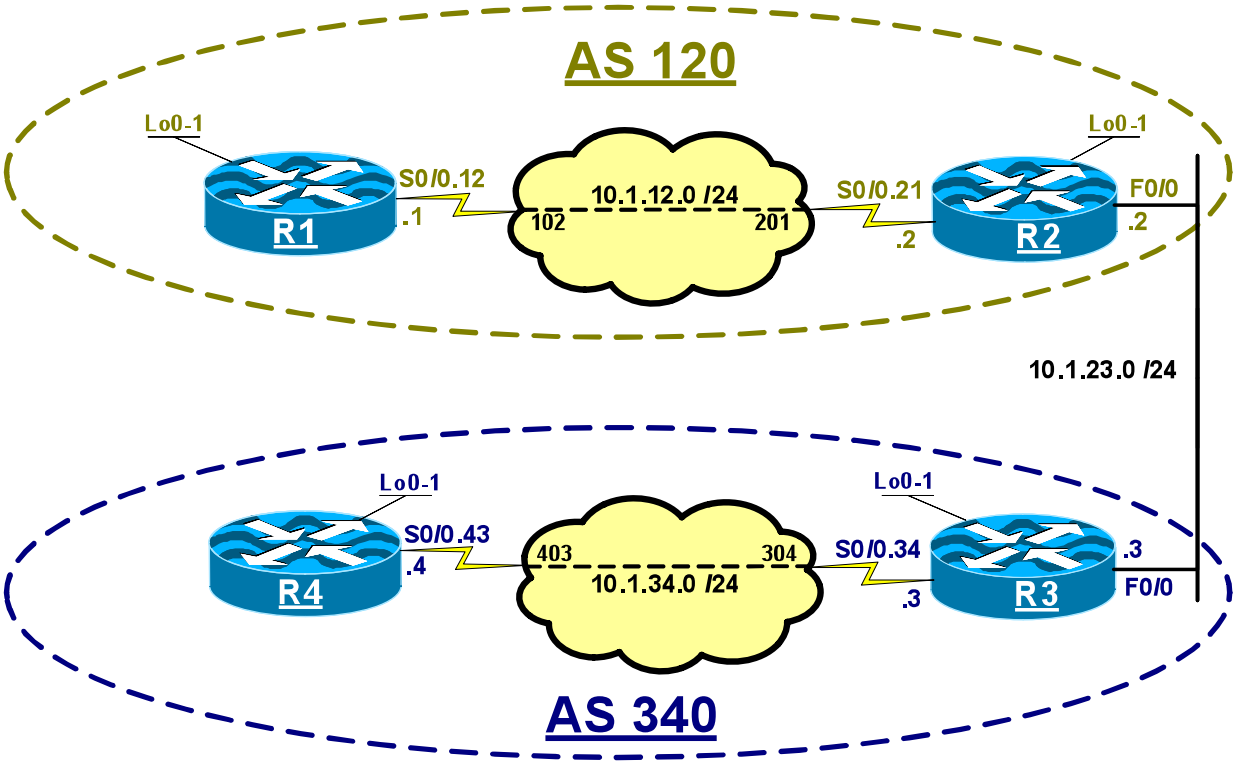
R3(config)#router bgp 65503
R3(config-router)#neighbor 10.1.34.4 default-originate route-map TEST
```

Task 11

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 19 - Administrative Distance



Lab Setup:

- Configure the routers that are connected to the frame-relay clouds in a point-to-point manner.
- R2 and R3's F0/0 interface should be configured in VLAN 23.
- Use the following IP addressing chart for IP address assignment.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP addressing:

Router	Interface	IP Address	AS number
R1	Lo0 Lo1 S0/0.12	1.1.1.1 /8 11.1.1.1 /8 10.1.12.1 /24	120
R2	**Lo0** **Lo1** **S0/0.21** **F0/0**	**2.2.2.2 /8** **22.2.2.2 /8** **10.1.12.2 /24** **10.1.23.2 /24**	**120**
R3	Lo0 Lo1 F0/0 S0/0.34	3.3.3.3 /8 33.3.3.3 /8 10.1.23.3 /24 10.1.34.3 /24	340
R4	**Lo0** **Lo1** **S0/0.43**	**4.4.4.4 /8** **44.4.4.4 /8** **10.1.34.4 /24**	**340**

Task 1

Configure the routers according to the above IP addressing chart; these routers should ONLY advertise their Loopback interfaces in BGP, the peering between the routers should be established as follows:
R1 in AS 120 should establish an IBGP peer session with R2 in AS 120.
R2 should establish an IBGP peer session with R1 in AS 120 and an EBGP peer session with R3 in AS 340.
R3 should establish an EBGP peer session with R2 in AS 120 and an IBGP peer session with R4 in AS 340.
R4 should establish an IBGP peer sessions with R3 in AS 340.
Provide NLRI for the links using RIPv2, disable automatic summarization.

On R1:

```
R1(config)#router bgp 120
R1(config-router)#No au
R1(config-router)#netw 1.0.0.0
R1(config-router)#netw 11.0.0.0
R1(config-router)#neighbor 10.1.12.2 remote-as 120
```

On R2:

```
R2(config)#router bgp 120
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config-router)#netw 2.0.0.0
R2(config-router)#netw 22.0.0.0
R2(config-router)#No au
R2(config-router)#neighbor 10.1.12.1 remote-as 120
R2(config-router)#neighbor 10.1.23.3 remote-as 340
```

On R3:

```
R3(config)#router bgp 340
R3(config-router)#No au
R3(config-router)#netw 3.0.0.0
R3(config-router)#netw 33.0.0.0
R3(config-router)#neighbor 10.1.23.2 remote-as 120
R3(config-router)#neighbor 10.1.34.4 remote-as 340
```

On R4:

```
R4(config)#router bgp 340
R4(config-router)#No au
R4(config-router)#netw 4.0.0.0
R4(config-router)#netw 44.0.0.0
R4(config-router)#neighbor 10.1.34.3 remote-as 340
```

On All Routers:

```
Rx(config)#router rip
Rx(config-router)#No au
Rx(config-router)#ver 2
Rx(config-router)#network 10.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*>i2.0.0.0          10.1.12.2                0    100      0 i
*>i3.0.0.0          10.1.23.3                0    100      0 340 i
*>i4.0.0.0          10.1.23.3                0    100      0 340 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
*>i22.0.0.0         10.1.12.2                0    100      0 i
*>i33.0.0.0         10.1.23.3                0    100      0 340 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
*>i44.0.0.0          10.1.23.3                0    100      0 340 i
```

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.12.1                0    100      0 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.0.0.0          10.1.23.3                0             0 340 i
*> 4.0.0.0          10.1.23.3                              0 340 i
*>i11.0.0.0         10.1.12.1                0    100      0 i
*> 22.0.0.0         0.0.0.0                  0         32768 i
*> 33.0.0.0         10.1.23.3                0             0 340 i
*> 44.0.0.0         10.1.23.3                              0 340 i
```

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.23.2                              0 120 i
*> 2.0.0.0          10.1.23.2                0             0 120 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
*> 11.0.0.0         10.1.23.2                              0 120 i
*> 22.0.0.0         10.1.23.2                0             0 120 i
*> 33.0.0.0         0.0.0.0                  0         32768 i
*>i44.0.0.0         10.1.34.4                0    100      0 i
```

On R4:

```
R4#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.23.2                0    100      0 120 i
*>i2.0.0.0          10.1.23.2                0    100      0 120 i
*>i3.0.0.0          10.1.34.3                0    100      0 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*>i11.0.0.0         10.1.23.2                0    100      0 120 i
*>i22.0.0.0         10.1.23.2                0    100      0 120 i
*>i33.0.0.0         10.1.34.3                0    100      0 i
*> 44.0.0.0         0.0.0.0                  0         32768 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 2

Configure R2 such that it changes the administrative distance of all prefixes received from R3 to 150.

On R2:

```
R2(config)#router bgp 120
R2(config-router)#distance 150 10.1.23.3 0.0.0.0
```

To verify the configuration:

On R2:

```
R2#Clear ip route *

R2#Show ip route bgp

B     1.0.0.0/8 [200/0] via 10.1.12.1, 00:00:02
B     33.0.0.0/8 [150/0] via 10.1.23.3, 00:00:02
B     3.0.0.0/8 [150/0] via 10.1.23.3, 00:00:02
B     4.0.0.0/8 [150/0] via 10.1.23.3, 00:00:02
B     11.0.0.0/8 [200/0] via 10.1.12.1, 00:00:02
B     44.0.0.0/8 [150/0] via 10.1.23.3, 00:00:02
```

Task 3

Remove the configuration commands from the previous task before proceeding to the next task.

On R2:

```
R2(config)#router bgp 120
R2(config-router)#No distance 150 10.1.23.3 0.0.0.0
```

To verify the configuration

On R2:

```
R2#Clear ip route *
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2#Sh ip route bgp

B     1.0.0.0/8 [200/0] via 10.1.12.1, 00:00:26
B     33.0.0.0/8 [20/0] via 10.1.23.3, 00:00:26
B     3.0.0.0/8 [20/0] via 10.1.23.3, 00:00:26
B     4.0.0.0/8 [20/0] via 10.1.23.3, 00:00:26
B     11.0.0.0/8 [200/0] via 10.1.12.1, 00:00:26
B     44.0.0.0/8 [20/0] via 10.1.23.3, 00:00:26
```

Task 4

Configure R2 such that it **only** changes the administrative distance of prefix 33.0.0.0 /8 to 150.

On R2:

```
R2(config)#access-list 33 permit 33.0.0.0

R2(config)#router bgp 120
R2(config-router)#distance 150 10.1.23.3 0.0.0.0 33
```

Note: The distance command changes the administrative distance to 150 from neighbor 10.1.23.3 followed by the inverse-mask of the IP address of the neighbor for whatever prefix that is specified in access-list 33.

To verify the configuration:

On R2:

```
R2#Clear ip route *

R2#Sh ip route bgp

B     1.0.0.0/8 [200/0] via 10.1.12.1, 00:00:33
B     33.0.0.0/8 [150/0] via 10.1.23.3, 00:00:33
B     3.0.0.0/8 [20/0] via 10.1.23.3, 00:00:33
B     4.0.0.0/8 [20/0] via 10.1.23.3, 00:00:33
B     11.0.0.0/8 [200/0] via 10.1.12.1, 00:00:33
B     44.0.0.0/8 [20/0] via 10.1.23.3, 00:00:33
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 5

Using minimum number of commands change the administrative distance of all IBGP prefixes to 90, EBGP prefixes to 60 and locally generated prefixes to 20. This should be performed on the routers in AS 340.

On R3: and R4

```
Rx(config)#router bgp 340
Rx(config-router)#distance bgp 60 90 20

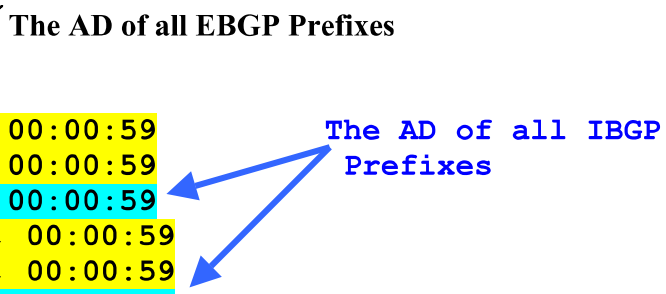
Rx#Clear ip bgp * in

Rx#Clear ip bgp * out
```

On R3:

```
R3#Show ip route bgp

B     1.0.0.0/8 [60/0] via 10.1.23.2, 00:00:59
B     2.0.0.0/8 [60/0] via 10.1.23.2, 00:00:59
B     4.0.0.0/8 [90/0] via 10.1.34.4, 00:00:59
B     22.0.0.0/8 [60/0] via 10.1.23.2, 00:00:59
B     11.0.0.0/8 [60/0] via 10.1.23.2, 00:00:59
B     44.0.0.0/8 [90/0] via 10.1.34.4, 00:00:59
```



To see the locally generated routes, an aggregate-address is configured:

On R3:

```
R3(config)#int lo100
R3(config-if)#ip address 100.1.0.3 255.255.255.0

R3(config)#int lo101
R3(config-if)#ip address 100.1.1.3 255.255.255.0

R3(config)#int lo102
R3(config-if)#ip address 100.1.2.3 255.255.255.0

R3(config)#int lo103
R3(config-if)#ip address 100.1.3.3 255.255.255.0

R3(config-if)#router bgp 340
R3(config-router)#netw 100.1.0.0 mask 255.255.255.0
R3(config-router)#netw 100.1.1.0 mask 255.255.255.0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-router)#netw 100.1.2.0 mask 255.255.255.0
R3(config-router)#netw 100.1.3.0 mask 255.255.255.0
R3(config-router)#aggregate-address 100.1.0.0 255.255.252.0 summary-only
```

To verify the configuration:

On R3:

```
R3#Show ip route bgp

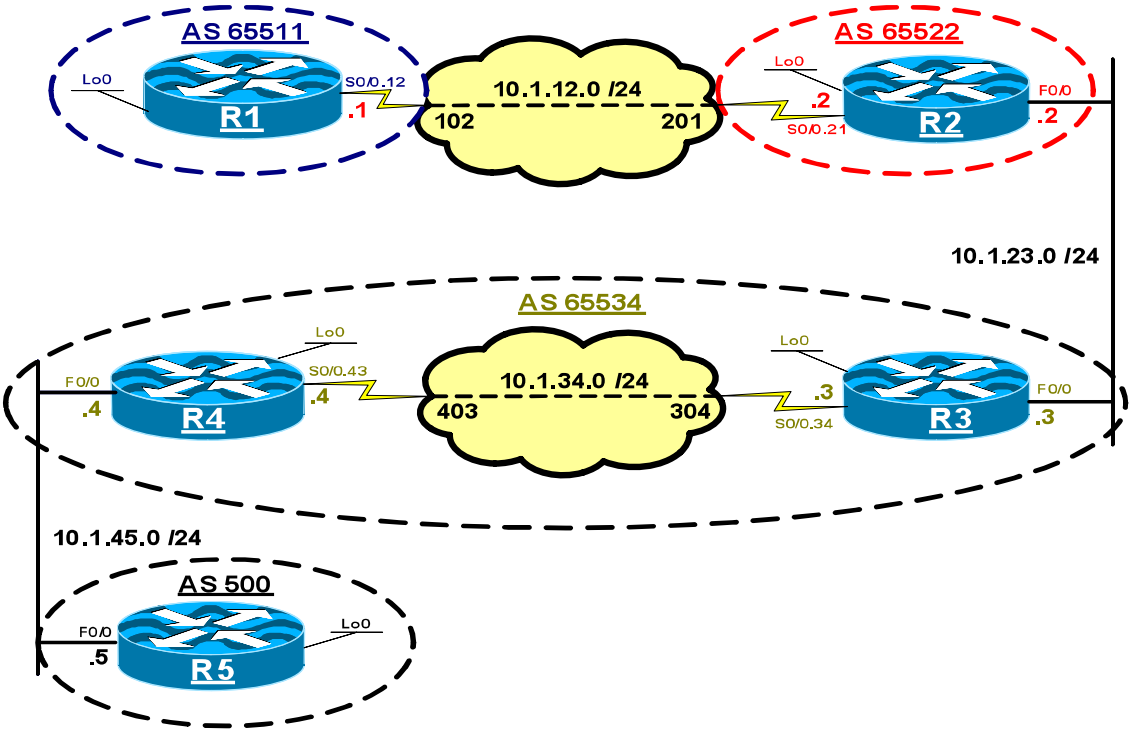
B     1.0.0.0/8 [60/0] via 10.1.23.2, 00:12:35
B     2.0.0.0/8 [60/0] via 10.1.23.2, 00:12:35
      100.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
B        100.1.0.0/22 [20/0] via 0.0.0.0, 00:07:56, Null0
B     4.0.0.0/8 [90/0] via 10.1.34.4, 00:12:35
B     22.0.0.0/8 [60/0] via 10.1.23.2, 00:12:35
B     11.0.0.0/8 [60/0] via 10.1.23.2, 00:12:35
B     44.0.0.0/8 [90/0] via 10.1.34.4, 00:12:35
```

Task 6

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 20 - BGP Confederation



Lab Setup:

- Configure the routers that are connected to the frame-relay clouds in a point-to-point manner.
- R2 and R3's F0/0 interface should be configured in VLAN 23.
- R4 and R5's F0/0 interface should be configured in VLAN 45.
- Use the following IP addressing chart for IP address assignment.

IP addressing:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Router	Interface	IP Address	AS number
R1	Lo0 S0/0.12	1.1.1.1 /8 10.1.12.1 /24	65511
R2	**Lo0** **S0/0.21** **F0/0**	**2.2.2.2 /8** **10.1.12.2 /24** **10.1.23.2 /24**	**65522**
R3	Lo0 S0/0.34 F0/0	3.3.3.3 /8 10.1.34.3 /24 10.1.23.3 /24	65534
R4	**Lo0** **S0/0.43** **F0/0**	**4.4.4.4 /8** **10.1.34.4 /24** **10.1.45.4 /24**	**65534**
R5	Lo0 F0/0	5.5.5.5 /8 10.1.45.5 /24	500

Task 1

Configure BGP peering on the routers as follows:

- R1 in AS 65511, should establish an EBGP peer session with R2 in AS 65522.
- R2 in AS 65522, should establish EBGP peer sessions with R1 and R3 in AS 65511 and 65534 respectively.
- R3 in AS 65534, should establish an EBGP peer sessions with R2 in AS 65522 and an IBGP peer session with R4 in AS 65534.
- R4 in AS 65534, should establish an IBGP peer sessions with R3 in AS 65534 and an EBGP peer session with R5 in AS 500.
- R5 in AS 500 should establish an EBGP peer session with R4 in AS 100.
- Provide NLRI to the links that connect the routers using RIPv2.
- These routers should advertise their Loopback interface in BGP.

On All Routers

```
Rx(config-router)#router rip
Rx(config-router)#NO au
Rx(config-router)#ver 2
Rx(config-router)#network 10.0.0.0
```

On R1:

```
R1(config)#router bgp 65511
R1(config-router)#NO au
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config-router)#netw 1.0.0.0
R1(config-router)#neighbor 10.1.12.2 remote-as 65522

R1(config-router)#bgp confederation identifier 100
```

The "BGP confederation identifier" command is used to configure a single AS number to identify a group of smaller ASes (Sub-ASes) as a single confederation. This command must be configured on all the routers within the confederation.

A confederation can be used to reduce the IBGP mesh by dividing a large single AS into multiple Sub-ASes and then grouping them into a single confederation.

```
R1(config-router)#bgp confederation peers 65522
```

The above command is used to identify the directly connected EBGP sub-confed peers.

On R2:

```
R2(config)#router bgp 65522
R2(config-router)#NO au
R2(config-router)#netw 2.0.0.0
R2(config-router)#neighbor 10.1.12.1 remote-as 65511
R2(config-router)#neighbor 10.1.23.3 remote-as 65534

R2(config-router)#bgp confederation identifier 100
R2(config-router)#bgp confederation peers 65511 65534
```

On R3:

```
R3(config)#router bgp 65534
R3(config-router)#NO au
R3(config-router)#neighbor 10.1.34.4 remote-as 65534
R3(config-router)#neighbor 10.1.23.2 remote-as 65522
R3(config-router)#netw 3.0.0.0

R3(config-router)#bgp confederation identifier 100
R3(config-router)#bgp confederation peers 65522
```

On R4:

```
R4(config)#router bgp 65534
R4(config-router)#NO au
R4(config-router)#netw 4.0.0.0

R4(config-router)#bgp confederation identifier 100
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config-router)#neighbor 10.1.45.5 remote-as 500
R4(config-router)#neighbor 10.1.34.3 remote-as 65534
```

Note: R4 does not have any directly connected EBGP sub-confed peers; therefore, the "BGP confederation peers" command is not configured.

On R5:

```
R5(config)#router bgp 500
R5(config-router)#NO au
R5(config-router)#netw 5.0.0.0

R5(config-router)#neighbor 10.1.45.4 remote-as 100
```

To verify the configuration:

On R5:

```
R5#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.45.4                              0 100 i
*> 2.0.0.0          10.1.45.4                              0 100 i
*> 3.0.0.0          10.1.45.4                              0 100 i
*> 4.0.0.0          10.1.45.4                0             0 100 i
*> 5.0.0.0          0.0.0.0                  0         32768 i
```

Note: From R5's perspective, all the prefixes are from AS 100.

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0    100      0 (65522) i
*> 3.0.0.0          10.1.23.3                0    100      0 (65522 65534) i
*> 4.0.0.0          10.1.34.4                0    100      0 (65522 65534) i
*> 5.0.0.0          10.1.45.5                0    100      0 (65522 65534) 500 i
```

Note: The AS-Path/s in the parenthesis are the private AS numbers within the confederation, AS 500 is outside of the parenthesis because it is not part of the confederation.

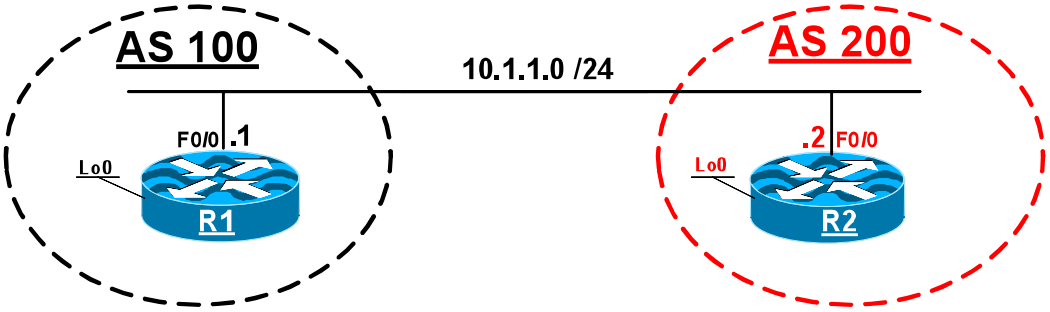
©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 2

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 21 – BGP Hiding Local AS number



Lab Setup:

- Configure the F0/0 interface of both routers in VLAN 100
- Use the IP addressing chart for IP address assignment

IP addressing:

Router	Interface / IP address
R1	F0/0 = 10.1.1.1 /24 Loopback0 = 1.1.1.1 /8
R2	**F0/0 = 10.1.1.2 /24** **Loopback0 = 2.2.2.2 /8**

Task 1

Configure R1 in AS 100 to establish an EBGP session with R2 in AS 200. These routers should advertise their loopback 0 interface. Disable auto summarization.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1:

```
R1(config)#router bgp 100
R1(config-router)#No auto-summary
R1(config-router)#network 1.0.0.0
R1(config-router)#neighbor 10.1.1.2 remote-as 200
```

On R2:

```
R2(config)#router bgp 200
R2(config-router)#No auto-summary
R2(config-router)#network 2.0.0.0
R2(config-router)#neighbor 10.1.1.1 remote-as 100
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.1.2                 0             0 200 i
```

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.1.1                 0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
```

Note: From R2's perspective prefix 1.0.0.0/8 was originated & advertised by AS 100, and from R1's perspective, prefix 2.0.0.0/8 was originated and advertised by AS 200.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 2

Configure R1 in AS 111 to establish an EBGP session with R2 in AS 200 such that the output of the "**show ip bgp**" command on these two routers will be identical to the follows:

On R1:

```
R1#Sh ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.1.2                 0             0 100 200 i
```

On R2:

```
R2#Sh ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.1.1                 0             0 100 111 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
```

On R1:

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

R1(config)#router bgp 111
R1(config-router)#No auto-summary
R1(config-router)#network 1.0.0.0
R1(config-router)#neighbor 10.1.1.2 remote-as 200
R1(config-router)#neighbor 10.1.1.2 local-as 100
```

You should see the following console message:

```
%BGP-5-ADJCHANGE: neighbor 10.1.1.2 Up
```

The "neighbor 10.1.1.2 local-as 100" command on R1 tells R2 that it's in AS 100; R2 will see R1's real AS of 111 originating the route and then AS 100 was the AS that advertised it. Note on R1, it shows that prefix 2.0.0.0/8 was originated by AS 200 but the advertising AS to R1 was AS 100. They both see this invisible AS (AS100).

To verify the configuration:

On R1:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.1.2                 0             0 100 200 i
```

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.1.1                 0             0 100 111 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
```

Note: Both routers see this invisible AS 100.

Task 3

Configure R1 such that when R2 advertises network 2.0.0.0/8, the output of the "**show ip bgp**" command on R1 resembles to the following:

On R1:

```
R1#sh ip bgp b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.1.2                 0             0 200 i
```

On R1:

```
R1(config)#router bgp 111
R1(config-router)#neighbor 10.1.1.2 local-as 100 no-prepend
```

Note: The "no-prepend" option tells the router not to prepend AS 100 to the advertised prefixes. This will only affect R1's BGP table.

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
*> 1.0.0.0              0.0.0.0                  0         32768 i
*> 2.0.0.0              10.1.1.2                 0             0 200 i
```

On R2:

```
R2#Show ip bgp | b Network
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.1.1                 0             0 100 111 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
```

Note: The above configuration only affected R1.

Task 4

Configure R1 such that the output of the "**show ip bgp**" command on R2 is identical to the following:

On R2:

```
R2#Show ip bgp | b Network
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.1.1                 0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
```

On R1:

```
R1(config)#router bgp 111
R1(config-router)#neighbor 10.1.1.2 local-as 100 no-prepend replace-as
```

Note: The "replace-as" option instructs the local router not to prepend the real AS number.

On R2:

```
R2#Show ip bgp | b Network
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.1.1                 0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
```

On R1:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.1.2                 0             0 200 i
```

Note: The above configuration did not affect R1 at all.

Task 5

Configure R1 such that R2 can establish an EBGP peer session with R1 using AS 111 or 100.

On R1:

```
R1(config)#router bgp 111
R1(config-router)#neighbor 10.1.1.2 local-as 100 no-prepend replace-as dual-as
```

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.1.2                 0             0 200 i
```

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.1.1                 0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
```

To see R2's configuration:

On R2:

```
R2#Show run | s router bgp 200

router bgp 200
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
 no synchronization
 bgp log-neighbor-changes
 network 2.0.0.0
 neighbor 10.1.1.1 remote-as 100
 no auto-summary
```

To see R1's configuration:

On R1:

```
R1#Show run | s router bgp 111

router bgp 111
 no synchronization
 bgp log-neighbor-changes
 network 1.0.0.0
 neighbor 10.1.1.2 remote-as 200
 neighbor 10.1.1.2 local-as 100 no-prepend replace-as dual-as
 no auto-summary
```

To test the configuration:

Note: The output of the above show command on R2, reveals that R2 is establishing a peer session with R1 using AS 100; the following verifies that R2 can also establish a peer session with R1 using AS 111:

On R2:

```
R2(config)#router bgp 200
R2(config-router)#No auto-summary
R2(config-router)#netw 2.0.0.0
R2(config-router)#neighbor 10.1.1.1 remote-as 111
```

To verify the configuration:

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.1.1                 0             0 111 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
```

Note: The only difference here is the AS that originated and advertised 1.0.0.0/8 prefix.

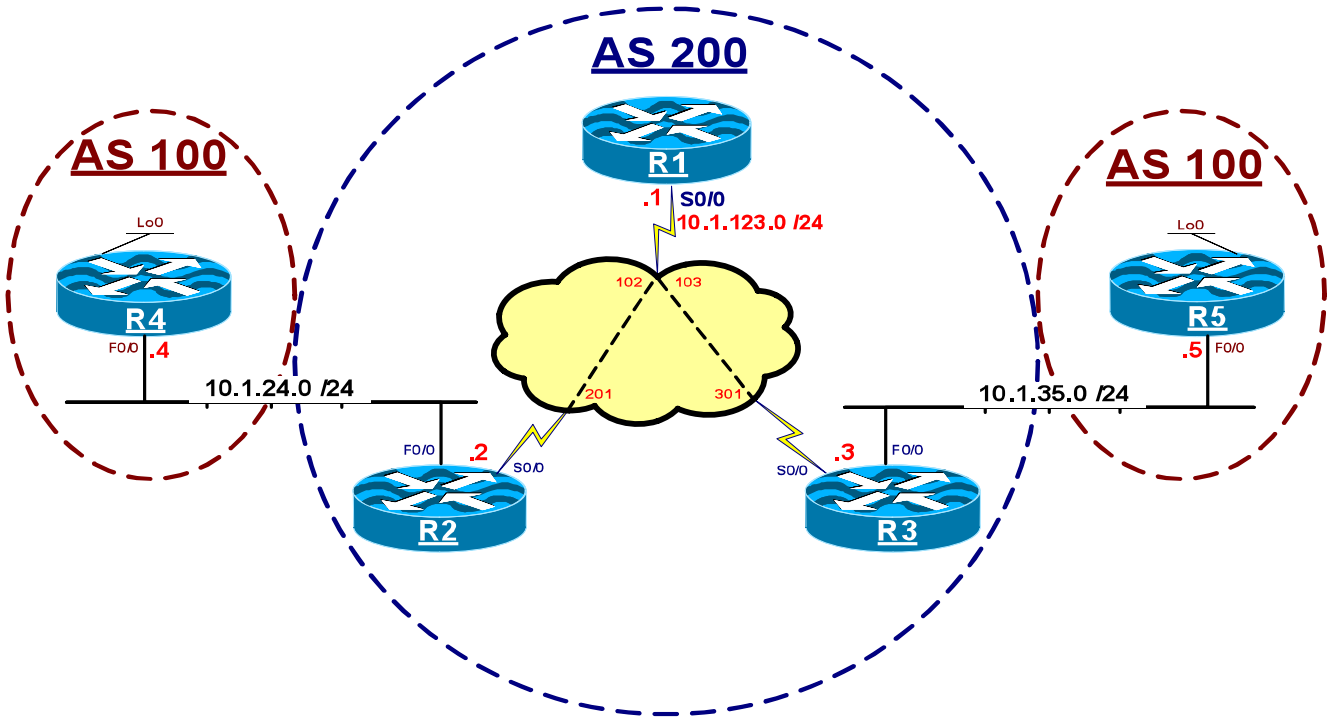
©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 6

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 22 – BGP Allow-as



Lab Setup:

- Configure Frame-Relay on R1, R2 and R3; this configuration should be done directly under the main interface. R2 and R3 should have a static frame-relay mapping to R1 and R1 should have static frame-relay mapping to R2 and R3.
- RIPv2 should be used to provide NLRI
- Configure the F0/0 interface of R2 and R4 in VLAN 24
- Configure the F0/0 interface R3 and R5 in VLAN 35
- Use the following chart for IP addressing and AS assignment.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

IP addressing & AS assignment:

Router	Interface / IP address	AS number
R1	**S0/0 = 10.1.123.1 /24**	**200**
R2	S0/0 = 10.1.123.2 /24 F0/0 = 10.1.24.2 /24	200
R3	**S0/0 = 10.1.123.3 /24** **F0/0 = 10.1.35.3 /24**	**200**
R4	F0/0 = 10.1.24.4 /24 Lo0 = 4.4.4.4 /8	100
R5	**F0/0 = 10.1.35.5 /24** **Lo0 = 5.5.5.5 /8**	**100**

Task 1

Configure the routers according to the above diagram, if this configuration is performed successfully, all routers should have network 4.0.0.0 /8 and 5.0.0.0 /8 in their BGP and routing table.

On R1:

```
R1(config)#router bgp 200
R1(config-router)#neighbor 10.1.123.2 remote-as 200
R1(config-router)#neighbor 10.1.123.3 remote-as 200
```

On R2:

```
R2(config)#router bgp 200
R2(config-router)#neighbor 10.1.123.1 remote-as 200
R2(config-router)#neighbor 10.1.123.3 remote-as 200
R2(config-router)#neighbor 10.1.24.4 remote-as 100
```

On R3:

```
R3(config)#router bgp 200
R3(config-router)#neighbor 10.1.123.1 remote-as 200
R3(config-router)#neighbor 10.1.123.2 remote-as 200
R3(config-router)#neighbor 10.1.35.5 remote-as 100
```

On R4:

```
R4(config)#router bgp 100
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config-router)#neighbor 10.1.24.2 remote-as 200
R4(config-router)#network 4.0.0.0
```

On R5:

```
R5(config)#router bgp 100
R5(config-router)#neighbor 10.1.35.3 remote-as 200
R5(config-router)#network 5.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i4.0.0.0          10.1.24.4                0    100      0 100 i
*>i5.0.0.0          10.1.35.5                0    100      0 100 i
```

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 4.0.0.0          10.1.24.4                0             0 100 i
*>i5.0.0.0          10.1.35.5                0    100      0 100 i
```

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i4.0.0.0          10.1.24.4                0    100      0 100 i
*> 5.0.0.0          10.1.35.5                0             0 100 i
```

On R4:

```
R4#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 4.0.0.0          0.0.0.0                  0         32768 i
```

Not:e R4 does not have prefix 5.0.0.0 /8 in its routing table, the question is did R2 advertise that prefix to R4? The output of the following show command can be used to verify:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2:

```
R2#Show ip bgp neighbor 10.1.24.4 advertised-routes | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i5.0.0.0          10.1.35.5                0    100      0 100 i

Total number of prefixes 1
```

Note: R2 has advertised prefix 5.0.0.0 /8 but this prefix came from AS 100 and since R4 is in the same AS (AS 100), it will discard that prefix, this is done by BGP as loop avoidance mechanism.

To fix this problem, R4 should be configured to allow prefix/es that have its own AS number in its AS, this is accomplished using the following configuration:

On R4:

```
R4(config)#router bgp 100
R4(config-router)#neighbor 10.1.24.2 allowas-in
```

To verify the configuration:

On R4:

```
R4#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 4.0.0.0          0.0.0.0                  0         32768 i
*> 5.0.0.0          10.1.24.2                              0 200 100 i

R4#Show ip route | inc B

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
B    5.0.0.0/8 [20/0] via 10.1.24.2, 00:01:11
```

Note: Prefix 5.0.0.0 /8 is in R4's BGP and routing table with AS number that matches its own. Since R5 has the same problem, it should also be configured the same way.

To verify R5's BGP table before configuration:

On R5:

```
R5#Show ip bgp | b Network
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 5.0.0.0          0.0.0.0                  0         32768 i
```

Note: Once again prefix 4.0.0.0 /8 is missing.

To fix the problem:

On R5:

```
R5(config)#router bgp 100
R5(config-router)#neighbor 10.1.35.3 allowas-in
```

To verify the configuration:

On R5:

```
R5#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 4.0.0.0          10.1.35.3                              0 200 100 i
*> 5.0.0.0          0.0.0.0                  0         32768 i
```

Task 2

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh