EIGRP Routing Protocol WorkBook 22 scenarios

Redouane MEDDANE

Table of Contents

Lab 1: EIGRP Different AS Distribute-List Offset List 2
Lab 2: Redistribution between two AS EIGRP Variance and Offset-List 10
Lab 3: EIGRP redistribution between two AS and the next-hop field 20
Lab 4: EIGRP route filtering using prefix-list and route-map 25
Lab 5: EIGRP Route Filtering using Route-Map Prefix-List ACL 30
Lab 6: EIGRP Route Leaking Maximum hop-count 38
Lab 7: EIGRP Default Route Summarization Authentication 47
Lab 8: EIGRP Named Mode for IPv4 and IPv6 56
Lab 9: EIGRP named mode over IPv6 66
Lab 10: EIGRP Named Mode Authentication 75
Lab 11: EIGRP Stub Routing with Named Mode 81
Lab 12: Maximum-prefix limit and Warning-Only Mode 87
Lab 13: Two EIGRP AS number and two equal paths to the same destination 91
Lab 14: Feasible Successor issues 96
Lab 15: EIGRP Summary and Leak Map 101
Lab 18: show ip eigrp topology all-links subcommand 104
Lab 19: Successor and Feasible Successor 110
Lab 20: Next-hop field between two AS eigrp IPv6 117
Lab 21: Filter EIGRP external routes 119
Lab 22: EIGRP Stub Feature 122

Lab 1: EIGRP Different AS Distribute-List Offset List

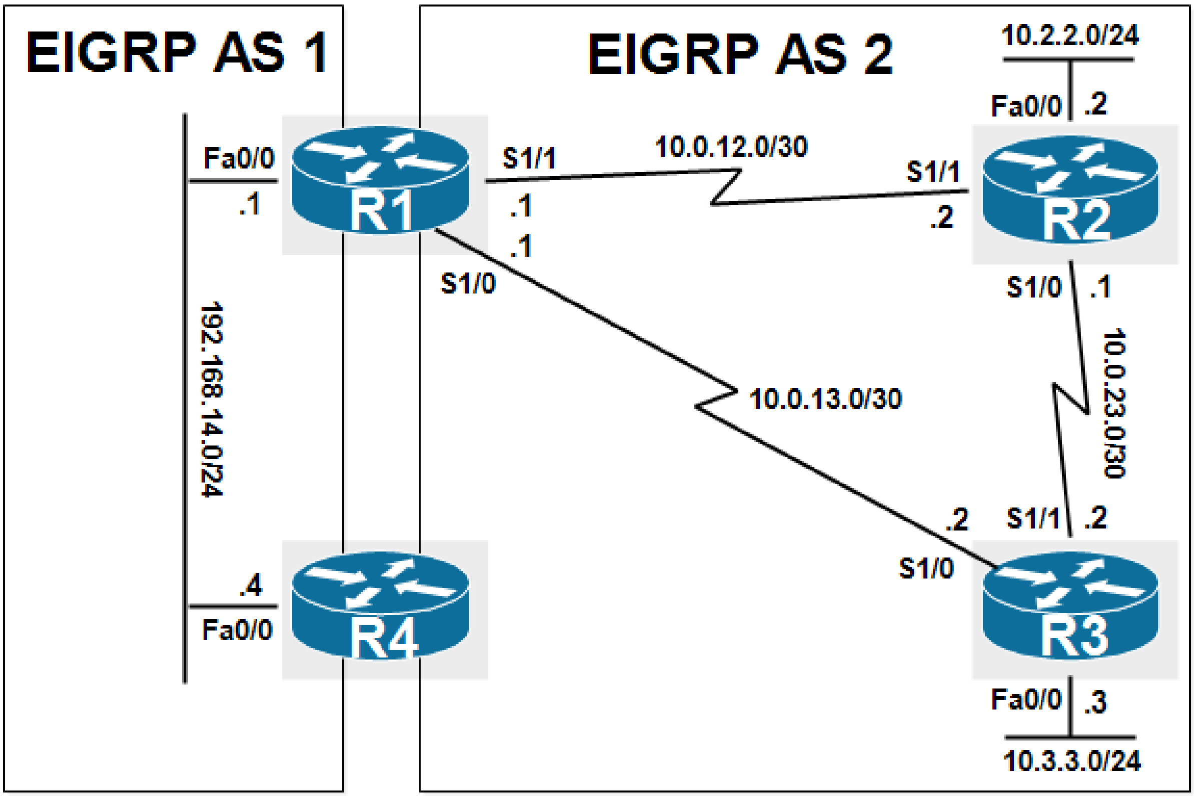

Task 1
Configure hostnames and IP addressing on all routers as illustrated in the network topology.

Task 2
Configure EIGRP for AS 1 as shown in the topology. R4 should be configured as an EIGRP stub router. R4 should NEVER advertise any routes. In addition to this, ensure that router R4 will only receive a default route from R1 even if external routes are redistributed into EIGRP 1.

Task 3
Configure EIGRP for AS 2 as illustrated in the topology. Verify your configuration using the appropriate commands for EIGRP.

Task 4
Configure EIGRP so that R4 can reach all other routers in the network and vice-versa. Ensure that only the 192.168.14.0/24 is allowed into the topology table for EIGRP 2. Verify your configuration and also ping to and from R4 from the 10.2.2.0/24 and 10.3.3.0/24 subnets.

Task 5
Assume that the WAN link between R1 and R3 is slow and should only be used when the WAN link between R1 and R2 is down. However, an EIGRP neighbor relationship should still be maintained across this link. Configure EIGRP so that neither routers R1 nor R3 use this link unless the WAN link between R1 and R2 is down. You are only allowed to configure R3.

Task 1:

Basic configuration of all routers:

R1

```
interface FastEthernet0/0
 ip address 192.168.14.1 255.255.255.0
 no shutd
!
interface Serial1/1
 ip address 10.0.12.1 255.255.255.252
 no shutd
!
interface Serial1/0
 ip address 10.0.13.1 255.255.255.252
 no shutd
```

R2

```
interface FastEthernet0/0
 ip address 10.2.2.2 255.255.255.0
 no shutd
!
interface Serial1/1
 ip address 10.0.12.2 255.255.255.252
 no shutd
!
interface Serial1/0
 ip address 10.0.23.1 255.255.255.252
 no shutd
```

R3

```
interface FastEthernet0/0
 ip address 10.3.3.3 255.255.255.0
 no shutd
!
interface Serial1/0
 ip address 10.0.13.2 255.255.255.252
 no shutd
!
interface Serial1/1
 ip address 10.0.23.2 255.255.255.252
 no shutd
```

R4

```
interface fastethernet 0/0
 ip address 192.168.14.4 255.255.255.0
 no shutdown
```

Task 2:

R4 should never advertise any routes therefore we will use the stub feature and the receive-only option.
In addition to this, we will ensure that router R4 will only receive a default route from R1 and filter all other routes.

To advertise a defaut route to R4, we can use the ip summary-address 0.0.0.0 0.0.0.0 command and to prevent all other routes we use the distribute-list command combined with the prefix-list.

The sequence 5 of the prefix-list called DEFAULT-ROUTE will match and allow the default-route.
All routes that do not match the prefix-list will be denied with an implicit deny.

R1:

```
ip prefix-list DEFAULT-ROUTE seq 5 permit 0.0.0.0/0
!
router eigrp 1
 network 192.168.14.0 0.0.0.255
 distribute-list prefix DEFAULT-ROUTE out fastethernet 0/0
!
interface fastethernet 0/0
 ip summary-address eigrp 1 0.0.0.0 0.0.0.0
```

R4

```
router eigrp 1
 network 192.168.14.0 0.0.0.255
 eigrp stub receive-only
```

The show ip eigrp neighbors detail shown that there are no routes advertised to the neighbor 192.168.14.4 and R1 suppresses the queries toward R4 because R4 is configured as a stub with receive-only option.

```
R1#show ip eigrp neighbors detail
EIGRP-IPv4 Neighbors for AS(1)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                   (sec)         (ms)       Cnt

Num
0   192.168.14.4            Fa0/0                    11 00:00:08  412  3708  0  1
   Version 11.0/2.0, Retrans: 1, Retries: 0
   Topology-ids from peer - 0
   Receive-Only Peer Advertising (No) Routes
   Suppressing queries
```

The show eigrp protocols shown that R4 is a stub router with receive-only option:

```
R4#show eigrp protocols
EIGRP-IPv4 Protocol for AS(1)
  Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  NSF-aware route hold timer is 240
  Router-ID: 192.168.14.4
  Stub, receive-only
  Topology : 0 (base)
    Active Timer: 3 min
    Distance: internal 90 external 170
    Maximum path: 4
    Maximum hopcount 100
    Maximum metric variance 1
```

R4 receives only one EIGRP route which is the default route from R1:

```
R4#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is 192.168.14.1 to network 0.0.0.0

D*    0.0.0.0/0 [90/30720] via 192.168.14.1, 00:02:50, FastEthernet0/0
R4#
```

Task 3:

Configuration of EIGRP AS 2:

R1

```
router eigrp 2
 no auto-summary
 network 10.0.12.0 0.0.0.3
 network 10.0.13.0 0.0.0.3
```

R2

```
router eigrp 2
 no auto-summary
 network 10.2.2.2 0.0.0.0
 network 10.0.12.2 0.0.0.0
 network 10.0.23.1 0.0.0.0
```

R3

```
router eigrp 2
 network 10.3.3.3 0.0.0.0
 network 10.0.13.2 0.0.0.0
 network 10.0.23.2 0.0.0.0
 no auto-summary
```

Task 4:

To ensure so that R4 can reach all other routes in the network and vice-versa.We should do the redistribution between EIGRP AS 1 and AS2 .
The requirement is that only the 192.168.14.0/24 is allowed in EIGRP AS 2.So we should configure filtering when redistributing so that only the 192.168.14.0/24 subnet is allowed into EIGRP AS 2.

First we define a prefix-list called LAN-NETWORK that matches the subnet 192.168.14.0/24.
Then we create a route-map called EIGRP1-to-EIGRP2 which allows all subnets matched by the prefix-list LAN-NETWORK and deny all other routes (sequence 20 of the route-map).

It is important to remember that a seed metric is not required when redistributing between different EIGRP processes.

R1

```
router eigrp 1
 redistribute eigrp 2
```

```
!
 ip prefix-list LAN-NETWORK seq 5 permit 192.168.14.0/24
!
route-map EIGRP1-to-EIGRP2 permit 10
 match ip address prefix-list LAN-NETWORK
!
route-map EIGRP1-to-EIGRP2 deny 20
!
router eigrp 2
 redistribute eigrp 1 route-map EIGRP1-to-EIGRP2
```

We can see that the external route toward 192.168.14.0/24 exists in the routing tables on R2 R3.

```
R2#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 8 subnets, 3 masks
D        10.0.13.0/30 [90/2681856] via 10.0.23.2, 00:09:56, Serial1/0
                      [90/2681856] via 10.0.12.1, 00:09:56, Serial1/1
D        10.3.3.0/24 [90/2172416] via 10.0.23.2, 00:00:10, Serial1/0
D EX  192.168.14.0/24 [170/2172416] via 10.0.12.1, 00:01:33, Serial1/1
R2#
```

```
R3#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 8 subnets, 3 masks
D        10.0.12.0/30 [90/2681856] via 10.0.23.1, 00:10:52, Serial1/1
                      [90/2681856] via 10.0.13.1, 00:10:52, Serial1/0
D        10.2.2.0/24 [90/2172416] via 10.0.23.1, 00:10:52, Serial1/1
D EX  192.168.14.0/24 [170/2172416] via 10.0.13.1, 00:02:28, Serial1/0
R3#
```

Because the filtering configured on R1 with the distribute-list command ,R4 did'nt learn external routes from EIGRP AS 2:

```
R4#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is 192.168.14.1 to network 0.0.0.0

D*    0.0.0.0/0 [90/30720] via 192.168.14.1, 00:08:09, FastEthernet0/0
R4#
```

R4 can reach the LANs subnet of R2 and R3 ,10.2.2.2 and 10.3.3.3 respectively:

```
R4#ping 10.2.2.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 72/103/132 ms
R4#
R4#ping 10.3.3.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/74/132 ms
R4#
```

Task 5:

The link between R1 and R3 should only be used when the link between R1 and R2 fails. Configure EIGRP so that neither routers R1 nor R3 use this link unless the link between R1 and R2 fails.

To complete this requirement we should configure an inbound and outbound offset-list on R3 so that routes received or sent via the link between R1 and R3 are not preferred:

Before doing the task let's check the topology table of R3:

```
R3#show ip eigrp topology
EIGRP-IPv4 Topology Table for AS(2)/ID(10.3.3.3)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 10.2.2.0/24, 1 successors, FD is 2172416
        via 10.0.23.1 (2172416/28160), Serial1/1
P 10.3.3.0/24, 1 successors, FD is 28160
        via Connected, FastEthernet0/0
P 192.168.14.0/24, 1 successors, FD is 2172416
        via 10.0.13.1 (2172416/28160), Serial1/0
P 10.0.13.0/30, 1 successors, FD is 2169856
        via Connected, Serial1/0
P 10.0.23.0/30, 1 successors, FD is 2169856
        via Connected, Serial1/1
P 10.0.12.0/30, 2 successors, FD is 2681856
        via 10.0.13.1 (2681856/2169856), Serial1/0
        via 10.0.23.1 (2681856/2169856), Serial1/1
```

R3

```
router eigrp 2
 offset-list 0 in 1000000 serial 1/0
 offset-list 0 out 1000000 serial 1/0
```

The inbound and the outbound offset-list on R3 ensure that the metrics of a routes received or sent via the link between R1 and R3 are increased by adding the value 1000000 **so the metrics via the link R1--R2 will be better:**

offset-list 0 with access-list 0 is the equivalent of permit any and it is applied to all routes .

Let's verify the routing table of R1 and R3:

On R1 All EIGRP routes pass through R2 (10.0.12.2):

```
R1#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is 0.0.0.0 to network 0.0.0.0

D*    0.0.0.0/0 is a summary, 00:28:30, Null0
      10.0.0.0/8 is variably subnetted, 7 subnets, 3 masks
D        10.0.23.0/30 [90/2681856] via 10.0.12.2, 00:03:00, Serial1/1
D        10.2.2.0/24 [90/2172416] via 10.0.12.2, 00:03:00, Serial1/1
D        10.3.3.0/24 [90/2684416] via 10.0.12.2, 00:03:00, Serial1/1
R1#
```

On R3 All EIGRP routes pass through R2 (10.0.23.1):

```
R3#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 8 subnets, 3 masks
D        10.0.12.0/30 [90/2681856] via 10.0.23.1, 00:00:03, Serial1/1
D        10.2.2.0/24 [90/2172416] via 10.0.23.1, 00:00:03, Serial1/1
D EX  192.168.14.0/24 [170/2684416] via 10.0.23.1, 00:00:03, Serial1/1
R3#
```

Finally notice the effect of the offset-list:
Let's take an example of the network 192.168.14.0/24,before configuring the offset-list the best metric is via R1 10.0.13.1 and it is equal to 2172416 **as shown above in the topology table of R3 before offset-list ,now after configuring the offset-list ,the Metric via R1 between the parenthesis for the network 192.168.14.0/24 which is equal to 1002195456 is calculated by adding the value configured in the offset-list** 1000000+2172416=3172416 **,Therfore the metric**

via R2 2684416 is less and better than the metric via R1 3172416, **this metric** 3172416 **is shown by the show ip eigrp topo all command which display all routes eigrp successor or no successor ,Feasible successor or not Feasible successor:**

```
R3#show ip eigrp topo all
EIGRP-IPv4 Topology Table for AS(2)/ID(10.3.3.3)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 10.2.2.0/24, 1 successors, FD is 2172416, serno 8
        via 10.0.23.1 (2172416/28160), Serial1/1
        via 10.0.13.1 (3684416/3172352), Serial1/0
P 10.3.3.0/24, 1 successors, FD is 28160, serno 6
        via Connected, FastEthernet0/0
P 192.168.14.0/24, 1 successors, FD is 2684416, serno 10
        via 10.0.23.1 (2684416/2172416), Serial1/1
        via 10.0.13.1 (3172416/1028096), Serial1/0
P 10.0.13.0/30, 1 successors, FD is 2169856, serno 4
        via Connected, Serial1/0
P 10.0.23.0/30, 1 successors, FD is 2169856, serno 5
        via Connected, Serial1/1
        via 10.0.13.1 (4193856/3681792), Serial1/0
P 10.0.12.0/30, 1 successors, FD is 2681856, serno 7
        via 10.0.23.1 (2681856/2169856), Serial1/1
        via 10.0.13.1 (3681856/3169792), Serial1/0

R3#
```

Lab 2: Redistribution between two AS EIGRP Variance and Offset-List

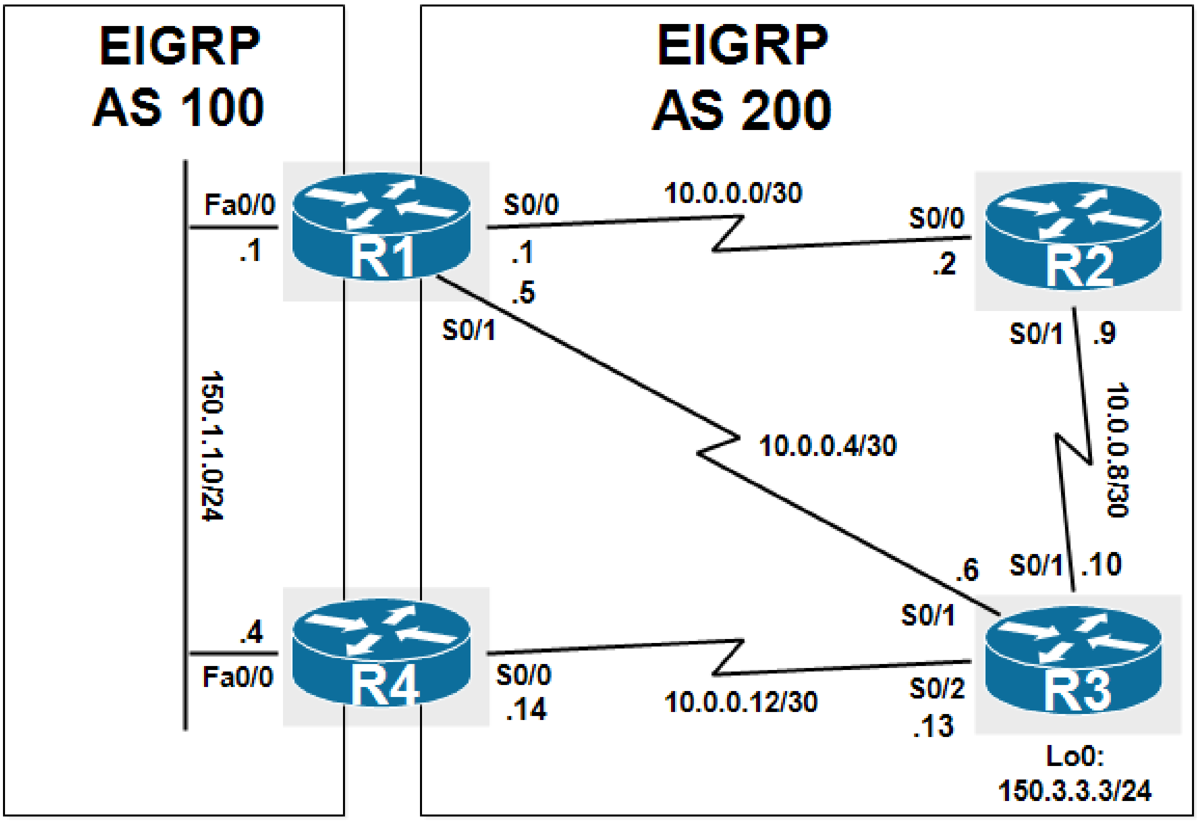

Task 1

Configure hostnames and IP addressing on all routers as illustrated in the network topology.

Task 2

Configure EIGRP AS 100 for the 150.1.1.0/24 subnet between R1 and R4. Configure EIGRP AS 200 as illustrated in the topology.

Task 3

Redistribute between EIGRP 100 and EIGRP 200. Use default metrics in your configuration. Verify that R1 has three paths to the 150.3.3.0/24 subnet.

Task 4

Without modifying interface bandwidth or delay values, configure routing and path control for EIGRP on R1 as shown below. Use the most simple configuration in your solution:

1. R1 should first prefer the path via R4 (FastEthernet 0/0) to reach the 150.3.3.0/24 subnet
2. If the path via R4 is down, R1 should prefer the path via R2 (Serial 0/0)
3. If the path via R4 and via R2 is down, R1 should use the direct path to R3 (Serial 0/1)

Task 5

Without modifying interface bandwidth or delay values, configure routing and path control for EIGRP on R1 as shown below. Use the most simple configuration in your solution:

1. R3 should first prefer the path via R4 (Serial 0/2) to reach the 150.1.1.0/24 subnet
2. If the path via R4 is down, R3 should prefer the path via R2 (Serial 0/1)
3. If the paths via R4 R2 are down, R3 should use the direct path to R1 (Serial 1/0)

Task 6

Configure routing on R3 so that all three paths are installed into the routing table; however, ONLY a single path - the path via R4 - should actually be used to forward packets destined to that subnet. The other two paths should not be used while this path is up. The path with the next-best route metric should then be used when the primary path is unavailable.

Task 1:

Basic configuration of all routers:

R1

```
interface FastEthernet0/0
 ip address 150.1.1.1 255.255.255.0
 no shutd
!
interface Serial0/0
 ip address 10.0.0.1 255.255.255.252
 no shutd
!
interface Serial0/1
 ip address 10.0.0.5 255.255.255.252
 no shutd
```

R2

```
interface Serial0/0
 ip address 10.0.0.2 255.255.255.252
 no shutd
!
interface Serial0/1
 ip address 10.0.0.9 255.255.255.252
 no shutd
```

R3

```
interface FastEthernet0/0
 ip address 150.3.3.3 255.255.255.0
 no shutd
!
interface Serial0/0
 ip address 10.0.0.6 255.255.255.252
 no shutd
!
interface Serial0/1
 ip address 10.0.0.10 255.255.255.252
```

```
 no shutd
!
interface Serial0/2
 ip address 10.0.0.13 255.255.255.252
 no shutd
```

R4

```
interface fastethernet 0/0
 ip address 150.1.1.4 255.255.255.0
 no shutdown
!
interface Serial0/0
 ip address 10.0.0.14 255.255.255.252
 no shutd
```

Task 2:

R1

```
router eigrp 100
 no auto-summary
 network 150.1.1.0 0.0.0.255
!
router eigrp 200
 network 10.0.0.0 0.0.0.3
 network 10.0.0.4 0.0.0.3
 no auto-summary
```

R2

```
router eigrp 200
 network 10.0.0.0 0.0.0.3
 network 10.0.0.8 0.0.0.3
```

R3

```
router eigrp 200
 network 10.0.0.4 0.0.0.3
 network 10.0.0.8 0.0.0.3
 network 10.0.0.12 0.0.0.3
 network 150.3.3.0 0.0.0.255
```

R4

```
router eigrp 200
 network 10.0.0.12 0.0.0.3
!
router eigrp 100
 no auto-summary
```

```
 network 150.1.1.0 0.0.0.255
```

Task 3:

R1 and R4 redistribute between the two EIGRP Autonomous Systems:

```
R1(config)#router eigrp 100
R1(config-router)#redistribute eigrp 200
R1(config)#router eigrp 200
R1(config-router)#redistribute eigrp 100
```

```
R4(config-router)#redistribute eigrp 200
R4(config-router)#exit
R4(config)#router eigrp 200
R4(config-router)#redistribute eigrp 100
```

show ip eigrp topology 150.3.3.0 255.255.255.0 command displays three routes, via R2 ,via R3 and via R4:

```
R1(config-router)#do show ip eigrp topology 150.3.3.0 255.255.255.0
IP-EIGRP (AS 100): Topology entry for 150.3.3.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2195456
  Routing Descriptor Blocks:
  10.0.0.6, from Redistributed, Send flag is 0x0
      Composite metric is (2195456/0), Route is External
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 21000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
      External data:
        Originating router is 150.1.1.1 (this system)
        AS number of route is 200
        External protocol is EIGRP, external metric is 2195456
        Administrator tag is 0 (0x00000000)
  150.1.1.4 (FastEthernet0/0), from 150.1.1.4, Send flag is 0x0
      Composite metric is (2221056/2195456), Route is External
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 22000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
      External data:
        Originating router is 150.1.1.4
        AS number of route is 200
        External protocol is EIGRP, external metric is 2195456
        Administrator tag is 0 (0x00000000)
IP-EIGRP (AS 200): Topology entry for 150.3.3.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2195456
  Routing Descriptor Blocks:
```

```
  10.0.0.6 (Serial0/1), from 10.0.0.6, Send flag is 0x0
      Composite metric is (2195456/281600), Route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 21000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
  10.0.0.2 (Serial0/0), from 10.0.0.2, Send flag is 0x0
      Composite metric is (2707456/2195456), Route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 41000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
```

And the best route installed in the routing table is through R3 which is the internal route learned from EIGRP 200 with the next-hop 10.1.1.6:

```
R1#show ip route  150.3.3.0 255.255.255.0
Routing entry for 150.3.3.0/24
  Known via "eigrp 200", distance 90, metric 2195456, type internal
  Redistributing via eigrp 100, eigrp 200
  Advertised by eigrp 100
  Last update from 10.0.0.6 on Serial0/1, 00:04:08 ago
  Routing Descriptor Blocks:
  * 10.0.0.6, from 10.0.0.6, 00:04:08 ago, via Serial0/1
      Route metric is 2195456, traffic share count is 1
      Total delay is 21000 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
```

Task 4:

To ensure that R1 prefers the path through R4 ,we will adjust the Administrative Distance by increasing the AD of internal and external routes learned from EIGRP 200 to be higher than the AD of external routes learned from EIGRP 100:

```
R1(config)#router eigrp 200
R1(config-router)#distance eigrp 180 200
```

Now we can see that the next-hop is R4 and the route has an AD of 170 and it's an external redistributed via eigrp 200:

```
R1(config-router)#do show ip route  150.3.3.0 255.255.255.0
Routing entry for 150.3.3.0/24
  Known via "eigrp 100", distance 170, metric 2221056, type external
  Redistributing via eigrp 100, eigrp 200
  Advertised by eigrp 200
  Last update from 150.1.1.4 on FastEthernet0/0, 00:00:05 ago
  Routing Descriptor Blocks:
```

```
  * 150.1.1.4, from 150.1.1.4, 00:00:05 ago, via FastEthernet0/0
      Route metric is 2221056, traffic share count is 1
      Total delay is 22000 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2
```

If the path via R4 is down, R1 should prefer the path via R2 so if the path via R4 fails , R1 goes through R2 instead of R3,thus we will increase the metric for routes learned from R3, we can see at the begining in the do show ip eigrp topology 150.3.3.0 255.255.255.0 command that:
-the metric to reach 150.3.3.0 through R3 is 2195456
-the metric to reach 150.3.3.0 through R2 is 2707456
The path via R3 is better:
to ensure R2 will become the preferred route if R4 fails we will increase the metric for the routes learned from R3 by using offset-list:

```
R1(config)#router eigrp 200
R1(config-router)#offset-list 0 in 1000000000 serial 0/1
```

Notice now:
-the metric to reach 150.3.3.0 through R3 is 1002195456(1000000000+2195456)
-the metric to reach 150.3.3.0 through R2 is 2707456
the path through R2 is better

```
R1#show ip eigrp topology 150.3.3.0 255.255.255.0 | include
metric|10.0.0.2|10.0.0.6
      Composite metric is (2221056/0), Route is External
      Vector metric:
        External protocol is EIGRP, external metric is 2221056
  10.0.0.6 (Serial0/1), from 10.0.0.6, Send flag is 0x0
      Composite metric is (1002195456/1000281600), Route is Internal
      Vector metric:
  10.0.0.2 (Serial0/0), from 10.0.0.2, Send flag is 0x0
      Composite metric is (2707456/2195456), Route is Internal
      Vector metric:
```

Task 5:

1. **R3 should prefer the path via R4 to reach the 150.1.1.0/24**
2. **If the path via R4 is down, R3 should prefer the path via R2**
3. **If the paths via R4 R2 are down, R3 should use the direct path to R1**

We have load balancing on R3 to reach 150.1.1.0, one path through R4 and another path through R1:

```
R3(config-router)# do show ip route 150.1.1.0 255.255.255.0
Routing entry for 150.1.1.0/24
  Known via "eigrp 200", distance 170, metric 2195456, type external
  Redistributing via eigrp 200
  Last update from 10.0.0.5 on Serial0/0, 00:12:46 ago
  Routing Descriptor Blocks:
  * 10.0.0.14, from 10.0.0.14, 00:12:46 ago, via Serial0/2
      Route metric is 2195456, traffic share count is 1
      Total delay is 21000 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
    10.0.0.5, from 10.0.0.5, 00:12:46 ago, via Serial0/0
      Route metric is 2195456, traffic share count is 1
      Total delay is 21000 microseconds, minimum bandwidth is 1544 Kbit
```

```
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
```

To ensure that the best path is through R4 and if R4 fails ,R3 prefers the path through R2 ,we will just increase the metric of R1 to be higher than the metric of R2 , we have already the route via R4 installed in the routing table (load balancing with the path R1):
before making any changes notice the metric through R1 10.0.0.5 is 2195456
and the metric through R2 10.0.0.9 is 2707456
R1 is better:

```
R3#show ip eigrp 200 topology 150.1.1.0 255.255.255.0 | include
metric|10.0.0.5|10.0.0.9|10.0.0.14
  10.0.0.5 (Serial0/0), from 10.0.0.5, Send flag is 0x0
      Composite metric is (2195456/281600), Route is External
      Vector metric:
        External protocol is EIGRP, external metric is 0
  10.0.0.14 (Serial0/2), from 10.0.0.14, Send flag is 0x0
      Composite metric is (2195456/281600), Route is External
      Vector metric:
        External protocol is EIGRP, external metric is 0
  10.0.0.9 (Serial0/1), from 10.0.0.9, Send flag is 0x0
      Composite metric is (2707456/2195456), Route is External
      Vector metric:
        External protocol is EIGRP, external metric is 0
```

We will increase the metric of the routes learned from R1 by configuring an offset-list as follow:

```
R3(config)#router eigrp 200
R3(config-router)# offset-list 0 in 604544 s0/0
```

offset-list 0 with access-list 0 is the equivalent of permit any and it is applied to all routes .
Now the metric through R1 is 2800000=604544+2195456
The metric through R2 is 2707456
R2 is better than R1:

```
R3#show ip eigrp 200 topology 150.1.1.0 255.255.255.0 | include
metric|10.0.0.5|10.0.0.9|10.0.0.14
  10.0.0.14 (Serial0/2), from 10.0.0.14, Send flag is 0x0
      Composite metric is (2195456/281600), Route is External
      Vector metric:
        External protocol is EIGRP, external metric is 0
  10.0.0.5 (Serial0/0), from 10.0.0.5, Send flag is 0x0
      Composite metric is (2800000/886144), Route is External
      Vector metric:
        External protocol is EIGRP, external metric is 0
  10.0.0.9 (Serial0/1), from 10.0.0.9, Send flag is 0x0
      Composite metric is (2707456/2195456), Route is External
      Vector metric:
        External protocol is EIGRP, external metric is 0
```

```
R3#show ip route eigrp
     10.0.0.0/30 is subnetted, 4 subnets
D       10.0.0.0 [90/2681856] via 10.0.0.9, 00:02:54, Serial0/1
     150.1.0.0/24 is subnetted, 1 subnets
D EX    150.1.1.0 [170/2195456] via 10.0.0.14, 00:02:54, Serial0/2
```

So R3 prefers the path through R4 (10.0.0.14) to reach 150.1.1.0 with the best metric 2195456, if R4 fails, R3 will use the path via R2 with the metric 2707456 and if R2 fails R3 will use the path via R1 with the metric 2800000.

Task 6:

Configure R3 so that it will install unequal load balancing for the three path show in the show ip eigrp 200 topology 150.1.1.0 255.255.255.0 command by using variance command:

```
R3(config)#router eigrp 200
R3(config-router)#variance 2
R3#show ip route eigrp
     10.0.0.0/30 is subnetted, 4 subnets
D       10.0.0.0 [90/2681856] via 10.0.0.9, 00:02:31, Serial0/1
     150.1.0.0/24 is subnetted, 1 subnets
D EX    150.1.1.0 [170/2195456] via 10.0.0.14, 00:00:06, Serial0/2
                  [170/2800000] via 10.0.0.5, 00:00:06, Serial0/0
```

Ok the result is not expected because we have only two path to perform unequal load balancing via R4 and via R1,the challenge tells us to install all the three path, why this behavior ? first by definition ,when configuring variance command ,any FS routes whose calculated metric is less than or equal to the product of variance times FD are added to the IP routing table ,Routes that are neither successor nor feasible successor can never be added to the IP routing table, regardless of the variance setting.
In this case before configuring the variance command on R3, the only feasible successor is R1 because the Reported Distance of R1 is 886144 less than the Feasible Distance of R4 which is 2195456,the path via R2 is not included because the Reported Distance of R2 is 2195456 and it is equal to the Feasible Distance of R4 which is 2195456 so the feasibility condition is not respected,R2 is not a FS:
To ensure that R2 meets the the feasibility condition in order to install this path for unequal load balancing, configure an offset-list on R4 :
with an offset-list configured with a value of 4 ,R4 will add it to the FD and RD and advertises out to its neighbor R3 ,(2195456/281600) becomes (2195460/281604)
in other word after configuring the offset-list , R4 calculates the new Feasible Distance FD and the Reported Distance RD as follow:

The new FD =2195456+4=2195460
The new RD =281600+4=281604

```
R4(config-router)#router eigrp 200
R4(config-router)#offset-list 0 out 4
```

Notice now R2 now meets the feasibility condition, the Reported Distance of R2 2195456 is less than the FD of R4 2195460:

```
R3#show ip eigrp 200 topology 150.1.1.0 255.255.255.0 | include
metric|10.0.0.5|10.0.0.9|10.0.0.14
  10.0.0.14 (Serial0/2), from 10.0.0.14, Send flag is 0x0
      Composite metric is (2195460/281604), Route is External
      Vector metric:
        External protocol is EIGRP, external metric is 0
  10.0.0.5 (Serial0/0), from 10.0.0.5, Send flag is 0x0
      Composite metric is (2800000/886144), Route is External
      Vector metric:
        External protocol is EIGRP, external metric is 0
  10.0.0.9 (Serial0/1), from 10.0.0.9, Send flag is 0x0
      Composite metric is (2707456/2195456), Route is External
      Vector metric:
```

```
        External protocol is EIGRP, external metric is 0
```

We should have three path in the routing table of R3 to perform unequal load balancing:

```
R3#show ip route eigrp
     10.0.0.0/30 is subnetted, 4 subnets
D       10.0.0.0 [90/2681856] via 10.0.0.9, 00:09:45, Serial0/1
     150.1.0.0/24 is subnetted, 1 subnets
D EX    150.1.1.0 [170/2195460] via 10.0.0.14, 00:05:11, Serial0/2
                  [170/2707456] via 10.0.0.9, 00:05:11, Serial0/1
                  [170/2800000] via 10.0.0.5, 00:05:11, Serial0/0
```

```
R3#show ip route 150.1.1.1 255.255.255.0
Routing entry for 150.1.1.0/24
  Known via "eigrp 200", distance 170, metric 2195460, type external
  Redistributing via eigrp 200
  Last update from 10.0.0.9 on Serial0/1, 00:12:16 ago
  Routing Descriptor Blocks:
  * 10.0.0.14, from 10.0.0.14, 00:12:16 ago, via Serial0/2
      Route metric is 2195460, traffic share count is 240
      Total delay is 21000 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
    10.0.0.9, from 10.0.0.9, 00:12:16 ago, via Serial0/1
      Route metric is 2707456, traffic share count is 195
      Total delay is 41000 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2
    10.0.0.5, from 10.0.0.5, 00:12:16 ago, via Serial0/0
      Route metric is 2800000, traffic share count is 188
      Total delay is 44615 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
```

R3 installs the three paths (via R4 ,via R2 and via R1) to reach 150.1.1.0 with unequal load balancing.Notice the traffic share count keyword in the show ip route 150.1.1.1 255.255.255.0 command for all three paths , it is equal to a non-zero :

240 packets routed to R4
195 packets routed to R2
188 packets routed to R1

We should install ONLY a single path - the path via R4 - this path will be used to forward packets destined to 150.1.1.0. The rest of two paths should not be used when the path through R4 is up. The path with the route with the best metric which is through R2 should then be used when the path via R4 fails:

```
R3(config)#router eigrp 200
R3(config-router)#traffic-share min across-interfaces
*Mar  1 02:07:28.367: %FIB-4-UNEQUAL: Range of unequal path weightings too large
for prefix 150.1.1.0/24. Some available paths may not be used.
```

```
R3#show ip route eigrp
     10.0.0.0/30 is subnetted, 4 subnets
D       10.0.0.0 [90/2681856] via 10.0.0.9, 00:00:39, Serial0/1
     150.1.0.0/24 is subnetted, 1 subnets
D EX    150.1.1.0 [170/2195460] via 10.0.0.14, 00:00:39, Serial0/2
                  [170/2707456] via 10.0.0.9, 00:00:39, Serial0/1
```

```
                    [170/2800000] via 10.0.0.5, 00:00:39, Serial0/0
```

Notice now the traffic share count is equal to 0 for R2(10.0.0.9) and R1(10.0.0.5):

```
R3#show ip route 150.1.1.1 255.255.255.0
Routing entry for 150.1.1.0/24
  Known via "eigrp 200", distance 170, metric 2195460, type external
  Redistributing via eigrp 200
  Last update from 10.0.0.9 on Serial0/1, 00:01:03 ago
  Routing Descriptor Blocks:
  * 10.0.0.14, from 10.0.0.14, 00:01:03 ago, via Serial0/2
      Route metric is 2195460, traffic share count is 1
      Total delay is 21000 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
    10.0.0.9, from 10.0.0.9, 00:01:03 ago, via Serial0/1
      Route metric is 2707456, traffic share count is 0
      Total delay is 41000 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2
    10.0.0.5, from 10.0.0.5, 00:01:03 ago, via Serial0/0
      Route metric is 2800000, traffic share count is 0
      Total delay is 44615 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
```

Lab 3: EIGRP redistribution between two AS and the next-hop field

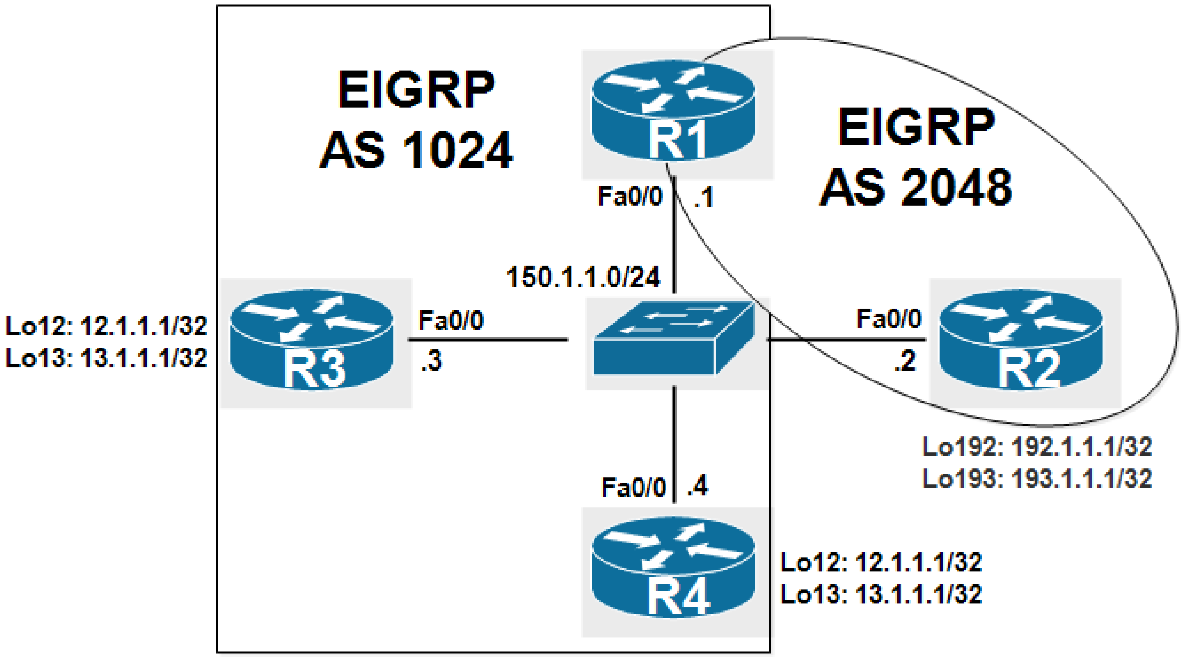

Task 1

Configure hostnames and IP addressing on all routers as illustrated in the network topology.

Task 2

Configure EIGRP for AS 1024 on R1, R3, and R4 as illustrated in the topology. EIGRP should be configured to use Unicast packets instead of Multicast packets.

Task 3

Configure EIGRP for AS 2048 on R1 and R2 as illustrated in the topology. EIGRP should be configured to use Unicast packets instead of Multicast packets.

Task 4

Configure and advertise the following subnets via EIGRP using the following metrics:

1. R2: Loopback192 - 192.1.1.1/32: Redistribution Metric - 100,000 - 1,000 - 255 - 1 - 1500
2. R2: Loopback193 - 193.1.1.1/32: Redistribution Metric - 10,000 - 500 - 255 - 1 - 1500

1. R3: Loopback12 - 12.1.1.1/32: Redistribution Metric - 30,000 - 300 - 255 - 1 - 1500
2. R3: Loopback13 - 13.1.1.1/32: Redistribution Metric - 40,000 - 400 - 255 - 1 - 1500

1. R4: Loopback12 - 12.1.1.1/32: Redistribution Metric - Default
2. R4: Loopback13 - 13.1.1.1/32: Redistribution Metric - Default

Task 5

Redistribute between EIGRP 1024 and EIGRP 2048. Ensure that the next hop IP address for the redistributed routes is the IP address of the router that originated the route and NOT the IP address of R1. Verify your configuration.

Task 6

Configure route filtering on R4 so that the router only accepts routes with a metric between 200,000 and 300,000. Verify your configuration using the appropriate commands.

Task 1:

Configuration of all routers:

R1:

```
interface FastEthernet0/0
 ip address 150.1.1.1 255.255.255.0
 no shut
```

R2:

```
interface Loopback192
 ip address 192.1.1.1 255.255.255.255
!
interface Loopback193
 ip address 193.1.1.1 255.255.255.255
!
interface FastEthernet0/0
 ip address 150.1.1.2 255.255.255.0
 no shut
```

R3:

```
interface Loopback12
 ip address 12.1.1.1 255.255.255.255
!
interface Loopback13
 ip address 13.1.1.1 255.255.255.255
!
interface FastEthernet0/0
 ip address 150.1.1.3 255.255.255.0
 no shut
```

R4:

```
interface Loopback12
 ip address 12.1.1.1 255.255.255.255
!
interface Loopback13
 ip address 13.1.1.1 255.255.255.255
!
interface FastEthernet0/0
 ip address 150.1.1.4 255.255.255.0
 no shut
```

Task 2:

For EIGRP we will use Unicast packets instead of Multicast packets to establish the adjacencies:

```
R1(config)#router eirgp 1024
R1(config-router)#neighbor 150.1.1.3 fastethernet 0/0
R1(config-router)#neighbor 150.1.1.4 fastethernet 0/0
R1(config-router)#network 150.1.1.1 0.0.0.0
```

```
R3(config)#router eigrp 1024
R3(config-router)#neighbor 150.1.1.1 fastethernet 0/0
R3(config-router)#neighbor 150.1.1.4 fastethernet 0/0
R3(config-router)#network 150.1.1.3 0.0.0.0
```

```
R4(config)#router eigrp 1024
R4(config-router)#neighbor 150.1.1.1 fastethernet 0/0
R4(config-router)#neighbor 150.1.1.3 fastethernet 0/0
R4(config-router)#network 150.1.1.4 0.0.0.0
```

Task 3:

```
R1(config)#router eigrp 2048
R1(config-router)#neighbor 150.1.1.2 fastethernet 0/0
R1(config-router)#network 150.1.1.1 0.0.0.0
```

```
R2(config)#router eigrp 2048
R2(config-router)#neighbor 150.1.1.1 fastethernet 0/0
R2(config-router)#network 150.1.1.2 0.0.0.0
```

Task 4:

Configure and advertise the following subnets via EIGRP with the requirements below:
1.R2: Loopback192 - 192.1.1.1/32: Redistribution Metric 100000 1000 255 1 1500
2.R2: Loopback193 - 193.1.1.1/32: Redistribution Metric 10000 500 255 1 1500

```
R2(config)#route-map LOOPBACK permit 10
R2(config-route-map)#match interface loopback 192
R2(config-route-map)#set metric 100000 1000 255 1 1500
R2(config)#route-map LOOPBACK permit 20
R2(config-route-map)#match interface loopback 193
R2(config-route-map)#set metric 10000 500 255 1 1500
R2(config)#router eigrp 2048
R2(config-router)#redistribute connected route-map LOOPBACK
```

1.R3: Loopback12 - 12.1.1.1/32: Redistribution Metric 30000 300 255 1 1500
2.R3: Loopback13 - 13.1.1./32: Redistribution Metric 40000 400 255 1 1500

```
R3(config)#route-map LOOPBACK permit 10
R3(config-route-map)#match interface loopback 12
R3(config-route-map)#set metric 30000 300 255 1 1500
R3(config)#route-map LOOPBACK permit 20
R3(config-route-map)#match interface loopback 13
R3(config-route-map)#set metric 40000 400 255 1 1500
R3(config)#router eigrp 1024
R3(config-router)#redistribute connected route-map LOOPBACK
```

1.R4: Loopback12 - 12.1.1.1/32: Redistribution Metric - Default
2.R4: Loopback13 - 13.1.1.1/32: Redistribution Metric - Default

```
R4(config-if)#exit
R4(config)#router eigrp 1024
R4(config-router)#redistribute connected
```

Task 5:

R1 Redistributes between EIGRP 1024 and EIGRP 2048:

```
R1(config)#router eigrp 1024
R1(config-router)#redistribute eigrp 2048
R1(config)#router eigrp 2048
R1(config-router)#redistribute eigrp 1024
```

Let's look the routing table of each router:

```
R2(config-route-map)#do show ip route eigrp
     12.0.0.0/32 is subnetted, 1 subnets
D EX    12.1.1.1 [170/384000] via 150.1.1.1, 00:20:14, FastEthernet0/0
     13.0.0.0/32 is subnetted, 1 subnets
D EX    13.1.1.1 [170/409600] via 150.1.1.1, 00:20:14, FastEthernet0/0
R2(config-route-map)#
```

```
R3(config-router)#do show ip route eigrp
     193.1.1.0/32 is subnetted, 1 subnets
D EX    193.1.1.1 [170/435200] via 150.1.1.1, 00:18:46, FastEthernet0/0
     192.1.1.0/32 is subnetted, 1 subnets
D EX    192.1.1.1 [170/563200] via 150.1.1.1, 00:20:34, FastEthernet0/0
R3(config-router)#
```

```
R4(config-router)#do show ip route eigrp
     193.1.1.0/32 is subnetted, 1 subnets
D EX    193.1.1.1 [170/435200] via 150.1.1.1, 00:19:10, FastEthernet0/0
     192.1.1.0/32 is subnetted, 1 subnets
D EX    192.1.1.1 [170/563200] via 150.1.1.1, 00:20:58, FastEthernet0/0
R4(config-router)#
```

We can deduce that the next-hop for each subnet in each router is always R1 (150.1.1.1), because EIGRP set the next-hop field to 0.0.0.0 so R1 when redistributing between the TWO eigrp as it leaves the next-hop field to 0.0.0.0 which tell to the other routers: use the router sendig the update as the next-hop,remember the FA in OSPF when it is set to 0.0.0.0 by an ASBR meaning use the ASBR as the next-hop.
To override the this extra-hop and route the packet more efficiently, we must change the next-hop field by configuring the ip next-hop-self command on R1:

```
R1(config)#interface fastethernet 0/0
R1(config-if)#no ip next-hop-self eigrp 1024
R1(config-if)#no ip next-hop-self eigrp 2048
```

Let' verify the routing table:
Now on R2 the next-hop to reach the loopback interfaces of R3 is the router R3 (150.1.1.3):

```
R2(config-route-map)#do show ip route eigrp
     12.0.0.0/32 is subnetted, 1 subnets
D EX    12.1.1.1 [170/384000] via 150.1.1.3, 00:00:07, FastEthernet0/0
     13.0.0.0/32 is subnetted, 1 subnets
```

```
D EX    13.1.1.1 [170/409600] via 150.1.1.3, 00:00:07, FastEthernet0/0
```

Here the next-hop to reach the loopback interfaces of R2 on both R3 and R4 is the router R2(150.1.1.2):

```
R3(config)#do show ip route eigrp
     193.1.1.0/32 is subnetted, 1 subnets
D EX    193.1.1.1 [170/435200] via 150.1.1.2, 00:00:38, FastEthernet0/0
     192.1.1.0/32 is subnetted, 1 subnets
D EX    192.1.1.1 [170/563200] via 150.1.1.2, 00:00:38, FastEthernet0/0
R3(config)#
```

```
R4(config)#do show ip route eigrp
     193.1.1.0/32 is subnetted, 1 subnets
D EX    193.1.1.1 [170/435200] via 150.1.1.2, 00:00:53, FastEthernet0/0
     192.1.1.0/32 is subnetted, 1 subnets
D EX    192.1.1.1 [170/563200] via 150.1.1.2, 00:00:53, FastEthernet0/0
R4(config)#
```

Task 6:

Configure on R4 so that the router installs only routes with a metric located between 500 000 and 600 000.
The routing table of R4 show the prefix 193.1.1.1/32 with the metric 435200 and the prefix 192.1.1.1/32 with the metric 563200.
to do this requirement so that only routes with a metric between 500,000 and 600,000 are installed on R4 we configure a route-map and we specify the metric range with the match metric command, we will use the command match metric 550 000 +- 50000 under the route-map which tells the router to match the routes with a metric that is 50000 less than 550 000 (500 000) or 50000 more than 550 000 (600 000) :

```
R4(config)#route-map EIGRP-FILTER permit 10
R4(config-route-map)#match metric 550000 +- 50000
R4(config)#route-map EIGRP-FILTER deny 20
R4(config)#router eigrp 1024
R4(config-router)#distribute-list route-map EIGRP-FILTER in
```

Now notice the subnet 193.1.1.1/32 is filtered and not installed in the routing table of R4 because this metric (435200) is outside the permitted range of 500,000 to 600,000:

```
R4(config-router)#do show ip route eigrp
     192.1.1.0/32 is subnetted, 1 subnets
D EX    192.1.1.1 [170/563200] via 150.1.1.2, 00:17:37, FastEthernet0/0
R4(config-router)#
```

Lab 4: EIGRP route filtering using prefix-list and route-map

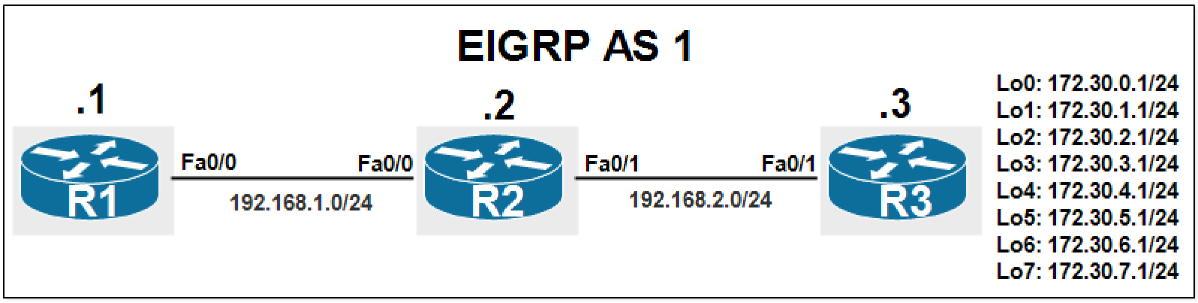

Task 1

Configure hostnames and IP addressing on all routers as illustrated in the network topology.

Task 2

On R2. Filter the routes that begin with 172.30 with the Prefix length between 16 and 24 using two different methods.

Task 3

On R2. Filter only the prefix 172.16.3.0/24 using two different methods.

Task 4

Add four loopback interfaces on R3:

1. **Lo10: 10.1.1.1/24**
2. **Lo20: 10.2.2.1/24**
3. **Lo30: 10.3.3.3/24**
4. **Lo40: 10.4.4.5/24**

On R2
-filter the routes with the Prefix length equal to 30.
-filter the routes that begin with 10. with the Prefix length between 20 and 22.
Use the prefix-list feature.

Task 1:

Configuration of all routers:

R1:

```
interface FastEthernet0/0
 ip address 192.168.1.1 255.255.255.0
 no shutdown
!
router eigrp 1
 network 192.168.1.0
 no auto-summary
```

R2:

```
interface FastEthernet0/0
 ip address 192.168.1.2 255.255.255.0
 no shutdown
!
interface FastEthernet0/1
 ip address 192.168.2.2 255.255.255.0
 no shutdown
!
router eigrp 1
 network 192.168.1.0
 network 192.168.2.0
 no auto-summary
```

R3:

```
interface Loopback0
 ip address 172.30.0.1 255.255.255.0
!
interface Loopback1
 ip address 172.30.1.1 255.255.255.0
!
interface Loopback2
 ip address 172.30.2.1 255.255.255.0
!
interface Loopback3
 ip address 172.30.3.1 255.255.255.0
!
interface Loopback4
 ip address 172.30.4.1 255.255.255.0
!
interface Loopback5
 ip address 172.30.5.1 255.255.255.0
!
interface Loopback6
 ip address 172.30.6.1 255.255.255.0
!
interface Loopback7
 ip address 172.30.7.1 255.255.255.0
!
interface FastEthernet0/1
 ip address 192.168.2.3 255.255.255.0
 no shutdown
!
router eigrp 1
 network 172.30.0.0
 network 192.168.2.0
 no auto-summary
```

Before configuring the filtering let's look the routing table of R1:

```
R1(config-router)#do show ip route eigrp
172.30.0.0/24 is subnetted, 8 subnets
D 172.30.2.0 [90/435200] via 192.168.1.2, 00:04:08, FastEthernet0/0
D 172.30.3.0 [90/435200] via 192.168.1.2, 00:04:08, FastEthernet0/0
D 172.30.0.0 [90/435200] via 192.168.1.2, 00:04:08, FastEthernet0/0
D 172.30.1.0 [90/435200] via 192.168.1.2, 00:04:08, FastEthernet0/0
D 172.30.6.0 [90/435200] via 192.168.1.2, 00:04:08, FastEthernet0/0
D 172.30.7.0 [90/435200] via 192.168.1.2, 00:04:08, FastEthernet0/0
D 172.30.4.0 [90/435200] via 192.168.1.2, 00:04:08, FastEthernet0/0
D 172.30.5.0 [90/435200] via 192.168.1.2, 00:04:08, FastEthernet0/0
D 192.168.2.0/24 [90/307200] via 192.168.1.2, 00:18:35, FastEthernet0/0
```

Task 1:

Filter the routes that begin with 172.30 with the Prefix length between 16 and 24.

First method using prefix-list :

```
R2(config-router)#ip prefix-list list1 seq 5 deny 172.30.0.0/16 le 24
R2(config)#ip prefix-list list1 seq 100 permit 0.0.0.0/0 le 32
R2(config)#router eigrp 1
R2(config-router)#distribute-list prefix list1 out fa0/0
```

```
R1#show ip route eigrp
D 192.168.2.0/24 [90/307200] via 192.168.1.2, 00:28:34, FastEthernet0/0
```

Second method using route-map:

```
R2(config)# ip prefix-list list1 seq 5 permit 172.30.0.0/16 le 24
R2(config)# route-map routemap1 deny 1
R2(config-route-map)# match ip address prefix-list list1
R2(config)# route-map routemap1 permit 10
R2(config)# router eigrp 1
R2(config-router)# distribute-list route-map routemap1 out fastEthernet 0/0
```

```
R1#show ip route eigrp
D 192.168.2.0/24 [90/307200] via 192.168.1.2, 00:00:21, FastEthernet0/0
```

Task 2:

Remove the previous configuration.

Filter only the prefix 172.16.3.0/24:

First method using prefix-list:

```
R2(config)#ip prefix-list net-172.30.3 seq 5 deny 172.30.3.0/24
R2(config)#ip prefix-list net-172.30.3 seq 10 permit 0.0.0.0/0 le 32
R2(config)#router eigrp 1
R2(config-router)#distribute-list prefix net-172.30.3 out fa0/0
```

```
R1#show ip route eigrp
172.30.0.0/24 is subnetted, 7 subnets
D 172.30.2.0 [90/435200] via 192.168.1.2, 00:03:36, FastEthernet0/0
D 172.30.0.0 [90/435200] via 192.168.1.2, 00:03:36, FastEthernet0/0
```

```
D 172.30.1.0 [90/435200] via 192.168.1.2, 00:03:36, FastEthernet0/0
D 172.30.6.0 [90/435200] via 192.168.1.2, 00:03:36, FastEthernet0/0
D 172.30.7.0 [90/435200] via 192.168.1.2, 00:03:36, FastEthernet0/0
D 172.30.4.0 [90/435200] via 192.168.1.2, 00:03:36, FastEthernet0/0
D 172.30.5.0 [90/435200] via 192.168.1.2, 00:03:36, FastEthernet0/0
D 192.168.2.0/24 [90/307200] via 192.168.1.2, 00:03:36, FastEthernet0/0
```

Second method using route-map:

```
R2(config)#ip prefix-list net-172.30.3 seq 5 permit 172.30.3.0/24
R2(config)#route-map routemap1 deny
R2(config-route-map)#match ip address prefix-list net-172.30.3
R2(config-router)#route-map routemap1 permit 20
R2(config-route-map)#router eigrp 1
R2(config-router)#distribute-list route-map routemap1 out fastEthernet 0/0
```

```
R1#show ip route eigrp
172.30.0.0/24 is subnetted, 7 subnets
D 172.30.2.0 [90/435200] via 192.168.1.2, 00:09:01, FastEthernet0/0
D 172.30.0.0 [90/435200] via 192.168.1.2, 00:09:01, FastEthernet0/0
D 172.30.1.0 [90/435200] via 192.168.1.2, 00:09:01, FastEthernet0/0
D 172.30.6.0 [90/435200] via 192.168.1.2, 00:09:01, FastEthernet0/0
D 172.30.7.0 [90/435200] via 192.168.1.2, 00:09:01, FastEthernet0/0
D 172.30.4.0 [90/435200] via 192.168.1.2, 00:09:01, FastEthernet0/0
D 172.30.5.0 [90/435200] via 192.168.1.2, 00:09:01, FastEthernet0/0
D 192.168.2.0/24 [90/307200] via 192.168.1.2, 00:09:01, FastEthernet0/0
```

Task 3:

Remove the previous configuration.

Let's add four loopback interfaces on R3 and enable EIGRP as follow :

```
interface Loopback10
 ip address 10.1.1.1 255.255.240.0
!
interface Loopback20
 ip address 10.2.2.1 255.255.248.0
!
interface Loopback30
 ip address 10.3.3.3 255.255.0.0
!
interface Loopback40
 ip address 10.4.4.5 255.255.255.252
```

```
R3(config-if)#router eigrp 1
R3(config-router)#network 10.0.0.0
R3(config-router)#no auto-summary
```

```
R1#show ip route eigrp
172.30.0.0/24 is subnetted, 8 subnets
D 172.30.2.0 [90/435200] via 192.168.1.2, 00:03:15, FastEthernet0/0
D 172.30.3.0 [90/435200] via 192.168.1.2, 00:03:15, FastEthernet0/0
D 172.30.0.0 [90/435200] via 192.168.1.2, 00:03:15, FastEthernet0/0
D 172.30.1.0 [90/435200] via 192.168.1.2, 00:03:15, FastEthernet0/0
```

```
D 172.30.6.0 [90/435200] via 192.168.1.2, 00:03:15, FastEthernet0/0
D 172.30.7.0 [90/435200] via 192.168.1.2, 00:03:15, FastEthernet0/0
D 172.30.4.0 [90/435200] via 192.168.1.2, 00:03:15, FastEthernet0/0
D 172.30.5.0 [90/435200] via 192.168.1.2, 00:03:15, FastEthernet0/0
10.0.0.0/8 is variably subnetted, 4 subnets, 4 masks
D 10.2.0.0/21 [90/435200] via 192.168.1.2, 00:03:15, FastEthernet0/0
D 10.3.0.0/16 [90/435200] via 192.168.1.2, 00:03:15, FastEthernet0/0
D 10.1.0.0/20 [90/435200] via 192.168.1.2, 00:03:15, FastEthernet0/0
D 10.4.4.4/30 [90/435200] via 192.168.1.2, 00:00:36, FastEthernet0/0
D 192.168.2.0/24 [90/307200] via 192.168.1.2, 00:03:16, FastEthernet0/0
```

Task 4:

Remove the previous configuration.

-Filter the routes with the Prefix length equal to 30.
-Filter the routes that begin with 10. with the Prefix length between 20 and 22.

```
R2(config)#ip prefix-list FILTER deny 10.0.0.0/8 ge 20 le 22
R2(config)#ip prefix-list FILTER deny 0.0.0.0/0 ge 30 le 30
R2(config)#ip prefix-list FILTER permit 0.0.0.0/0 le 32
R2(config)#router eigrp 1
R2(config-router)#distribute-list prefix FILTER out fastEthernet 0/0
```

```
R1#show ip route eigrp
172.30.0.0/24 is subnetted, 8 subnets
D 172.30.2.0 [90/435200] via 192.168.1.2, 00:00:02, FastEthernet0/0
D 172.30.3.0 [90/435200] via 192.168.1.2, 00:00:02, FastEthernet0/0
D 172.30.0.0 [90/435200] via 192.168.1.2, 00:00:02, FastEthernet0/0
D 172.30.1.0 [90/435200] via 192.168.1.2, 00:00:02, FastEthernet0/0
D 172.30.6.0 [90/435200] via 192.168.1.2, 00:00:02, FastEthernet0/0
D 172.30.7.0 [90/435200] via 192.168.1.2, 00:00:02, FastEthernet0/0
D 172.30.4.0 [90/435200] via 192.168.1.2, 00:00:02, FastEthernet0/0
D 172.30.5.0 [90/435200] via 192.168.1.2, 00:00:02, FastEthernet0/0
10.0.0.0/16 is subnetted, 1 subnets
D 10.3.0.0 [90/435200] via 192.168.1.2, 00:00:02, FastEthernet0/0
D 192.168.2.0/24 [90/307200] via 192.168.1.2, 00:00:02, FastEthernet0/0
```

Lab 5: EIGRP Route Filtering using Route-Map Prefix-List ACL

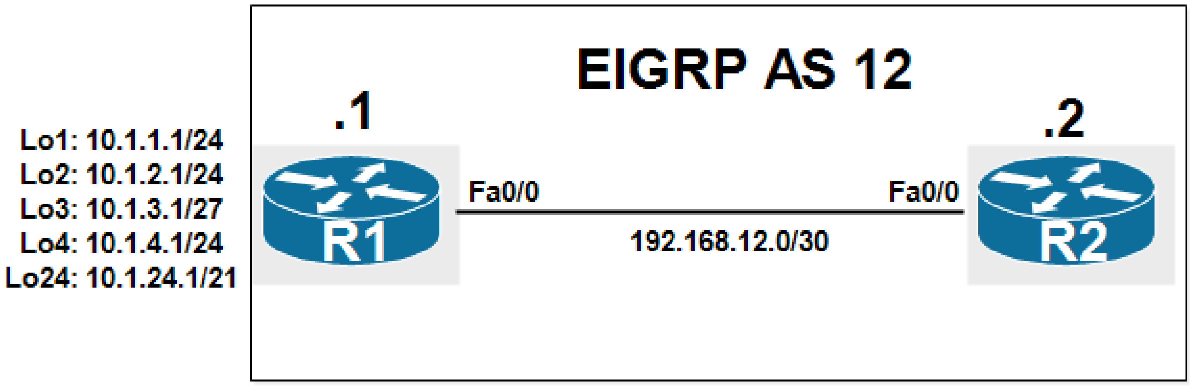

Task 1

Configure hostnames and IP addressing on all routers as illustrated in the network topology.

Task 2

Advertise all networks 10.1.0.0/21 using an ACL.

Task 3

Advertise only 10.1.1.0 and 10.1.3.0 subnets using an ACL.

Task 4

Advertise only 10.1.1.0/24 and 10.1.3.0/24 using prefix-list..

Task 5

Advertise all prefixes in the range 10.1.0.0/16 with a prefix length <= 24 (less or equal than 24) using prefix-list.

Task 6

Advertise all prefixes in the range 10.1.0.0/16 with a prefix length >= 24 (greater or equal than 24) using prefix-list.

Task 7

Advertise all prefixes with a prefix length <= 21 and >=27 (–> 10.1.3.0/27 and 10.1.24.0/21), but not a default-route using prefix-list.

Task 8

Advertise all prefixes with a prefix length >= 21 and <=27.

Task 9

Advertise only 10.1.1.0 and 10.1.3.0 using Route-map combined with standard access-list.

Task 10

Advertise all prefixes in the range 10.1.0.0/16 with a prefix length >= 24 using Route-map combined with ip prefix-list.

Task 11

Advertise all prefixes in the range 10.1.0.0/16 with a prefix length >= 24 and 10.1.3.0 (without a specific prefix) using Route-map combined with ip prefix-list and access-list.

Task 12

Advertise the prefixes configured on Loopback1 Loopback2 Loopback24 interfaces, use only a route-map. Do not use an ACL or a prefix-list.

Task 1

R1:

```
int fa0/0
 ip add 192.168.12.1 255.255.255.0
 no shutdown
!
interface Loopback1
 ip address 10.1.1.1 255.255.255.0
!
interface Loopback2
 ip address 10.1.2.1 255.255.255.0
!
interface Loopback3
 ip address 10.1.3.1 255.255.255.224
!
interface Loopback4
 ip address 10.1.4.1 255.255.255.0
!
interface Loopback30
 ip address 10.1.24.1 255.255.248.0
!
router eigrp 12
 network 192.168.12.0
 redistribute connected
```

R2:

```
int fa0/0
 ip add 192.168.12.2 255.255.255.0
 no shutdown
!
router eigrp 12
 network 192.168.12.0
```

R2 receives 5 external eigrp routes from R1 as shown by its routing table*:*

```
R2#show ip route eigrp | begin 10.0.0.0/8
      10.0.0.0/8 is variably subnetted, 5 subnets, 3 masks
D EX     10.1.1.0/24 [170/156160] via 192.168.12.1, 00:01:13, FastEthernet0/0
D EX     10.1.2.0/24 [170/156160] via 192.168.12.1, 00:01:13, FastEthernet0/0
D EX     10.1.3.0/27 [170/156160] via 192.168.12.1, 00:01:13, FastEthernet0/0
D EX     10.1.4.0/24 [170/156160] via 192.168.12.1, 00:01:13, FastEthernet0/0
D EX     10.1.24.0/21 [170/156160] via 192.168.12.1, 00:01:13, FastEthernet0/0
```

Task 2

Advertise all networks 10.1.0.0/21 using an ACL:

R1:

```
router eigrp 12
 distribute-list 1 out connected
 no auto-summary
!
access-list 1 permit 10.1.0.0 0.0.7.255
access-list 1 deny any
```

The subnet 10.1.24.0/21 is not installed in the routing table of R2 because it does not match the first instruction of the ACL so it is denied by the second instruction (deny any):

```
R2#show ip route eigrp | begin 10.0.0.0/8
      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
D EX     10.1.1.0/24 [170/156160] via 192.168.12.1, 00:08:30, FastEthernet0/0
D EX     10.1.2.0/24 [170/156160] via 192.168.12.1, 00:08:30, FastEthernet0/0
D EX     10.1.3.0/27 [170/156160] via 192.168.12.1, 00:08:30, FastEthernet0/0
D EX     10.1.4.0/24 [170/156160] via 192.168.12.1, 00:08:30, FastEthernet0/0
R2#
```

Task 3

Advertise only 10.1.1.0 and 10.1.3.0 subnets using still an ACL:

```
router eigrp 12
 distribute-list 1 out connected
 no auto-summary
!
access-list 1 permit 10.1.1.0 0.0.0.255
access-list 1 permit 10.1.3.0 0.0.0.255
access-list 1 deny   any
```

R2 installs only 10.1.1.0 and 10.1.3.0 ,all other routes are filtered by R1:

```
R2#show ip route eigrp | begin 10.0.0.0/8
      10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
D EX     10.1.1.0/24 [170/156160] via 192.168.12.1, 00:12:40, FastEthernet0/0
D EX     10.1.3.0/27 [170/156160] via 192.168.12.1, 00:12:40, FastEthernet0/0
```

As you can see, filtering with standard ACLs is very easy but it don't care about the prefix length.
prefix-lists are are more flexible way for filtering.

By using IP prefix lists the router can examine both the prefix and the prefix length, and a range of prefixes or a range of prefix lengths. The command then sets either a deny or permit action for each matched prefix/length like an ACL. To use the prefix-list, the configuration simply refers to the prefix-list with the same distribute-list command seen earlier.
Using IP prefix lists for route filtering has several advantages. First, IP prefix lists allow matching of the prefix length, whereas the ACLs used by the EIGRP distribute-list command cannot. (Some other route filtering configurations can match both the prefix and prefix length using extended ACLs.) Many people find IP prefix lists more intuitive for configuring route filtering. Finally, the internal processing of the IP prefix lists uses an internal tree structure that results in faster matching of routes as compared with ACLs.

Task 4

Advertise only 10.1.1.0/24 and 10.1.3.0/24 using prefix-list.

On R1:

```
router eigrp 12
 distribute-list prefix FILTER  out connected
 no auto-summary
!
ip prefix-list FILTER  seq 5 permit 10.1.1.0/24
ip prefix-list FILTER  seq 10 permit 10.1.3.0/24
```

Because 10.1.3.0 has a prefix-length of 27, it is not allowed ,if we have configured the loopback with 10.1.3.1/24 ,it will be allowed. So Only The subnet 10.1.1.1 with a /24 netmask is included in the routes updates :

Thus, R2 installs only the subnet 10.1.1.0/24 as shown by the following output:

```
R2#show ip route eigrp | begin 10
      10.0.0.0/24 is subnetted, 1 subnets
D EX     10.1.1.0 [170/156160] via 192.168.12.1, 00:20:21, FastEthernet0/0
R2#
```

Task 5

Advertise all prefixes in the range 10.1.0.0/16 with a prefix length <= 24 (less or equal than 24):

```
router eigrp 12
 distribute-list prefix FILTER out connected
!
ip prefix-list FILTER seq 5 permit 10.1.0.0/16 le 24
```

R2 does not have the subnet 10.1.3.0/27 in its routing table because it has a prefix-length of 27 which is greater than 24, so it does not meet the criteria defined in the prefix-list (less or equal than 24):

```
R2#show ip route eigrp | begin 10
      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
D EX     10.1.1.0/24 [170/156160] via 192.168.12.1, 00:29:33, FastEthernet0/0
D EX     10.1.2.0/24 [170/156160] via 192.168.12.1, 00:05:29, FastEthernet0/0
D EX     10.1.4.0/24 [170/156160] via 192.168.12.1, 00:05:29, FastEthernet0/0
D EX     10.1.24.0/21 [170/156160] via 192.168.12.1, 00:00:15, FastEthernet0/0
```

Task 6

Advertise all prefixes in the range 10.1.0.0/16 with a prefix length >= 24 (greater or equal than 24):

```
router eigrp 12
 distribute-list prefix FILTER out connected
!
ip prefix-list FILTER seq 5 permit 10.1.0.0/16 ge 24
```

In this case the subnet 10.1.24.0/21 is not learned by R2 because it has a prefix-length of 21 which is less than 24, so it does not meet the criteria (greater or equal than 24):

```
R2#show ip route eigrp | begin 10
      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
D EX     10.1.1.0/24 [170/156160] via 192.168.12.1, 00:33:09, FastEthernet0/0
D EX     10.1.2.0/24 [170/156160] via 192.168.12.1, 00:09:05, FastEthernet0/0
D EX     10.1.3.0/27 [170/156160] via 192.168.12.1, 00:00:35, FastEthernet0/0
D EX     10.1.4.0/24 [170/156160] via 192.168.12.1, 00:09:05, FastEthernet0/0
```

Task 7

Advertise all prefixes with a prefix length <= 21 and >=27 (–> 10.1.3.0/27 and 10.1.24.0/21), but not a default-route.

0.0.0.0/0 represents a default route.
0.0.0.0/0 le 21 represents all prefixes with a prefix length between 0 and 21 (only 10.1.24.0/21 matches this criteria).
0.0.0.0/0 ge 27 represents all prefixes with a prefix length between 27 and 32 (only 10.1.3.0/27 matches this criteria).

```
router eigrp 12
 distribute-list prefix FILTER out connected
!
ip prefix-list FILTER seq 5 deny 0.0.0.0/0
ip prefix-list FILTER seq 10 permit 0.0.0.0/0 le 21
ip prefix-list FILTER seq 15 permit 0.0.0.0/0 ge 27
```

R2 installs only two routes toward the prefixes 10.1.24.0/21 and 10.1.3.0/27 :

Task 8

Advertise all prefixes with a prefix length >= 21 and <=27:

```
router eigrp 12
 distribute-list prefix FILTER out connected
!
ip prefix-list FILTER seq 10 deny 0.0.0.0/0 ge 28
ip prefix-list FILTER seq 15 deny 0.0.0.0/0 le 20
ip prefix-list FILTER seq 20 permit 0.0.0.0/0 le 27
```

The prefix length 21 24 and 27 are included in the space >= 21 and <=27, so R2 receives all routes eigrp from R1:

```
R2#show ip route eigrp | begin 10
      10.0.0.0/8 is variably subnetted, 5 subnets, 3 masks
```

```
D EX     10.1.1.0/24 [170/156160] via 192.168.12.1, 00:00:06, FastEthernet0/0
D EX     10.1.2.0/24 [170/156160] via 192.168.12.1, 00:00:06, FastEthernet0/0
D EX     10.1.3.0/27 [170/156160] via 192.168.12.1, 00:07:50, FastEthernet0/0
D EX     10.1.4.0/24 [170/156160] via 192.168.12.1, 00:00:06, FastEthernet0/0
D EX     10.1.24.0/21 [170/156160] via 192.168.12.1, 00:02:49, FastEthernet0/0
```

Route map is the third EIGRP route filtering tool that can be referenced with the distribute-list command.
Route maps can be used for many functions besides being used to filter routes for a single routing protocol like EIGRP. Route maps can be used to filter routes during the route redistribution process, and to set BGP Path Attributes (PAs) for the purpose of influencing the choice of the best routes in an internetwork.
When used for filtering EIGRP routes, route maps do provide a few additional features beyond what can be configured using ACLs and prefix lists.

Task 9

Route-map combined with standard access-list.
Advertise only 10.1.1.0 and 10.1.3.0.

```
router eigrp 12
 redistribute connected route-map CONNECTED
!
access-list 1 permit 10.1.1.0 0.0.0.255
access-list 1 permit 10.1.3.0 0.0.0.255
access-list 1 deny   any
!
route-map CONNECTED permit 10
 match ip address 1
```

We can see that R2 receives only the two routes allowed in the route-map 10.1.1.0 and 10.1.3.0:

```
R2#show ip route eigrp | begin 10
      10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
D EX     10.1.1.0/24 [170/156160] via 192.168.12.1, 00:00:27, FastEthernet0/0
D EX     10.1.3.0/27 [170/156160] via 192.168.12.1, 00:00:14, FastEthernet0/0
```

Task 10

Route-map combined with ip prefix-list
Advertise all prefixes in the range 10.1.0.0/16 with a prefix length >= 24

```
router eigrp 12
 redistribute connected route-map CONNECTED
!
ip prefix-list FILTER seq 5 permit 10.1.0.0/16 le 24
!
route-map CONNECTED permit 10
 match ip address prefix-list FILTER
```

R2 does not have the subnet 10.1.3.0/27 in its routing table because it has a prefix-length of 27 so R1 filters this route:

```
R2#show ip route eigrp | begin 10
      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
D EX     10.1.1.0/24 [170/156160] via 192.168.12.1, 00:00:23, FastEthernet0/0
```

```
D EX     10.1.2.0/24 [170/156160] via 192.168.12.1, 00:00:23, FastEthernet0/0
D EX     10.1.4.0/24 [170/156160] via 192.168.12.1, 00:00:23, FastEthernet0/0
D EX     10.1.24.0/21 [170/156160] via 192.168.12.1, 00:00:23, FastEthernet0/0
```

Task 11

Route-map combined with ip prefix-list and access-list
Advertise all prefixes in the range 10.1.0.0/16 with a prefix length >= 24 and 10.1.3.0 (without a specific prefix).

```
router eigrp 12
 redistribute connected route-map CONNECTED
!
ip prefix-list FILTER seq 5 permit 10.1.0.0/16 le 24
!
access-list 1 permit 10.1.3.0 0.0.0.255
access-list 1 deny any
!
route-map CONNECTED permit 10
 match ip address prefix-list FILTER
!
route-map CONNECTED permit 20
 match ip address 1
```

10.1.3.0/27 is allowed by the sequence 20 of the route-map because the ACL 2 referenced in this route-map seq 20 identifies the network 10.1.3.0/24 inlcuding all subnets such as 10.1.3.0/27.
All other subnets are allowed by the sequence 10 of route-map, these subnets are identified by the prefix-list named FILTER which matches all subnets that begin with 10.1 and with a prefix length greater than 24 and less than 16:

```
R2#show ip route eigrp | begin 10
      10.0.0.0/8 is variably subnetted, 5 subnets, 3 masks
D EX     10.1.1.0/24 [170/156160] via 192.168.12.1, 00:00:12, FastEthernet0/0
D EX     10.1.2.0/24 [170/156160] via 192.168.12.1, 00:00:12, FastEthernet0/0
D EX     10.1.3.0/27 [170/156160] via 192.168.12.1, 00:00:12, FastEthernet0/0
D EX     10.1.4.0/24 [170/156160] via 192.168.12.1, 00:00:12, FastEthernet0/0
D EX     10.1.24.0/21 [170/156160] via 192.168.12.1, 00:00:12, FastEthernet0/0
```

Task 12

Advertise the prefixes configured on the interface matching the route-map,this route-map should match Loopback1 Loopback2 Loopback24 interfaces:

```
router eigrp 12
 redistribute connected route-map CONNECTED
!
route-map CONNECTED permit 10
 match interface Loopback1 Loopback2 Loopback24
```

Only the subnets of the loopback1 ,loopback2 and loopback24 are sent to R2 as shown by the routing table of R2:

```
R2#show ip route eigrp | begin 10
      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
```

```
D EX     10.1.1.0/24 [170/156160] via 192.168.12.1, 00:00:49, FastEthernet0/0
D EX     10.1.2.0/24 [170/156160] via 192.168.12.1, 00:00:49, FastEthernet0/0
D EX     10.1.24.0/21 [170/156160] via 192.168.12.1, 00:00:49, FastEthernet0/0
```

Lab 6: EIGRP Route Leaking Maximum hop-count

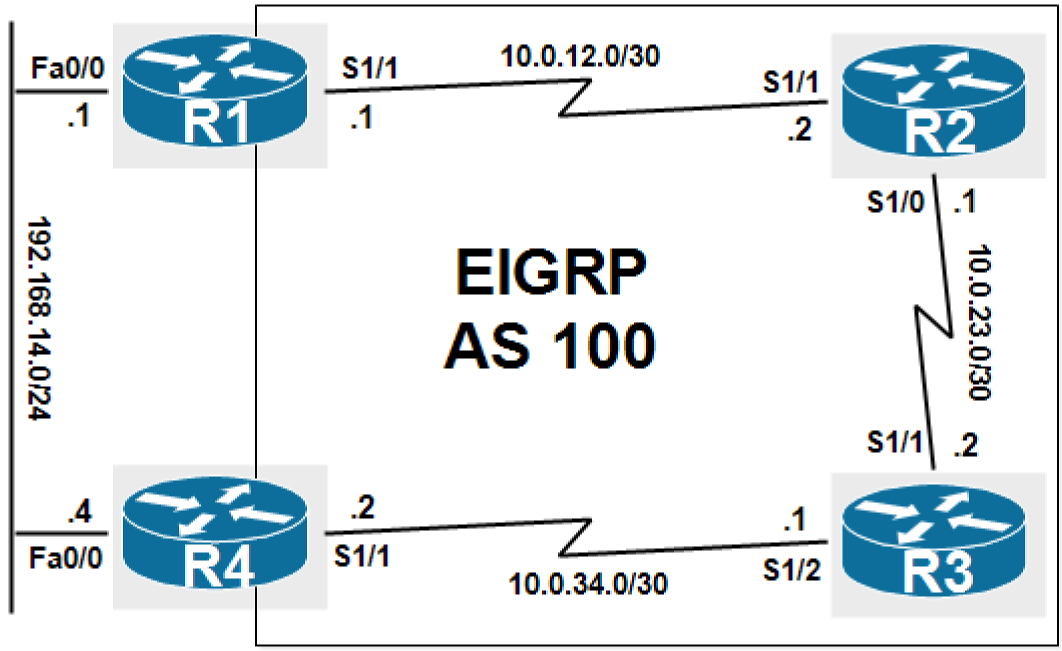

Task 1

Configure hostnames and IP addressing on all routers as illustrated in the network topology.

Task 2

Configure EIGRP for AS 100 as illustrated in the topology. Ensure that the 192.168.14.0/24 subnet is advertised as an external EIGRP route with a route tag of 11 from R1 and 44 from R4.

Ensure that neither R1 nor R4 has a duplicate entry for this subnet in the EIGRP topology table. You are NOT allowed to use IP ACLs or IP prefix lists to implement this solution. Verify your configuration using the appropriate commands.

Task 3

Configure the 10.5.0.0/24, 10.5.1.0/24, 10.5.2.0/24, and 10.5.3.0/24 secondary subnets on the LAN segment between R1 and R4. Use the .1 address for R1 and the .4 address for R4. These networks should be advertised via EIGRP as internal EIGRP routes.

Task 4

Configure EIGRP on routers R1 and R4 to advertise a single summary route for the 10.14.0.0/24, 10.14.1.0/24, 10.14.2.0/24, and 10.14.3.0/24 secondary subnets. Verify this on R2 and R3.

Task 5

Without removing the summary configuration on routers R1 and R4, adding any static routes, or using PBR; configure your network so that R2 and R3 prefer the path via R1 to reach the 10.5.0.0/24 subnet; and the path via R4 to reach the 10.5.3.0/24 subnet. As a hint, use an EIGRP feature to complete this task. Verify your configuration using the appropriate commands.

Task 6

Your EIGRP network will grow to at least 200 routers and these will incorporate the EIGRP network. Configure the current EIGRP implementation to support the future EIGRP network.

Task 1

R1

```
interface FastEthernet0/0
 ip address 192.168.14.4 255.255.255.0
 no shutd
!
interface Serial1/1
 ip address 10.0.12.1 255.255.255.252
 no shutd
!
router eig 100
 net 10.0.12.1 0.0.0.0
 no au
```

R2

```
interface Serial1/1
 ip address 10.0.12.2 255.255.255.252
 no shutd
!
interface Serial1/0
 ip address 10.0.23.1 255.255.255.252
 no shutd
!
router eig 100
 net 10.0.23.1 0.0.0.0
 net 10.0.12.2 0.0.0.0
 no au
```

R3

```
interface Serial1/1
 ip address 10.0.23.2 255.255.255.252
 no shutd
!
interface Serial1/2
 ip address 10.0.34.1 255.255.255.252
 no shutd
!
router eig 100
 net 10.0.23.2 0.0.0.0
 net 10.0.34.1 0.0.0.0
 no au
```

R4

```
interface fastethernet 0/0
 ip address 192.168.14.1 255.255.255.0
 no shutdown
!
interface Serial1/0
 ip address 10.0.34.2 255.255.255.252
 no shutd
```

```
!
router eig 100
 net 10.0.34.2 0.0.0.0
 no au
```

The adjacencies are successful as shown by the following outputs:

```
R1#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                   (sec)         (ms)       Cnt
Num
0   10.0.12.2               Se1/1                    14 00:01:37  159    954  0   9
R1#
```

```
R2#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                   (sec)         (ms)       Cnt
Num
1   10.0.12.1               Se1/1                    14 00:03:30  187   1122  0  3
0   10.0.23.2               Se1/0                    12 00:04:25  211   1266  0  11
R2#
```

```
R3#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                   (sec)         (ms)       Cnt
Num
1   10.0.34.2               Se1/2                    12 00:04:35  150    900  0   5
0   10.0.23.1               Se1/1                    11 00:04:39   86    516  0  10
R3#
```

```
R4#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                   (sec)         (ms)       Cnt
Num
0   10.0.34.1               Se1/0                    12 00:04:49   74    444  0  12
R4#
```

Task 2

The subnet 192.168.14.0/24 is advertised as an external EIGRP Route by setting a route tag of 11 by R1 and 44 by R4:

R1:
```
route-map CONNECTED permit 10
 match interface fastethernet 0/0
 set tag 11
!
```

```
router eigrp 100
 redistribute connected route-map CONNECTED
```

```
R4:
route-map CONNECTED permit 10
 match interface fastethernet 0/0
 set tag 44
router eigrp 100
 redistribute connected route-map CONNECTED
```

R2 and R3 receive an external EIGRP Route for the subnet 192.168.14.0/24 as shown by the routing table displayed below:

```
R2#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
+ - replicated route, % - next hop override
Gateway of last resort is not set

10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
D 10.0.34.0/30 [90/2681856] via 10.0.23.2, 00:23:51, Serial1/0
D EX 192.168.14.0/24 [170/2172416] via 10.0.12.1, 00:09:19, Serial1/1
```

```
R3#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
+ - replicated route, % - next hop override
Gateway of last resort is not set

10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
D 10.0.12.0/30 [90/2681856] via 10.0.23.1, 00:25:56, Serial1/1
D EX 192.168.14.0/24 [170/2172416] via 10.0.34.2, 00:12:17, Serial1/2
```

Task 3

Configure four secondary ip addresses on R1's Fa0/0 and R4's Fa0/0 and avertise these subnets via EIGRP as internal EIGRP Routes as follow:

```
R1:
interface fastethernet 0/0
 ip address 10.14.0.1 255.255.255.0 secondary
 ip address 10.14.1.1 255.255.255.0 secondary
 ip address 10.14.2.1 255.255.255.0 secondary
 ip address 10.14.3.1 255.255.255.0 secondary
!
```

```
router eigrp 100
 network 10.14.0.0 0.0.3.255
```

R4:

```
interface fastethernet 0/0
 ip address 10.14.0.4 255.255.255.0 secondary
 ip address 10.14.1.4 255.255.255.0 secondary
 ip address 10.14.2.4 255.255.255.0 secondary
 ip address 10.14.3.4 255.255.255.0 secondary
!
router eigrp 100
 network 10.14.0.0 0.0.3.255
```

We can verify that the prefixes representing the secondary ip addresses are advertised to R2 and R3 and installed in the routing table as shown by the following output:

```
R2#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
+ - replicated route, % - next hop override
Gateway of last resort is not set

10.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
D 10.0.34.0/30 [90/2681856] via 10.0.23.2, 00:00:24, Serial1/0
D 10.14.0.0/24 [90/2172416] via 10.0.12.1, 00:00:27, Serial1/1
D 10.14.1.0/24 [90/2172416] via 10.0.12.1, 00:00:27, Serial1/1
D 10.14.2.0/24 [90/2172416] via 10.0.12.1, 00:00:27, Serial1/1
D 10.14.3.0/24 [90/2172416] via 10.0.12.1, 00:00:27, Serial1/1
D EX 192.168.14.0/24 [170/2172416] via 10.0.12.1, 00:15:49, Serial1/1
```

```
R3#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
+ - replicated route, % - next hop override
Gateway of last resort is not set

10.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
D 10.0.12.0/30 [90/2681856] via 10.0.23.1, 00:00:32, Serial1/1
D 10.14.0.0/24 [90/2172416] via 10.0.34.2, 00:00:34, Serial1/2
D 10.14.1.0/24 [90/2172416] via 10.0.34.2, 00:00:34, Serial1/2
D 10.14.2.0/24 [90/2172416] via 10.0.34.2, 00:00:34, Serial1/2
D 10.14.3.0/24 [90/2172416] via 10.0.34.2, 00:00:34, Serial1/2
D EX 192.168.14.0/24 [170/2172416] via 10.0.34.2, 00:15:57, Serial1/2
```

Task 4

Ensure that R1 and R2 will advertise a summary route for the 10.14.0.0/24, 10.14.1.0/24, 10.14.2.0/24, and 10.14.3.0/24 secondary subnets:

R1

```
interface serial 1/1
 ip summary-address eigrp 100 10.14.0.0 255.255.252.0
!
interface serial 1/0
 ip summary-address eigrp 100 10.14.0.0 255.255.252.0
```

The following output shown that R2 and R3 install a single summary route for the 10.14.0.0/24, 10.14.1.0/24, 10.14.2.0/24, and 10.14.3.0/24 secondary subnets:
-R2 prefers R1 as a next-hop to reach the secondary subnets.
-R4 prefers R4 as a next-hop to reach the secondary subnets.

```
R2#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
+ - replicated route, % - next hop override
Gateway of last resort is not set

10.0.0.0/8 is variably subnetted, 6 subnets, 3 masks
D 10.0.34.0/30 [90/2681856] via 10.0.23.2, 00:03:41, Serial1/0
D 10.14.0.0/22 [90/2172416] via 10.0.12.1, 00:00:16, Serial1/1
D EX 192.168.14.0/24 [170/2172416] via 10.0.12.1, 00:19:06, Serial1/1
```

```
R3#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
+ - replicated route, % - next hop override
Gateway of last resort is not set

10.0.0.0/8 is variably subnetted, 6 subnets, 3 masks
D 10.0.12.0/30 [90/2681856] via 10.0.23.1, 00:03:46, Serial1/1
D 10.14.0.0/22 [90/2172416] via 10.0.34.2, 00:00:21, Serial1/2
D EX 192.168.14.0/24 [170/2172416] via 10.0.34.2, 00:19:11, Serial1/2
R3#
```

Now if we want that R2 and R3 prefer the path via R1 to reach the 10.14.0.0/24 subnet, and the path via R4 to reach the 10.14.3.0/24 subnet,we will use an EIGRP feature called route leaking. When summarizing, EIGRP allows you to 'leak' or advertise more specific routes encompassed by that summary, ensuring that the traffic destined to the leaked subnet is sent to the router that leaked the route based on the longest-match rule:

R1:

```
interface serial 1/1
 ip summary-a eigrp 100 10.14.0.0 255.255.252.0 leak-map LEAK
!
```

```
access-list 1 permit 10.14.0.0 0.0.0.255
 route-map LEAK permit 10
 match ip address 1
!
route-map LEAK deny 20
```

R4:

```
interface serial 1/0
 ip summary- eigrp 100 10.14.0.0 255.255.252.0 leak-map LEAK
!
access-list 4 permit 10.14.3.0 0.0.0.255
 route-map LEAK permit 10
 match ip address 4
!
route-map LEAK deny 20
```

As a result:
-R2 and R3 prefer the path via R1 to reach the 10.14.0.0/24 subnet
On R2 the next-hop:10.0.12.1 is the ip address of R1,R2 passes through R1.
On R3 the next-hop:10.0.23.1 is the ip address of R2,R3 passes through R2--R1

-R2 and R3 prefer the path via R4 to reach the 10.14.3.0/24 subnet
On R3 the next-hop:10.0.34.2 is the ip address of R4,R3 passes through R4
On R2 the next-hop:10.0.23.2 is the ip address of R4,R3 passes through R3--R4

```
R2#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
+ - replicated route, % - next hop override
Gateway of last resort is not set

10.0.0.0/8 is variably subnetted, 8 subnets, 4 masks
D 10.0.34.0/30 [90/2681856] via 10.0.23.2, 00:12:39, Serial1/0
D 10.14.0.0/22 [90/2172416] via 10.0.12.1, 00:09:14, Serial1/1
D 10.14.0.0/24 [90/2172416] via 10.0.12.1, 00:00:37, Serial1/1
D 10.14.3.0/24 [90/2684416] via 10.0.23.2, 00:00:21, Serial1/0
D EX 192.168.14.0/24 [170/2172416] via 10.0.12.1, 00:28:04, Serial1/1
R2#
```

```
R3#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
+ - replicated route, % - next hop override
Gateway of last resort is not set

10.0.0.0/8 is variably subnetted, 8 subnets, 4 masks
D 10.0.12.0/30 [90/2681856] via 10.0.23.1, 00:13:09, Serial1/1
D 10.14.0.0/22 [90/2172416] via 10.0.34.2, 00:09:44, Serial1/2
```

```
D 10.14.0.0/24 [90/2684416] via 10.0.23.1, 00:01:06, Serial1/1
D 10.14.3.0/24 [90/2172416] via 10.0.34.2, 00:00:51, Serial1/2
D EX 192.168.14.0/24 [170/2172416] via 10.0.34.2, 00:28:34, Serial1/2
```

Task 5

Now the purpose is to increase the default EIGRP hop count limitation of 100 to 200 for scalability in future.
To display the default hop count we will use the show ip protocols command,notice the line "Maximum hopcount 100":

```
R1#show ip protocols
*** IP Routing is NSF aware ***
Routing Protocol is "eigrp 100"
Outgoing update filter list for all interfaces is not set
Incoming update filter list for all interfaces is not set
Default networks flagged in outgoing updates
Default networks accepted from incoming updates
Redistributing: connected
EIGRP-IPv4 Protocol for AS(100)
Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
NSF-aware route hold timer is 240
Router-ID: 192.168.14.1
Topology : 0 (base)
Active Timer: 3 min
Distance: internal 90 external 170
Maximum path: 4
Maximum hopcount 100
Maximum metric variance 1
Default redistribution metric is 1 1 1 1 1500
Automatic Summarization: disabled
Address Summarization:
10.14.0.0/22 for Se1/1
Summarizing 4 components with metric 28160
Maximum path: 4
Routing for Networks:
10.0.12.1/32
10.14.0.0/22
Routing Information Sources:
Gateway Distance Last Update
10.0.12.2 90 00:12:14
192.168.14.4 90 00:21:07
Distance: internal 90 external 170
```

to configure a hop count limit of 200 for ALL routers ,we will use the metric maximum-hops command:

Rx:

```
router eigrp 100
 metric maximum-hops 200
```

The show ip protocols command tell us that the hop count limit is set to 200:

```
R1#show ip protocols
*** IP Routing is NSF aware ***
Routing Protocol is "eigrp 100"
Outgoing update filter list for all interfaces is not set
```

```
Incoming update filter list for all interfaces is not set
Default networks flagged in outgoing updates
Default networks accepted from incoming updates
Redistributing: connected
EIGRP-IPv4 Protocol for AS(100)
Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
NSF-aware route hold timer is 240
Router-ID: 192.168.14.1
Topology : 0 (base)
Active Timer: 3 min
Distance: internal 90 external 170
Maximum path: 4
Maximum hopcount 200
Maximum metric variance 1
Default redistribution metric is 1 1 1 1 1500
Automatic Summarization: disabled
Address Summarization:
10.14.0.0/22 for Se1/1
Summarizing 4 components with metric 28160
Maximum path: 4
Routing for Networks:
10.0.12.1/32
10.14.0.0/22
Routing Information Sources:
Gateway Distance Last Update
10.0.12.2 90 00:00:17
192.168.14.4 90 00:00:18
Distance: internal 90 external 170
```

Lab 7: EIGRP Default Route Summarization Authentication

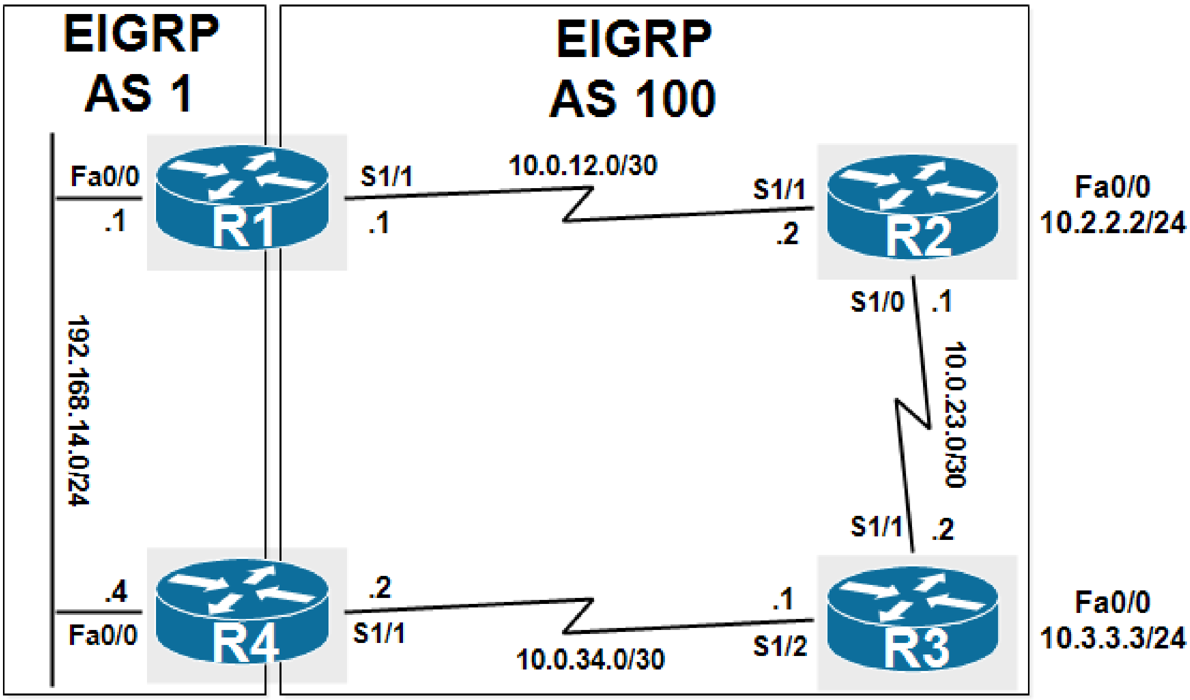

Task 1

Configure hostnames and IP addressing on all routers as illustrated in the network topology.

Task 2

Configure EIGRP for AS 1 as illustrated in the topology. However, do NOT advertise the 10.2.2.0/24 and 10.3.3.0/24 subnets connected to R2 and R3 via EIGRP, use redistribution. Additionally, authenticate EIGRP protocol updates using the password 'CCNP'.

Task 3

Configure routers R2 and R3 to advertise ONLY the default route to R1 and R4. Ensure that R2 and R3 can ping each others' 10.x.x.x/24 subnets. Verify that R1 and R4 can reach the 10.2.2.0/24 and 10.3.3.0/24 subnets even if their own WAN links are down.

Task 4

Configure the following secondary subnets on the LAN segment between R1 and R4:

1. Network: 192.168.0.0/24 - Assign 192.168.0.1/24 to R1 and 192.168.0.4/24 to R4
2. Network: 192.168.1.0/24 - Assign 192.168.1.1/24 to R1 and 192.168.1.4/24 to R4
3. Network: 192.168.2.0/24 - Assign 192.168.2.1/24 to R1 and 192.168.2.4/24 to R4
4. Network: 192.168.3.0/24 - Assign 192.168.3.1/24 to R1 and 192.168.3.4/24 to R4

Next, configure R1 and R4 to advertise only a single route for these subnets to R2 and R3.

Task 1

Basic configuration of all routers:

R1

```
interface FastEthernet0/0
 ip address 192.168.14.1 255.255.255.0
 no shutdown
!
interface Serial1/1
 ip address 10.0.12.1 255.255.255.252
 no shutdown
```

R2

```
interface FastEthernet0/0
 ip address 10.2.2.2 255.255.255.0
 no shutdown
!
interface Serial1/0
 ip address 10.0.23.1 255.255.255.252
 no shutdown
!
interface Serial1/1
 ip address 10.0.12.2 255.255.255.252
 no shutdown
```

R3

```
interface FastEthernet0/0
 ip address 10.3.3.3 255.255.255.0
 no shutdown
!
interface Serial1/1
 ip address 10.0.23.2 255.255.255.252
 no shutdown
!
interface Serial1/2
 ip address 10.0.34.1 255.255.255.252
 no shutdown
```

R4

```
interface FastEthernet0/0
 ip address 192.168.14.4 255.255.255.0
 no shutdown
!
interface Serial1/0
 ip address 10.0.34.2 255.255.255.252
 no shutdown
```

Task 2

Configure the routers to run EIGRP AS 100 .
The LANs Subnet 10.2.2.0/24 and 10.3.3.0/24 connected to R2 and R3 respectively will be redistributed as external routes.

R1

```
router eigrp 100
 network 192.168.14.1 0.0.0.0
 network 10.0.12.1 0.0.0.0
```

```
 no auto-summary
```

R2

```
router eigrp 100
 network 10.0.12.2 0.0.0.0
 network 10.0.23.1 0.0.0.0
 no auto-summary
 redistribute connected
```

R3

```
router eigrp 100
 no auto-summary
 network 10.0.23.2 0.0.0.0
 network 10.0.34.1 0.0.0.0
 redistribute connected
```

R4

```
router eigrp 100
 no auto-summary
 network 192.168.14.4 0.0.0.0
 network 10.0.34.2 0.0.0.0
```

Let's verify the neighbor relationship:

```
R1#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                   (sec)         (ms)       Cnt
Num
1   192.168.14.4            Fa0/0                    13 00:03:10  137    822  0   3
0   10.0.12.2               Se1/1                    13 00:03:40  257   1542  0   7
R1#
```

```
R2#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                   (sec)         (ms)       Cnt
Num
1   10.0.23.2               Se1/0                    14 00:05:01  289   1734  0   3
0   10.0.12.1               Se1/1                    12 00:05:19  270   1620  0   8
R2#
```

```
R3#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                   (sec)         (ms)       Cnt
Num
0   10.0.23.1               Se1/1                    11 00:06:28  467   2802  0   6
R3#
```

```
R4#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
```

```
                                                  (sec)         (ms)       Cnt
Num
0   192.168.14.1            Fa0/0                   13 00:09:50 1579  5000  0  7
R4#
```

Configure authentication of EIGRP in order to secure the Updates.

We will use the password CISCO:

R1

```
key chain CISCO-CHAIN
 key 1
 key-string CISCO
!
interface fastethernet 0/0
 ip authentication mode eigrp 100 md5
 ip authentication key-chain eigrp 100 CISCO-CHAIN
!
interface serial 1/1
 ip authentication mode eigrp 100 md5
 ip authentication key-chain eigrp 100 CISCO-CHAIN
```

R2

```
key chain CISCO-CHAIN
 key 1
 key-string CISCO
!
interface serial 1/1
 ip authentication mode eigrp 100 md5
 ip authentication key-chain eigrp 100 CISCO-CHAIN
!
interface serial 1/0
 ip authentication mode eigrp 100 md5
 ip authentication key-chain eigrp 100 CISCO-CHAIN
```

R3

```
key chain CISCO-CHAIN
 key 1
 key-string CISCO
!
interface serial 1/1
 ip authentication mode eigrp 100 md5
 ip authentication key-chain eigrp 100 CISCO-CHAIN
!
interface serial 1/2
 ip authentication mode eigrp 100 md5
 ip authentication key-chain eigrp 100 CISCO-CHAIN
```

R4

```
key chain CISCO-CHAIN
 key 1
 key-string CISCO
!
interface fastethernet 0/0
 ip authentication mode eigrp 100 md5
 ip authentication key-chain eigrp 100 CISCO-CHAIN
!
```

```
interface serial 1/1
 ip authentication mode eigrp 100 md5
 ip authentication key-chain eigrp 100 CISCO-CHAIN
```

Let's check the authentication:

```
R1#show key chain CISCO-CHAIN
Key-chain CISCO-CHAIN:
    key 1 -- text "CISCO"
        accept lifetime (always valid) - (always valid) [valid now]
        send lifetime (always valid) - (always valid) [valid now]
```

The show ip eigrp interfaces details command display the interfaces configured to authenticate the EIGRP updates:

```
R1#show ip eigrp interfaces detail
EIGRP-IPv4 Interfaces for AS(100)
                              Xmit Queue   PeerQ        Mean   Pacing Time
Multicast    Pending
Interface               Peers  Un/Reliable  Un/Reliable  SRTT   Un/Reliable   Flow
Timer    Routes
Fa0/0                     1        0/0        0/0         137       0/0          560
0
  Hello-interval is 5, Hold-time is 15
  Split-horizon is enabled
  Next xmit serial <none>
  Packetized sent/expedited: 2/0
  Hello's sent/expedited: 237/2
  Un/reliable mcasts: 0/2  Un/reliable ucasts: 2/2
  Mcast exceptions: 0  CR packets: 0  ACKs suppressed: 0
  Retransmissions sent: 1  Out-of-sequence rcvd: 1
  Topology-ids on interface - 0
  Authentication mode is md5,  key-chain is "CISCO-CHAIN"
Se1/1                     1        0/0        0/0         257       0/16         924
0
  Hello-interval is 5, Hold-time is 15
  Split-horizon is enabled
  Next xmit serial <none>
  Packetized sent/expedited: 4/0
  Hello's sent/expedited: 237/2
  Un/reliable mcasts: 0/0  Un/reliable ucasts: 4/5
  Mcast exceptions: 0  CR packets: 0  ACKs suppressed: 0
  Retransmissions sent: 0  Out-of-sequence rcvd: 0
  Topology-ids on interface - 0
  Authentication mode is md5,  key-chain is "CISCO-CHAIN"
```

Task 3

Configure the routers R2 and R3 to advertise the default route to R1 and R4.

R2

```
interface serial 1/1
 ip summary-address eigrp 100 0.0.0.0 0.0.0.0
```

R3

```
interface serial 1/2
 ip summary-address eigrp 100 0.0.0.0 0.0.0.0
```

R1 and R4 are not allowed to receive the external route EIGRP to 10.2.2.0/24 and 10.3.3.0/24 therefore We will use prefix-list called DEFAULT-ROUTE to allow only a default route to R1 and R4 and a distribute-list command is configured under router eigrp 100 with the prefix-list out the interfaces s1/1 and s1/2 of R2 and R3 respectively. All other routes that do not match the sequence 5 of the prefix-list are denied by default with an implicit deny (in this case (10.2.2.0/24 and 10.3.3.0/24 ,also 10.0.23/24).

R2

```
ip prefix-list DEFAULT-ROUTE seq 5 permit 0.0.0.0/0
!
router eigrp 100
 distribute-list prefix DEFAULT-ROUTE out serial 1/1
```

R3

```
ip prefix-list DEFAULT-ROUTE seq 5 permit 0.0.0.0/0
!
router eigrp 100
 distribute-list prefix DEFAULT-ROUTE out serial 1/2
```

Let's check the routing tables of all routers:

R1 receives only a default route from R2 and an EIGRP route to 10.0.34.0/24 from R4, the subnets 10.2.2.0/24,10.3.3.0/24 and 10.0.23/24 are filtered by the prefix-list:

```
R1#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is 10.0.12.2 to network 0.0.0.0

D*    0.0.0.0/0 [90/2172416] via 10.0.12.2, 00:00:34, Serial1/1
      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D        10.0.34.0/30 [90/2172416] via 192.168.14.4, 00:00:42, FastEthernet0/0
```

R2 receives an external EIGRP route to 10.3.3.0/24 from R3 as expected because R3 filters this subnet only out the interface s1/2 toward R1:

```
R2#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is 0.0.0.0 to network 0.0.0.0

D*    0.0.0.0/0 is a summary, 00:01:43, Null0
      10.0.0.0/8 is variably subnetted, 8 subnets, 3 masks
D        10.0.34.0/30 [90/2681856] via 10.0.23.2, 00:31:20, Serial1/0
```

```
D EX     10.3.3.0/24 [170/2172416] via 10.0.23.2, 00:01:37, Serial1/0
D     192.168.14.0/24 [90/2172416] via 10.0.12.1, 00:31:57, Serial1/1
R2#
```

R3 receives an external EIGRP route to 10.2.2.0/24 from R2 as expected because R2 filters this subnet only out the interface s1/1 toward R1:

```
R3#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is 0.0.0.0 to network 0.0.0.0

D*    0.0.0.0/0 is a summary, 00:03:32, Null0
      10.0.0.0/8 is variably subnetted, 8 subnets, 3 masks
D        10.0.12.0/30 [90/2681856] via 10.0.23.1, 00:33:45, Serial1/1
D EX     10.2.2.0/24 [170/2172416] via 10.0.23.1, 00:03:58, Serial1/1
D     192.168.14.0/24 [90/2684416] via 10.0.23.1, 00:33:45, Serial1/1
R3#
```

R4 receives only a default route from R2 and an EIGRP route to 10.0.12.0/24 from R4, the subnets 10.2.2.0/24,10.3.3.0/24 and 10.0.23/24 are filtered by the prefix-list:

```
R4#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is 192.168.14.1 to network 0.0.0.0

D*    0.0.0.0/0 [90/2174976] via 192.168.14.1, 00:05:35, FastEthernet0/0
      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D        10.0.12.0/30 [90/2172416] via 192.168.14.1, 00:35:18, FastEthernet0/0
R4#
```

Task 4

Add secondary ip addresses on the Lan interfaces of R1 and R4 and let's advertise these subnet into EIGRP.Then we configure R1 and R4 to advertise only a summarized route for these subnets to R2 and R3:

1.Network: 192.168.0.0/24 - 192.168.0.1/24 to R1 and 192.168.0.4/24 to R4
2.Network: 192.168.1.0/24 - 192.168.1.1/24 to R1 and 192.168.1.4/24 to R4
3.Network: 192.168.2.0/24 - 192.168.2.1/24 to R1 and 192.168.2.4/24 to R4
4.Network: 192.168.3.0/24 - 192.168.3.1/24 to R1 and 192.168.3.4/24 to R4

R1

```
interface fastethernet 0/0
 ip address 192.168.0.1 255.255.255.0 secondary
 ip address 192.168.1.1 255.255.255.0 secondary
 ip address 192.168.2.1 255.255.255.0 secondary
 ip address 192.168.3.1 255.255.255.0 secondary
!
interface serial 1/1
 ip summary-address eigrp 100 192.168.0.0 255.255.252.0
!
router eigrp 100
 network 192.168.0.0 0.0.3.255
```

R2

```
interface fastethernet 0/0
 ip address 192.168.0.4 255.255.255.0 secondary
 ip address 192.168.1.4 255.255.255.0 secondary
 ip address 192.168.2.4 255.255.255.0 secondary
 ip address 192.168.3.4 255.255.255.0 secondary
!
interface serial 1/0
 ip summary-address eigrp 100 192.168.0.0 255.255.252.0
!
router eigrp 100
 network 192.168.0.0 0.0.3.255
```

R2 and R3 are receiving a summary route toward the secondaries ip addresses:

```
R2#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is 0.0.0.0 to network 0.0.0.0

D*    0.0.0.0/0 is a summary, 00:35:51, Null0
      10.0.0.0/8 is variably subnetted, 8 subnets, 3 masks
D        10.0.34.0/30 [90/2681856] via 10.0.23.2, 01:05:28, Serial1/0
D EX     10.3.3.0/24 [170/2172416] via 10.0.23.2, 00:35:45, Serial1/0
D     192.168.0.0/22 [90/2172416] via 10.0.12.1, 00:19:58, Serial1/1
D     192.168.14.0/24 [90/2172416] via 10.0.12.1, 01:06:05, Serial1/1
R2#
```

Let's test connectivity to these subnets:

```
R2#ping 192.168.0.0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.0.0, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/94/132 ms
R2#
R2#ping 192.168.1.0
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 192.168.1.0, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 68/97/144 ms
R2#
R2#ping 192.168.2.0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.2.0, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 88/113/132 ms
R2#
R2#ping 192.168.3.0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.3.0, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 96/122/168 ms
R2#
```

```
R3#ping 192.168.0.0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.0.0, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 100/134/188 ms
R3#
R3#ping 192.168.1.0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.0, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 80/120/152 ms
R3#
R3#ping 192.168.2.0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.2.0, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 116/137/172 ms
R3#
R3#ping 192.168.3.0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.3.0, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/148/180 ms
R3#
```

Lab 8: EIGRP Named Mode for IPv4 and IPv6

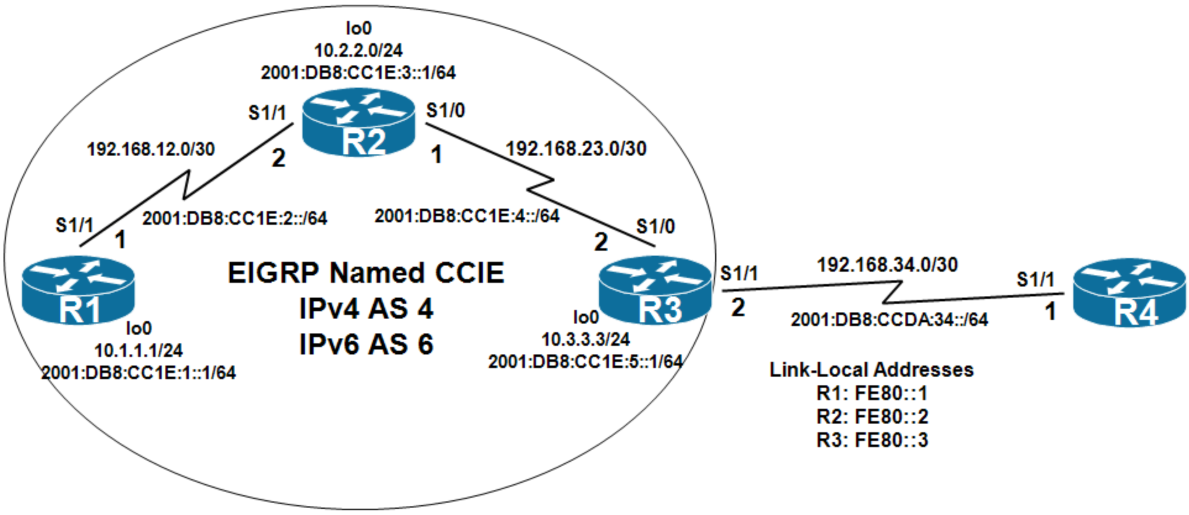

Basic configuration of all routers:

R1:

```
interface lo0
 ip address 10.1.1.1 255.255.255.0
 ipv6 address FE80::1 link-local
 ipv6 address 2001:DB8:CC1E:1::1/64
 no shutdown
!
interface Serial1/1
 ip address 192.168.12.1 255.255.255.252
 ipv6 address FE80::1 link-local
 ipv6 address 2001:DB8:CC1E:2::1/64
 no shutdown
```

R2:

```
interface lo0
 ip address 10.2.2.2 255.255.255.0
 ipv6 address FE80::2 link-local
 ipv6 address 2001:DB8:CC1E:3::1/64
 no shutdown
!
interface Serial1/1
 ip address 192.168.12.2 255.255.255.252
 ipv6 address FE80::2 link-local
 ipv6 address 2001:DB8:CC1E:2::2/64
 no shutdown
!
interface Serial1/0
 ip address 192.168.23.1 255.255.255.252
 ipv6 address FE80::2 link-local
 ipv6 address 2001:DB8:CC1E:4::1/64
 no shutdown
```

R3:

```
interface lo0
```

```
 ip address 10.3.3.3 255.255.255.0
 ipv6 address FE80::3 link-local
 ipv6 address 2001:DB8:CC1E:5::1/64
 no shutdown
!
interface Serial1/0
 ip address 192.168.23.2 255.255.255.252
 ipv6 address FE80::3 link-local
 ipv6 address 2001:DB8:CC1E:4::2/64
 no shutdown
!
interface Serial1/1
 ip address 192.168.34.2 255.255.255.0
 ipv6 address FE80::3 link-local
 ipv6 address 2001:DB8:CCDA:34::2/64
 no shutdown
```

R4:

```
interface Serial1/1
 ip address 192.168.34.1 255.255.255.0
 ipv6 address FE80::4 link-local
 ipv6 address 2001:DB8:CCDA:34::1/64
 no shutdown
!
ipv6 route 2001:DB8:CC1E::/48 2001:DB8:FEED:77::2
ip route 0.0.0.0 0.0.0.0 192.168.34.2
```

To configure the IPv4 and IPv6 address families and autonomous system we use the address-family ipv4 unicast autonomous-system 4 and the address-family ipv6 unicast autonomous-system 6 commands.
The Named mode is called CCIE.
The IPv4 address family use AS 4 and the IPv6 address family use AS 6.
Configure the EIGRP router ID for the IPv4 and IPv6 address families using the eigrp router-id command.
Enable EIGRP on the interfaces using the network command in the address family configuration mode for IPv4.
For IPv6 address family by default, all IPv6 interfaces are automatically enabled for EIGRP for IPv6.

```
R1(config)#router eigrp CCIE
R1(config-router)#address-family ipv4 unicast autonomous-system 4
R1(config-router-af)#eigrp router-id 1.1.1.1
R1(config-router-af)#network 10.1.1.0 0.0.0.255
R1(config-router-af)#network 192.168.12.0 0.0.0.3
```

```
R1(config)#router eigrp CCIE
R1(config-router)#address-family ipv6 unicast autonomous-system 6
R1(config-router-af)#eigrp router-id 1.1.1.1
```

```
R2(config)# router eigrp CCIE
R2(config-router)# address-family ipv4 unicast autonomous-system 4
R2(config-router-af)# eigrp router-id 2.2.2.2
R2(config-router-af)# network 192.168.12.0 0.0.0.3
R2(config-router-af)# network 10.2.2.0 0.0.0.255
R2(config-router-af)# network 192.168.23.0 0.0.0.3
```

```
R2(config)#router eigrp CCIE
R2(config-router)#address-family ipv6 unicast autonomous-system 6
```

```
R2(config-router-af)#eigrp router-id 2.2.2.2
```

```
R3(config)# router eigrp CCIE
R3(config-router)# address-family ipv4 unicast autonomous-system 4
R3(config-router-af)# eigrp router-id 3.3.3.3
R3(config-router-af)# network 192.168.23.0 0.0.0.3
R3(config-router-af)# network 10.3.3.0 0.0.0.255
R3(config-router-af)# exit-address-family
R3(config-router)# address-family ipv6 unicast autonomous-system 6
R3(config-router-af)# eigrp router-id 3.3.3.3
```

Verify the neighbor relationship:

```
R2#show ip eigrp neighbors
EIGRP-IPv4 VR(CCIE) Address-Family Neighbors for AS(4)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                   (sec)         (ms)       Cnt
Num
1   192.168.23.2            Se1/0                    13 00:31:02  282   1692  0   4
0   192.168.12.1            Se1/1                    11 00:31:02  314   1884  0   4
R2#
```

```
R2#show ipv6 eigrp neighbors
EIGRP-IPv6 VR(CCIE) Address-Family Neighbors for AS(6)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                   (sec)         (ms)       Cnt
Num
1   Link-local address:     Se1/1                    14 00:31:17 1429   5000  0   6
    FE80::1
0   Link-local address:     Se1/0                    14 00:31:17 1385   5000  0   4
    FE80::3
R2#
```

Within each IPv4 and IPv6 AF is the address family interface configuration mode. This mode is used to configure EIGRP specific parameters on an interface, such as the hello timer and summarization. From address family configuration mode, use the af-interface interface-type interface-number command to enter address family interface configuration mode. Then use the passive-interface command in order to configure the loopback interface as passive for both the IPv4 and IPv6 EIGRP address families:

```
R1(config)# router eigrp CCIE
R1(config-router)# address-family ipv4 unicast autonomous-system 4
R1(config-router-af)# af-interface Lo0
R1(config-router-af-interface)# passive-interface
R1(config-router-af)# exit-address-family
R1(config-router)# address-family ipv6 unicast autonomous-system 6
R1(config-router-af)# af-interface lo0
R1(config-router-af-interface)# passive-interface
```

```
R2(config)# router eigrp CCIE
R2(config-router)# address-family ipv4 unicast autonomous-system 4
R2(config-router-af)# af-interface lo0
R2(config-router-af-interface)# passive-interface
R2(config-router-af-interface)# exit-af-interface
R2(config-router-af)# exit-address-family
R2(config-router)# address-family ipv6 unicast autonomous-system 6
```

```
R2(config-router-af)# af-interface lo0
R2(config-router-af-interface)# passive-interface
```

```
R3(config)# router eigrp CCIE
R3(config-router)# address-family ipv4 unicast autonomous-system 4
R3(config-router-af)# af-interface lo0
R3(config-router-af-interface)# passive-interface
R3(config-router-af-interface)# exit-af-interface
R3(config-router-af)# exit-address-family
R3(config-router)# address-family ipv6 unicast autonomous-system 6
R3(config-router-af)# af-interface lo0
R3(config-router-af-interface)# passive-interface
```

Disable named EIGRP on a specific IPv6 interface.

As we said previously by default, all IPv6 interfaces are enabled for EIGRP for IPv6. This happens when enabling the IPv6 address family with the address-family ipv6 unicast autonomous-system command.
We can see that on R3 all three of its IPv6 interfaces are enabled for EIGRP for IPv6. Notice that the S1/1 interface is also included using the show ipv6 protocols command on R3:

```
R3#show ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"
IPv6 Routing Protocol is "eigrp 6"
EIGRP-IPv6 VR(CCIE) Address-Family Protocol for AS(6)
  Metric weight K1=1, K2=0, K3=1, K4=0, K5=0 K6=0
  Metric rib-scale 128
  Metric version 64bit
  NSF-aware route hold timer is 240
  Router-ID: 3.3.3.3
  Topology : 0 (base)
    Active Timer: 3 min
    Distance: internal 90 external 170
    Maximum path: 16
    Maximum hopcount 100
    Maximum metric variance 1
    Total Prefix Count: 6
    Total Redist Count: 0

  Interfaces:
    Serial1/0
    Serial1/1
    Loopback0 (passive)
  Redistribution:
    None
R3#
```

In this topology R3's S1/1 interface does not need to be included in the EIGRP updates instead a default route will be configured in order to have a reachability beyond the EIGRP routing domain.
With the IPv4 AF we excluded the network command for interface s1/1. However, the same interface is automatically included when configuring the IPv6 AF. The shutdown command under the address family interface is used to disable EIGRP on a specific interface. This does not disable the physical interface, but only removes it from participating in EIGRP.

```
R3(config)# router eigrp CCIE
R3(config-router)# address-family ipv6 unicast autonomous-system 6
R3(config-router-af)# af-interface serial 1/1
```

```
R3(config-router-af-interface)# shutdown
```

We can verify that R3 is no longer including S1/1 in EIGRP for IPv6 using the show ipv6 protocols:

```
R3#show ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"
IPv6 Routing Protocol is "eigrp 6"
EIGRP-IPv6 VR(CCIE) Address-Family Protocol for AS(6)
  Metric weight K1=1, K2=0, K3=1, K4=0, K5=0 K6=0
  Metric rib-scale 128
  Metric version 64bit
  NSF-aware route hold timer is 240
  Router-ID: 3.3.3.3
  Topology : 0 (base)
    Active Timer: 3 min
    Distance: internal 90 external 170
    Maximum path: 16
    Maximum hopcount 100
    Maximum metric variance 1
    Total Prefix Count: 5
    Total Redist Count: 0

  Interfaces:
    Serial1/0
    Loopback0 (passive)
  Redistribution:
    None
R3#
```

To have the reachability beyond the EIGRP routing domain configure a static default route and redistribute it into EIGRP domain:
In named EIGRP redistribution of static routes is done in topology configuration mode. Topology configuration mode is a subset of an address family. By default, EIGRP has a base topology for each address family.
For each address family, use the topology base command to enter the base EIGRP topology. In topology configuration mode use the redistribute static command to redistribute the default static route into EIGRP.

```
R3(config)# ip route 0.0.0.0 0.0.0.0 192.168.34.1
R3(config)# ipv6 route ::/0 serial1/1 2001:db8:CCDA:34::1
```

```
R3(config)# router eigrp CCIE
R3(config-router)# address-family ipv4 unicast autonomous-system 4
R3(config-router-af)# topology base
R3(config-router-af-topology)# redistribute static
R3(config-router-af-topology)# exit-af-topology
R3(config-router-af)# exit-address-family
R3(config-router)# address-family ipv6 unicast autonomous-system 6
R3(config-router-af)# topology base
R3(config-router-af-topology)# redistribute static
R3(config-router-af-topology)# exit-af-topology
R3(config-router-af)# exit-address-family
```

To verify that EIGRP is redistributing the static route use the show ip protocols and show ipv6 protocols commands on R3:

```
R3#show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "eigrp 4"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  Redistributing: static
  EIGRP-IPv4 VR(CCIE) Address-Family Protocol for AS(4)
    Metric weight K1=1, K2=0, K3=1, K4=0, K5=0 K6=0
    Metric rib-scale 128
    Metric version 64bit
    NSF-aware route hold timer is 240
    Router-ID: 3.3.3.3
    Topology : 0 (base)
      Active Timer: 3 min
      Distance: internal 90 external 170
      Maximum path: 4
      Maximum hopcount 100
      Maximum metric variance 1
      Total Prefix Count: 6
      Total Redist Count: 1

  Automatic Summarization: disabled
  Maximum path: 4
  Routing for Networks:
    10.3.3.0/24
    192.168.23.0/30
  Passive Interface(s):
    Loopback0
  Routing Information Sources:
    Gateway         Distance      Last Update
    192.168.23.1          90      00:24:02
  Distance: internal 90 external 170

R3#
```

```
R3#show ipv proto
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"
IPv6 Routing Protocol is "eigrp 6"
EIGRP-IPv6 VR(CCIE) Address-Family Protocol for AS(6)
  Metric weight K1=1, K2=0, K3=1, K4=0, K5=0 K6=0
  Metric rib-scale 128
  Metric version 64bit
  NSF-aware route hold timer is 240
  Router-ID: 3.3.3.3
  Topology : 0 (base)
    Active Timer: 3 min
    Distance: internal 90 external 170
    Maximum path: 16
    Maximum hopcount 100
    Maximum metric variance 1
    Total Prefix Count: 6
    Total Redist Count: 1

  Interfaces:
    Serial1/0
```

```
    Loopback0 (passive)
  Redistribution:
    Redistributing protocol static
IPv6 Routing Protocol is "static"
R3#
```

Verify the IPv4 and IPv6 routing tables on R1 to verify that it is receiving the default static route:

```
R1#show ip route eigrp | beg Gate
Gateway of last resort is 192.168.12.2 to network 0.0.0.0

D*EX  0.0.0.0/0 [170/34036062] via 192.168.12.2, 00:06:49, Serial1/1
      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
D        10.2.2.0/24 [90/13556702] via 192.168.12.2, 00:31:01, Serial1/1
D        10.3.3.0/24 [90/23796702] via 192.168.12.2, 00:30:00, Serial1/1
      192.168.23.0/30 is subnetted, 1 subnets
D        192.168.23.0 [90/23796062] via 192.168.12.2, 00:31:01, Serial1/1
R1#
```

```
R1#show ipv route eigrp
IPv6 Routing Table - default - 9 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, l - LISP
EX  ::/0 [170/23796574]
     via FE80::2, Serial1/1
D   2001:DB8:CC1E:3::/64 [90/13556702]
     via FE80::2, Serial1/1
D   2001:DB8:CC1E:4::/64 [90/23796062]
     via FE80::2, Serial1/1
D   2001:DB8:CC1E:5::/64 [90/23796702]
     via FE80::2, Serial1/1
R1#
```

Examine R1's EIGRP topology tables for IPv4 and IPv6 using the show ip eigrp topology and show ipv6 eigrp topology commands:
We can see that the FDs of the EIGRP learned routes displayed in the topology tables are not the same FDs displayed in the routing tables.

```
R1#show ip eigrp topology
EIGRP-IPv4 VR(CCIE) Topology Table for AS(4)/ID(1.1.1.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 192.168.23.0/30, 1 successors, FD is 3045895958
        via 192.168.12.2 (3045895958/1735175958), Serial1/1
P 10.2.2.0/24, 1 successors, FD is 1735257878
        via 192.168.12.2 (1735257878/163840), Serial1/1
P 192.168.12.0/30, 1 successors, FD is 1735175958
        via Connected, Serial1/1
P 10.3.3.0/24, 1 successors, FD is 3045977878
        via 192.168.12.2 (3045977878/1735257878), Serial1/1
P 0.0.0.0/0, 1 successors, FD is 4356615958
        via 192.168.12.2 (4356615958/3045895958), Serial1/1
P 10.1.1.0/24, 1 successors, FD is 163840
```

```
        via Connected, Loopback0

R1#
```

```
R1#show ipv eigrp topology
EIGRP-IPv6 VR(CCIE) Topology Table for AS(6)/ID(1.1.1.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 2001:DB8:CC1E:5::/64, 1 successors, FD is 3045977878
        via FE80::2 (3045977878/1735257878), Serial1/1
P 2001:DB8:CC1E:4::/64, 1 successors, FD is 3045895958
        via FE80::2 (3045895958/1735175958), Serial1/1
P 2001:DB8:CC1E:3::/64, 1 successors, FD is 1735257878
        via FE80::2 (1735257878/163840), Serial1/1
P ::/0, 1 successors, FD is 3045961494
        via FE80::2 (3045961494/1735241494), Serial1/1
P 2001:DB8:CC1E:1::/64, 1 successors, FD is 163840
        via Connected, Loopback0
P 2001:DB8:CC1E:2::/64, 1 successors, FD is 1735175958
        via Connected, Serial1/1

R1#
```

Focus on the EIGRP route to prefixes 10.3.3.0/24 and 2001:DB8:CC1E:5::/64 .

In the routing tables for these prefixes displayed below we can see that the local router's cost to get the prefixes 10.3.3.0/24 and 2001:DB8:CC1E:5::/64 is "23796702":

```
R1#show ip route 10.3.3.0
Routing entry for 10.3.3.0/24
  Known via "eigrp 4", distance 90, metric 23796702, type internal
  Redistributing via eigrp 4
  Last update from 192.168.12.2 on Serial1/1, 01:03:06 ago
  Routing Descriptor Blocks:
  * 192.168.12.2, from 192.168.12.2, 01:03:06 ago, via Serial1/1
      Route metric is 23796702, traffic share count is 1
      Total delay is 40001 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2
R1#
```

```
R1#show ipv6 route 2001:DB8:CC1E:5::/64
Routing entry for 2001:DB8:CC1E:5::/64
  Known via "eigrp 6", distance 90, metric 23796702, type internal
  Route count is 1/1, share count 0
  Routing paths:
    FE80::2, Serial1/1
      Last updated 00:03:25 ago

R1#
```

let's check the topology table for these prefixes:

The feasible Distance is 3045977878 but the RIB is 23796702, what is going on?
In the EIGRP named mode configuration we can see a "Metric rib-scale" command and by default it is set to 128 as shown by the show ip protocols and show ipv6 protocols:

```
R1#show ip protocols | s Metric rib-scale
    Metric rib-scale 128
R1#
R1#show ipv6 protocols | s Metric rib-scale
  Metric rib-scale 128
R1#
```

Therefore, if the Feasible Distance is divided by the RIB-Scale value we should see what is entered in the routing table, let's verify:

3045977878 / 128 = 23796702

The reason is that the variable in the routing table that allows the composite metric to be placed is a 32 bit value and therefore, larger numbers can't be installed, therefore, the RIB-Scale command can reduce this number so it can fit in.

```
R1#show ip eigrp topology 10.3.3.0/24
EIGRP-IPv4 VR(CCIE) Topology Entry for AS(4)/ID(1.1.1.1) for 10.3.3.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 3045977878, RIB
is 23796702
  Descriptor Blocks:
  192.168.12.2 (Serial1/1), from 192.168.12.2, Send flag is 0x0
      Composite metric is (3045977878/1735257878), route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 40001250000 picoseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
        Originating router is 3.3.3.3
R1#
```

```
R1#show ipv6 eigrp topology 2001:DB8:CC1E:5::/64
EIGRP-IPv6 VR(CCIE) Topology Entry for AS(6)/ID(1.1.1.1) for 2001:DB8:CC1E:5::/64
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 3045977878, RIB
is 23796702
  Descriptor Blocks:
  FE80::2 (Serial1/1), from FE80::2, Send flag is 0x0
      Composite metric is (3045977878/1735257878), route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 40001250000 picoseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
        Originating router is 3.3.3.3
R1#
```

Finally verify the connectivity:

```
R1#ping 192.168.34.1 sou lo0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.34.1, timeout is 2 seconds:
Packet sent with a source address of 10.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 176/197/248 ms
R1#
```

```
R1#ping 2001:DB8:CCDA:34::1 sou lo0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:CCDA:34::1, timeout is 2 seconds:
Packet sent with a source address of 2001:DB8:CC1E:1::1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 116/135/160 ms
R1#
```

Lab 9: EIGRP named mode over IPv6

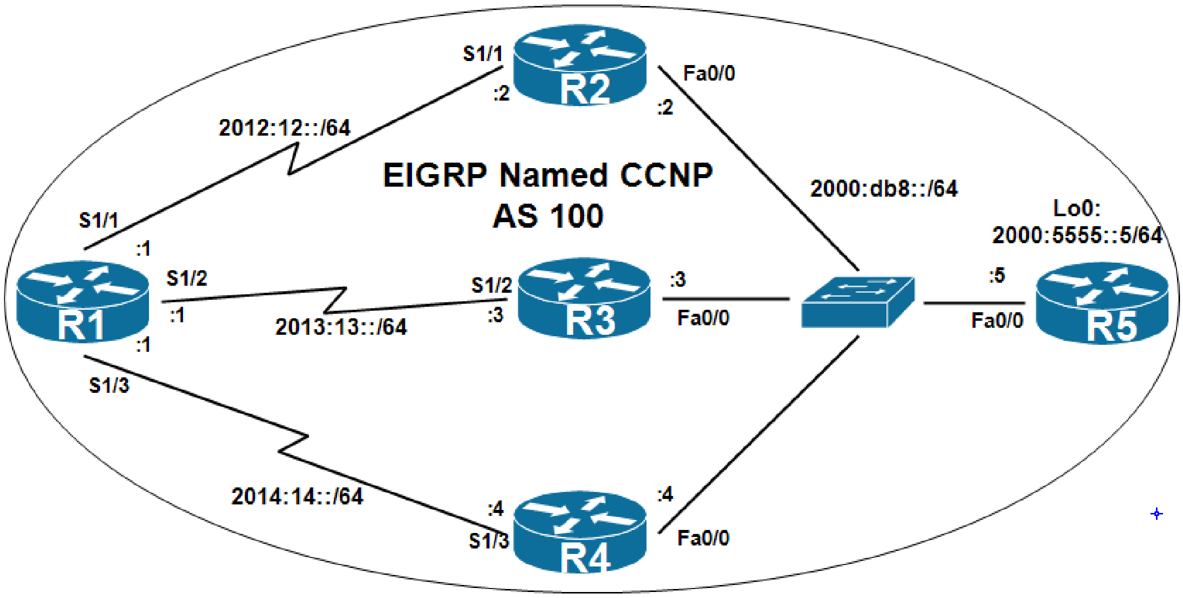

Basic configuration of all routers:

R1:

```
interface Serial1/1
 ipv6 address 2012:12::1/64
 no shutdown
!
interface Serial1/2
 ipv6 address 2013:13::1/64
 no shutdown
!
interface Serial1/3
 ipv6 address 2014:14::1/64
 no shutdown
```

R2:

```
interface FastEthernet0/0
 ipv6 address 2000:db8::2/64
 no shutdown
!
interface Serial1/1
 ipv6 address 2012:12::2/64
 no shutdown
```

R3:

```
interface FastEthernet0/0
 ipv6 address 2000:db8::3/64
 no shutdown
!
interface Serial1/2
 ipv6 address 2013:13::3/64
 no shutdown
```

R4:

```
interface FastEthernet0/0
```

```
 ipv6 address 2000:db8::4/64
 no shutdown
!
interface Serial1/3
 ipv6 address 2014:14::4/64
 no shutdown
```

R5:

```
interface loopback0
 ipv6 address 2000:5555::5/64
!
interface FastEthernet0/0
 ipv6 address 2000:db8::5/64
 no shutdown
```

To configure EIGRP named mode we should use the router eigrp command followed by an arbitrary name (in this case CCNP)

One the router eigrp CCNP command is entered, we need to specify which address-family we should use, either IPv4 or IPv6 and followed by the autonomous-systeme number 100:

In classic mode with IPv6, EIGRP is activated directly in the interfaces with all specific EIGRP commands such as bandwitdh and delay command, the hello and the hold timers, split-horizon etc...

In named mode, all these commands except the bandwitdh and delay, are applied under the EIGRP process by entering the af-interface command followed by the specific interface number. With af-interface command we will enter in the Per-AF-interface section.

By default each interface on which IPv6 is enabled , the EIGRP address-family automatically activates EIGRP on these interfaces.
Now if we need to disable EIGRP on all interfaces, we should configure the af-interface default command and enter the shutdown command.
As a result we need to select interfaces individually to activate EIGRP by entering in their specific af-interface sections followed by the no shutdown command.
This method is helpful when we need to activate EIGRP in some interfaces only.

R1:

```
router eigrp CCNP
 !
 address-family ipv6 unicast autonomous-system 100
  !
  af-interface default
   shutdown
  exit-af-interface
  !
  af-interface Serial1/1
   no shutdown
  exit-af-interface
  !
  af-interface Serial1/2
   no shutdown
  exit-af-interface
  !
  af-interface Serial1/3
   no shutdown
  exit-af-interface
```

R2:

```
router eigrp CCNP
 !
 address-family ipv6 unicast autonomous-system 100
  !
  af-interface default
   shutdown
  exit-af-interface
  !
  af-interface Serial1/1
   no shutdown
  exit-af-interface
  !
  af-interface FastEthernet0/0
   no shutdown
  exit-af-interface
```

R3:

```
router eigrp CCNP
 !
 address-family ipv6 unicast autonomous-system 100
  !
  af-interface default
   shutdown
  exit-af-interface
  !
  af-interface Serial1/2
   no shutdown
  exit-af-interface
  !
  af-interface FastEthernet0/0
   no shutdown
  exit-af-interface
```

R4:

```
router eigrp CCNP
 !
 address-family ipv6 unicast autonomous-system 100
  !
  af-interface default
   shutdown
  exit-af-interface
  !
  af-interface Serial1/3
   no shutdown
  exit-af-interface
  !
  af-interface FastEthernet0/0
   no shutdown
  exit-af-interface
```

R5:

```
router eigrp CCNP
 !
 address-family ipv6 unicast autonomous-system 100
  !
  af-interface default
   shutdown
  exit-af-interface
```

```
 !
 af-interface FastEthernet0/0
  no shutdown
 exit-af-interface
 !
 af-interface Loopback0
  no shutdown
 exit-af-interface
```

Now to configure the router-ID on all routers, we should go to the address-family section using still the address-family ipv6 unicast autonomous-system 100 command:

R1:

```
router eigrp CCNP
 !
 address-family ipv6 unicast autonomous-system 100
  !
  eigrp router-id 1.1.1.1
```

R2:

```
router eigrp CCNP
 !
 address-family ipv6 unicast autonomous-system 100
  !
  eigrp router-id 2.2.2.2
```

R3:

```
router eigrp CCNP
 !
 address-family ipv6 unicast autonomous-system 100
  !
  eigrp router-id 3.3.3.3
```

R4:

```
router eigrp CCNP
 !
 address-family ipv6 unicast autonomous-system 100
  !
  eigrp router-id 4.4.4.4
```

R5:

```
router eigrp CCNP
 !
 address-family ipv6 unicast autonomous-system 100
  !
  eigrp router-id 5.5.5.5
```

Let's verify the neighbor relationships:

```
R1#show eigrp address-family ipv6 neighbors
EIGRP-IPv6 VR(CCNP) Address-Family Neighbors for AS(100)
H   Address                  Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                    (sec)         (ms)       Cnt
Num
2   Link-local address:      Se1/3                    13 00:00:34   86    516  0   11
    FE80::C803:DFF:FE3C:8
1   Link-local address:      Se1/2                    10 00:00:46   93    558  0   14
```

```
    FE80::C802:DFF:FE3C:8
0   Link-local address:     Se1/1                   13 00:00:56   85    510  0  17
    FE80::C801:17FF:FE40:8
R1#
```

```
R2#show eigrp address-family ipv6 neighbors
EIGRP-IPv6 VR(CCNP) Address-Family Neighbors for AS(100)
H   Address                 Interface             Hold Uptime    SRTT    RTO  Q
Seq
                                                  (sec)          (ms)        Cnt
Num
3   Link-local address:     Fa0/0                   10 00:00:53   44    264  0  8
    FE80::C804:26FF:FED4:8
2   Link-local address:     Fa0/0                   12 00:01:03   96    576  0  12
    FE80::C803:DFF:FE3C:8
1   Link-local address:     Fa0/0                   12 00:01:15   91    546  0  13
    FE80::C802:DFF:FE3C:8
0   Link-local address:     Se1/1                   13 00:01:25   76    456  0  20
    FE80::C800:17FF:FE40:8
R2#
```

```
R3#show eigrp address-family ipv6 neighbors
EIGRP-IPv6 VR(CCNP) Address-Family Neighbors for AS(100)
H   Address                 Interface             Hold Uptime    SRTT    RTO  Q
Seq
                                                  (sec)          (ms)        Cnt
Num
3   Link-local address:     Fa0/0                   11 00:01:10  141    846  0  8
    FE80::C804:26FF:FED4:8
2   Link-local address:     Fa0/0                   14 00:01:20   78    468  0  12
    FE80::C803:DFF:FE3C:8
1   Link-local address:     Se1/2                   10 00:01:32  124    744  0  21
    FE80::C800:17FF:FE40:8
0   Link-local address:     Fa0/0                   11 00:01:32   82    492  0  16
    FE80::C801:17FF:FE40:8
R3#
```

```
R4#show eigrp address-family ipv6 neighbors
EIGRP-IPv6 VR(CCNP) Address-Family Neighbors for AS(100)
H   Address                 Interface             Hold Uptime    SRTT    RTO  Q
Seq
                                                  (sec)          (ms)        Cnt
Num
3   Link-local address:     Fa0/0                   14 00:01:30   79    474  0  8
    FE80::C804:26FF:FED4:8
2   Link-local address:     Fa0/0                   14 00:01:40   78    468  0  16
    FE80::C801:17FF:FE40:8
1   Link-local address:     Se1/3                   10 00:01:40   78    468  0  19
    FE80::C800:17FF:FE40:8
0   Link-local address:     Fa0/0                   11 00:01:40   72    432  0  13
    FE80::C802:DFF:FE3C:8
R4#
```

```
R5#show eigrp address-family ipv6 neighbors
EIGRP-IPv6 VR(CCNP) Address-Family Neighbors for AS(100)
H   Address                 Interface             Hold Uptime    SRTT    RTO  Q
Seq
```

```
                                                        (sec)         (ms)        Cnt
Num
0   Link-local address:     Fa0/0                        11 00:00:47    93    558  0   42
    FE80::C802:DFF:FE3C:8
2   Link-local address:     Fa0/0                        14 00:01:08   399   2394  0   38
    FE80::C801:17FF:FE40:8
1   Link-local address:     Fa0/0                        12 00:01:08    75    450  0   29
    FE80::C803:DFF:FE3C:8
R5#
```

Now let's verify and focus the routing table of R1:
R1 receives an EIGRP route to 2000:5555::/64 prefix with the metric 13607262 and because we are the default parameters values of bandwitdh and delay , R1 performs a

Load-Balancing through R2, R3 and R4:

```
R1#show ipv route eigrp
IPv6 Routing Table - default - 9 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, l - LISP
D   2000:DB8::/64 [90/13607262]
     via FE80::C801:17FF:FE40:8, Serial1/1
     via FE80::C803:DFF:FE3C:8, Serial1/3
     via FE80::C802:DFF:FE3C:8, Serial1/2
D   2000:5555::/64 [90/13607902]
     via FE80::C803:DFF:FE3C:8, Serial1/3
     via FE80::C801:17FF:FE40:8, Serial1/1
     via FE80::C802:DFF:FE3C:8, Serial1/2
R1#
```

```
R1#show ipv route 2000:5555::/64
Routing entry for 2000:5555::/64
  Known via "eigrp 100", distance 90, metric 13607902, type internal
  Route count is 3/3, share count 0
  Routing paths:
    FE80::C803:DFF:FE3C:8, Serial1/3
      Last updated 00:12:14 ago
    FE80::C801:17FF:FE40:8, Serial1/1
      Last updated 00:12:12 ago
    FE80::C802:DFF:FE3C:8, Serial1/2
      Last updated 00:11:53 ago

R1#
```

Now let's check the topology table of R1:
We can see that the topology table confirms that R1 has three Successors to 2000:5555::/64 prefix.
But the Feasible Distance FD (in other the total metric) is 1741811478 which is different than the metric listed in the routing table shown above 13607262:

```
R1#show eigrp address-family ipv6 topology
EIGRP-IPv6 VR(CCNP) Topology Table for AS(100)/ID(1.1.1.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status
```

```
P 2013:13::/64, 1 successors, FD is 1735175958
        via Connected, Serial1/2
P 2012:12::/64, 1 successors, FD is 1735175958
        via Connected, Serial1/1
P 2000:DB8::/64, 3 successors, FD is 1741729558
        via FE80::C801:17FF:FE40:8 (1741729558/13107200), Serial1/1
        via FE80::C802:DFF:FE3C:8 (1741729558/13107200), Serial1/2
        via FE80::C803:DFF:FE3C:8 (1741729558/13107200), Serial1/3
P 2014:14::/64, 1 successors, FD is 1735175958
        via Connected, Serial1/3
P 2000:5555::/64, 3 successors, FD is 1741811478
        via FE80::C801:17FF:FE40:8 (1741811478/13189120), Serial1/1
        via FE80::C802:DFF:FE3C:8 (1741811478/13189120), Serial1/2
        via FE80::C803:DFF:FE3C:8 (1741811478/13189120), Serial1/3

R1#
```

The reason is that EIGRP named mode automatically uses wide metric when talking with another EIGRP named mode.
In classis mode it uses the old computation.

In the named mode the new metric is designed to be able to differentiate paths above 10GB. below the changes of the metric computation with the named mode:

- Delay is now measured in picoseconds instead of microseconds with the classic mode.
- Bandwidth's scaling factor is now much larger, the calculation is now 10^7 * 65536 / Min-Bandwidth, rather that the old formula in the classic mode
10^7 * 256 / Min Bandwidth.
- The overall metric is now 64 bit.
- The K6 value has been added "for future use", but Cisco indicated this will be used for accumulated energy or accumulated jitter.

In the named mode with wide metrics, the metric is no longer fits into the RIB or the routing table.

Below the details of the topology tables of R1 for the prefix 2000:5555::/64:
The show eigrp address-family ipv6 topology 2000:5555::/64 command shown more informations about the new metric and the old metric.

The EIGRP topology table indicates the FD 1741811478, the RIB says 13607902.

By default The RIB's metric can't exceed 32-bits, therefore with EIGRP named mode which uses a metric 64 bits, more granular metrics won't fit into the RIB. So all

metrics, regardless if the value will be displayed with 32-bits using the old metric calculation in the classic mode,

The FD of the topology table which uses 64 bits is divided by the rib-scale value. Since by default the rib-scale is 128, therefore R1 divides the FD 1741811478 by 128 and finds the metric 13607902 listed in the routing table:

1741811478 / 128 = 13607902

```
R1#show eigrp address-family ipv6 topology 2000:5555::/64
EIGRP-IPv6 VR(CCNP) Topology Entry for AS(100)/ID(1.1.1.1) for 2000:5555::/64
  State is Passive, Query origin flag is 1, 3 Successor(s), FD is 1741811478, RIB
is 13607902
  Descriptor Blocks:
  FE80::C801:17FF:FE40:8 (Serial1/1), from FE80::C801:17FF:FE40:8, Send flag is
0x0
```

```
      Composite metric is (1741811478/13189120), route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 20101250000 picoseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
        Originating router is 5.5.5.5
  FE80::C802:DFF:FE3C:8 (Serial1/2), from FE80::C802:DFF:FE3C:8, Send flag is 0x0
      Composite metric is (1741811478/13189120), route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 20101250000 picoseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
        Originating router is 5.5.5.5
  FE80::C803:DFF:FE3C:8 (Serial1/3), from FE80::C803:DFF:FE3C:8, Send flag is 0x0
      Composite metric is (1741811478/13189120), route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 20101250000 picoseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
        Originating router is 5.5.5.5
R1#
```

We can change the value of the rib-scale from 1 to 255:

```
R1(config)# router eigrp CCNP
R1(config-router)#address-family ipv6 unicast autonomous-system 100
R1(config-router-af)#metric rib-scale 200
```

Now the FD listed in the topology table is still 1741811478, but the metric listed in the RIB (Routing Table) is now changed 8709057 and calculated as follow:

1741811478 / 200 = 8709057

```
R1#show eigrp address-family ipv6 topology 2000:5555::/64
EIGRP-IPv6 VR(CCNP) Topology Entry for AS(100)/ID(1.1.1.1) for 2000:5555::/64
  State is Passive, Query origin flag is 1, 3 Successor(s), FD is 1741811478, RIB
is 8709057
  Descriptor Blocks:
  FE80::C802:BFF:FE18:8 (Serial1/1), from FE80::C802:BFF:FE18:8, Send flag is 0x0
      Composite metric is (1741811478/13189120), route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 20101250000 picoseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
        Originating router is 5.5.5.5
  FE80::C803:16FF:FE50:8 (Serial1/3), from FE80::C803:16FF:FE50:8, Send flag is
0x0
```

```
      Composite metric is (1741811478/13189120), route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 20101250000 picoseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
        Originating router is 5.5.5.5
  FE80::C804:16FF:FE50:8 (Serial1/2), from FE80::C804:16FF:FE50:8, Send flag is
0x0
      Composite metric is (1741811478/13189120), route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 20101250000 picoseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
        Originating router is 5.5.5.5
R1#
```

```
R1#show ipv route  2000:5555::/64
Routing entry for 2000:5555::/64
  Known via "eigrp 100", distance 90, metric 8709057, type internal
  Route count is 3/3, share count 0
  Routing paths:
    FE80::C802:BFF:FE18:8, Serial1/1
      Last updated 00:50:32 ago
    FE80::C804:16FF:FE50:8, Serial1/2
      Last updated 00:50:31 ago
    FE80::C803:16FF:FE50:8, Serial1/3
      Last updated 00:50:31 ago

R1#
```

Lab 10: EIGRP Named Mode Authentication

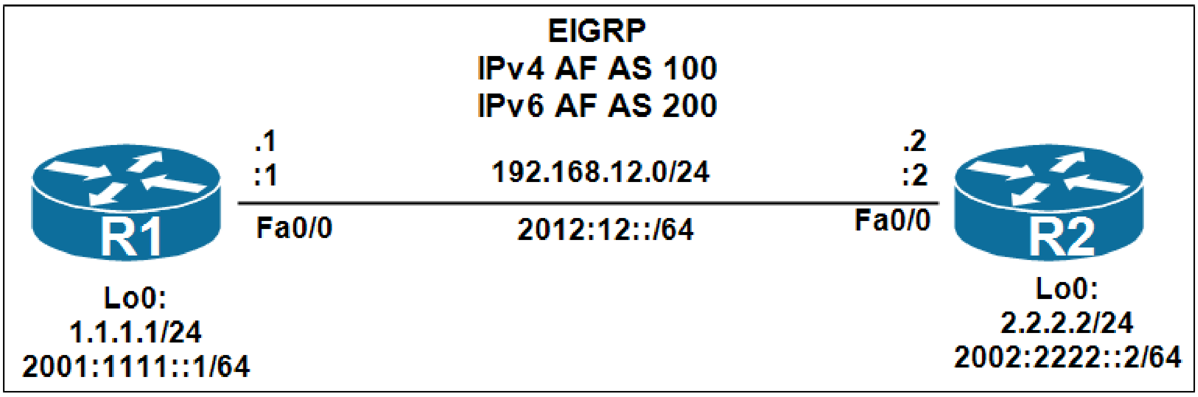

R1 and R2 are using EIGRP Named Mode for both IPv4 and IPv6 Address Families using the router eigrp CCIE command.
IPv4 EIGRP address family for AS number 100 is enabled.
IPv6 EIGRP address family for AS number 200 is enabled.

For IPv4, the network commands are used to enable EIGRP on selected interfaces in the AF section. Unlike with IPv6, when issuing the address-family ipv6 unicast

autonomous-system 200 command, EIGRP is enabled on all interfaces.

Below the output of the show run | s CCIE command to see the details of the commands on R1 and R2:

```
R1#show run | s CCIE
router eigrp CCIE
 !
 address-family ipv4 unicast autonomous-system 100
  !
  topology base
  exit-af-topology
  network 1.1.1.1 0.0.0.0
  network 192.168.12.0
 exit-address-family
 !
 address-family ipv6 unicast autonomous-system 200
  !
  topology base
  exit-af-topology
 exit-address-family
R1#
```

```
R2#show run | s CCIE
router eigrp CCIE
 !
 address-family ipv4 unicast autonomous-system 100
  !
  topology base
  exit-af-topology
  network 2.2.2.2 0.0.0.0
  network 192.168.12.0
 exit-address-family
```

```
 !
 address-family ipv6 unicast autonomous-system 200
  !
  topology base
  exit-af-topology
 exit-address-family
R2#
```

To verify the neighbor relationship, use the show eigrp address-family ipv4 neighbors and show eigrp address-family ipv6 neighbors commands:

```
R1#show eigrp address-family ipv4 neighbors
EIGRP-IPv4 VR(CCIE) Address-Family Neighbors for AS(100)
H   Address                 Interface              Hold Uptime    SRTT    RTO  Q
Seq
                                                   (sec)          (ms)        Cnt
Num
0   192.168.12.2            Fa0/0                    14 00:06:51   117    702  0   4
R1#
```

```
R1#show eigrp address-family ipv6 neighbors
EIGRP-IPv6 VR(CCIE) Address-Family Neighbors for AS(200)
H   Address                 Interface              Hold Uptime    SRTT    RTO  Q
Seq
                                                   (sec)          (ms)        Cnt
Num
0   Link-local address:     Fa0/0                    12 00:10:09   122    732  0   3
    FE80::C802:17FF:FE34:0
R1#
```

Let's verify the IPv4 and IPv6 routing tables of R1 and R2 to confirm that they are receiving the IPv4 networks and IPv6 prefixes of the loopback interfaces:

```
R1#show ip route eigr | beg Gate
Gateway of last resort is not set

      2.0.0.0/24 is subnetted, 1 subnets
D        2.2.2.0 [90/103040] via 192.168.12.2, 00:03:46, FastEthernet0/0
R1#
```

```
R1#show ipv6 route eig
IPv6 Routing Table - default - 6 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, l - LISP
D   2002:2222::/64 [90/103040]
     via FE80::C802:17FF:FE34:0, FastEthernet0/0
R1#
```

```
R2#show ip route eigr | beg Gate
Gateway of last resort is not set

      1.0.0.0/24 is subnetted, 1 subnets
D        1.1.1.0 [90/103040] via 192.168.12.1, 00:05:20, FastEthernet0/0
R2#
```

```
R2#show ipv6 route eigr
IPv6 Routing Table - default - 6 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, l - LISP
D   2001:1111::/64 [90/103040]
     via FE80::C801:18FF:FEC0:0, FastEthernet0/0
R2#
```

Now let's see how to configure EIGRP authentication in Named Mode.

First we define a Key chain KEY-CCIE with a single key:

```
R1(config)#key chain KEY-CCIE
R1(config-keychain)#key 1
R1(config-keychain-key)#key-string CISCO-CCIE
```

In EIGRP Named Mode we have the ability to configure the MD5 or hmac-sha-256 authentication, another option is to enable authentication on all EIGRP enabled interfaces using the KEY-CCIE key chain in the af-interface default section as follow:

```
R1(config)# router eigrp CCIE
R1(config-router)# address-family ipv4 autonomous-system 100
R1(config-router-af)# af-interface default
R1(config-router-af-interface)# authentication key-chain KEY-CCIE
R1(config-router-af-interface)#authentication mode ?
  hmac-sha-256  HMAC-SHA-256 Authentication
  md5           Keyed message digest

R1(config-router-af-interface)#authentication mode md5
```

Now if we specify the hmac-sha-256 authentication on Fa0/0 of R1 and R2, as a result the authentication type is overridden to SHA-256, using the key configured in the

KEY-CCIE key chain. But, the authentication mode hmac-sha-256 command tell us that a password should be specified even if a key chain was configured above.
And if we use another password "PASSWORD-CCIE" as we have done below, both will be used for authentication (CISCO-CCIE configured in the KEY-CCIE key chain and per-interface configured password "PASSWORD-CCIE”).

```
R1(config)#router eigrp CCIE
R1(config-router)#address-family ipv4 autonomous-system 100
R1(config-router-af)#af-interface fa0/0
R1(config-router-af-interface)#authentication ?
  key-chain  key-chain
  mode       authentication mode

R1(config-router-af-interface)#authentication key-chain KEY-CCIE
R1(config-router-af-interface)#authentication mode ?
  hmac-sha-256  HMAC-SHA-256 Authentication
  md5           Keyed message digest

R1(config-router-af-interface)#authentication mode hmac-sha-256 ?
  <0-7>  Encryption type (0 to disable encryption, 7 for proprietary)
  LINE   password
```

```
R1(config-router-af-interface)#authentication mode hmac-sha-256 PASSWORD-CCIE
R1(config-router-af-interface)#
```

The same steps are used for IPv6 AF using the same key-chain and a different password for the fa0/0 interface:

```
R1(config)#router eigrp CCIE
R1(config-router)#address-family ipv6 autonomous-system 200
R1(config-router-af)#af-interface fa0/0
R1(config-router-af-interface)#authentication key-chain KEY-CCIE
R1(config-router-af-interface)#authentication mode hmac-sha-256 password-ccie
```

The same configuration is applied on R2:

```
R2(config)#key chain KEY-CCIE
R2(config-keychain)#key 1
R2(config-keychain-key)#key-string CISCO-CCIE

R2(config)#router eigrp CCIE
R2(config-router)#address-family ipv4 autonomous-system 100
R2(config-router-af)#af-interface fa0/0
R2(config-router-af-interface)#authentication key-chain KEY-CCIE
R2(config-router-af-interface)#authentication mode hmac-sha-256 PASSWORD-CCIE

R2(config)#router eigrp CCIE
R2(config-router)#address-family ipv6 autonomous-system 200
R2(config-router-af)#af-interface fa0/0
R2(config-router-af-interface)#authentication key-chain KEY-CCIE
R2(config-router-af-interface)#authentication mode hmac-sha-256 password-ccie
```

Let's verify all the commands using the show run | s CCIE command:

```
R1#show run | s CCIE
key chain KEY-CCIE
 key 1
   key-string CISCO-CCIE
router eigrp CCIE
 !
 address-family ipv4 unicast autonomous-system 100
  !
  af-interface FastEthernet0/0
   authentication mode hmac-sha-256 PASSWORD-CCIE
   authentication key-chain KEY-CCIE
  exit-af-interface
  !
  topology base
  exit-af-topology
  network 1.1.1.1 0.0.0.0
  network 192.168.12.0
 exit-address-family
 !
 address-family ipv6 unicast autonomous-system 200
  !
  af-interface FastEthernet0/0
   authentication mode hmac-sha-256 password-ccie
   authentication key-chain KEY-CCIE
  exit-af-interface
  !
```

```
  topology base
  exit-af-topology
 exit-address-family
R1#
```

```
R2#show run | s CCIE
key chain KEY-CCIE
 key 1
   key-string CISCO-CCIE
router eigrp CCIE
 !
 address-family ipv4 unicast autonomous-system 100
  !
  af-interface FastEthernet0/0
   authentication mode hmac-sha-256 PASSWORD-CCIE
   authentication key-chain KEY-CCIE
  exit-af-interface
  !
  topology base
  exit-af-topology
  network 2.2.2.2 0.0.0.0
  network 192.168.12.0
 exit-address-family
 !
 address-family ipv6 unicast autonomous-system 200
  !
  af-interface FastEthernet0/0
   authentication mode hmac-sha-256 password-ccie
   authentication key-chain KEY-CCIE
  exit-af-interface
  !
  topology base
  exit-af-topology
 exit-address-family
R2#
```

The neighbor relationships confirm that the authentication is successfull for both IPv4 and IPv6 AFs:

```
R1#show eigrp address-family ipv4 neighbors
EIGRP-IPv4 VR(CCIE) Address-Family Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                   (sec)         (ms)       Cnt
Num
0   192.168.12.2            Fa0/0                    11 00:04:11  117   702  0  7
R1#
```

```
R1#show eigrp address-family ipv6 neighbors
EIGRP-IPv6 VR(CCIE) Address-Family Neighbors for AS(200)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                   (sec)         (ms)       Cnt
Num
0   Link-local address:     Fa0/0                    11 00:03:50  109   654  0  6
    FE80::C802:17FF:FE34:0
R1#
```

We can verify the authentication for fa0/0 interface by using the show eigrp address-family ipv4 int detail fa0/0 command:

```
R1#show eigrp address-family ipv4 int detail fa0/0
EIGRP-IPv4 VR(CCIE) Address-Family Interfaces for AS(100)
                              Xmit Queue   PeerQ        Mean   Pacing Time
Multicast    Pending
Interface              Peers  Un/Reliable  Un/Reliable  SRTT   Un/Reliable   Flow
Timer   Routes
Fa0/0                    1        0/0       0/0          117       0/1           496
0
  Hello-interval is 5, Hold-time is 15
  Split-horizon is enabled
  Next xmit serial <none>
  Packetized sent/expedited: 4/0
  Hello's sent/expedited: 380/3
  Un/reliable mcasts: 0/4  Un/reliable ucasts: 5/4
  Mcast exceptions: 0  CR packets: 0  ACKs suppressed: 0
  Retransmissions sent: 2  Out-of-sequence rcvd: 0
  Topology-ids on interface - 0
  Authentication mode is HMAC-SHA-256, key-chain is "KEY-CCIE"
R1#
```

```
R1#show eigrp address-family ipv6 int detail fa0/0
EIGRP-IPv6 VR(CCIE) Address-Family Interfaces for AS(200)
                              Xmit Queue   PeerQ        Mean   Pacing Time
Multicast    Pending
Interface              Peers  Un/Reliable  Un/Reliable  SRTT   Un/Reliable   Flow
Timer   Routes
Fa0/0                    1        0/0       0/0          109       0/1           448
0
  Hello-interval is 5, Hold-time is 15
  Split-horizon is enabled
  Next xmit serial <none>
  Packetized sent/expedited: 3/0
  Hello's sent/expedited: 443/3
  Un/reliable mcasts: 0/4  Un/reliable ucasts: 5/4
  Mcast exceptions: 0  CR packets: 0  ACKs suppressed: 0
  Retransmissions sent: 2  Out-of-sequence rcvd: 1
  Topology-ids on interface - 0
  Authentication mode is HMAC-SHA-256, key-chain is "KEY-CCIE"
R1#
```

Verify the test with the ping command:

```
R1#ping 2.2.2.2 sou 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 = 64/72/76 ms
R1#
R1#ping 2002:2222::2 sou 2001:1111::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2002:2222::2, timeout is 2 seconds:
Packet sent with a source address of 2001:1111::1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/52/92 ms
R1#
```

Lab 11: EIGRP Stub Routing with Named Mode

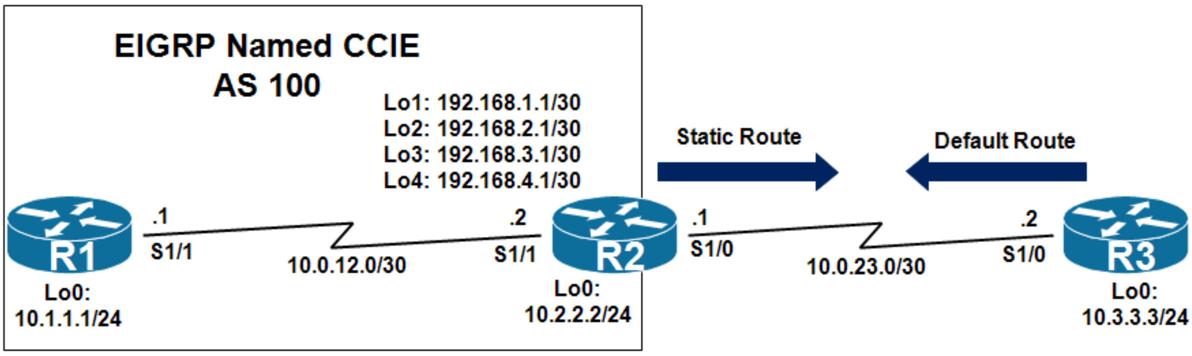

Basic configuration of all routers:

R1:

```
interface Loopback0
 ip address 10.1.1.1 255.255.255.0
 no shutdown
!
interface Serial1/1
 ip address 10.0.12.1 255.255.255.252
 no shutdown
```

R2:

```
interface Loopback1
 ip address 192.168.1.1 255.255.255.252
!
interface Loopback2
 ip address 192.168.2.1 255.255.255.252
!
interface Loopback3
 ip address 192.168.3.1 255.255.255.252
!
interface Loopback4
 ip address 192.168.4.1 255.255.255.252
!
interface Loopback0
 ip address 10.2.2.2 255.255.255.0
 no shutdown
!
interface Serial1/1
 ip address 10.0.12.2 255.255.255.0
 no shutdown
!
interface Serial1/0
 ip address 10.0.23.1 255.255.255.252
 no shutdown
```

R3:

```
interface Serial1/0
 ip address 10.0.23.2 255.255.255.252
 no shutdown
!
interface Loopback0
```

```
 ip address 10.3.3.3 255.255.255.0
 no shutdown
```

Enable EIGRP between R1 and R2 and advertise connected prefixes using EIGRP Named Mode:

```
R1(config)#router eigrp CCIE
R1(config-router)#address-family ipv4 autonomous-system 100
R1(config-router-af)#network 10.0.0.0
```

```
R2(config)#router eigrp CCIE
R2(config-router)#address-family ipv4 autonomous-system 100
R2(config-router-af)#network 10.0.0.0
R2(config-router-af)#network 192.168.0.0 0.0.255.255
```

We configure manual summarization of R2's loopback interfaces Lo1, Lo2, Lo3 and Lo4:

```
R2(config)#router eig CCIE
R2(config-router)#address-family ipv4 autonomous-system 100
R2(config-router-af)#af-interface s1/1
R2(config-router-af-interface)#address-family ipv4 autonomous-system 100
```

We configure a static route on R2 to R3'loopback0 and a default static route on R3 forwarding all traffic to R2:

```
R2(config)#ip route 10.3.3.0 255.255.255.0 10.0.23.2
```

```
R3(config)#ip route 0.0.0.0 0.0.0.0 10.0.23.1
```

Verify connectivity:

```
R2#ping 10.3.3.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/69/76 ms
R2#
```

```
R3#ping 10.2.2.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 80/194/240 ms
R3#
```

Verify the neighbor relationship:

```
R2#show ip eigrp neighbors
EIGRP-IPv4 VR(CCIE) Address-Family Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  
Seq
                                                   (sec)         (ms)       Cnt 
Num
0   10.0.12.1               Se1/1                    12 00:00:07  122    732  0  3
R2#
```

Verify that R1 is receiving a summary route for R2's loopback prefixes:

```
R1#show ip route eigrp | beg Gate
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 6 subnets, 3 masks
D        10.0.23.0/30 [90/23796062] via 10.0.12.2, 00:03:19, Serial1/1
D        10.2.2.0/24 [90/13556702] via 10.0.12.2, 00:03:19, Serial1/1
D     192.168.0.0/21 [90/13556702] via 10.0.12.2, 00:03:19, Serial1/1
R1#
```

Configure R2 as a stub router using the default eigrp stub command.

```
R2(config-router-af)#router eig CCIE
R2(config-router)#address-family ipv4 autonomous-system 100
R2(config-router-af)#eigrp stub
```

By default, the connected and summary options are used as shown by the show running-config | section eigrp command on R2:

```
R2#show running-config | section eigrp
router eigrp CCIE
 !
 address-family ipv4 unicast autonomous-system 100
  !
  af-interface Serial1/1
   summary-address 192.168.0.0 255.255.248.0
  exit-af-interface
  !
  topology base
  exit-af-topology
  network 10.0.0.0
  network 192.168.0.0 0.0.255.255
  eigrp stub connected summary
 exit-address-family
R2#
```

Let's verify the routing table of R1, notice that R1 shows EIGRP routes for R2's connected prefixes and R2's 192.16.0.0/21 summary route:

```
R1#show ip route eigrp | beg Gate
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 6 subnets, 3 masks
D        10.0.23.0/30 [90/23796062] via 10.0.12.2, 00:03:19, Serial1/1
D        10.2.2.0/24 [90/13556702] via 10.0.12.2, 00:03:19, Serial1/1
D     192.168.0.0/21 [90/13556702] via 10.0.12.2, 00:03:19, Serial1/1
R1#
```

R1 sees R2 as a stub router as shown by the show ip eigrp neighbors detail command:

```
R1#show ip eigrp neighbors detail
EIGRP-IPv4 VR(CCIE) Address-Family Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                   (sec)         (ms)       Cnt
Num
0   10.0.12.2               Se1/1                    10 00:04:28  114   684  0  8
```

```
   Version 11.0/2.0, Retrans: 0, Retries: 0, Prefixes: 3
   Topology-ids from peer - 0
   Stub Peer Advertising (CONNECTED SUMMARY ) Routes
   Suppressing queries
R1#
```

Let's redistribute the static route ip route 10.3.3.0 255.255.255.0 10.0.23.2 using the redistribute static command and configure R2's stub routing to also include its static route in its EIGRP update to R1 using the eigrp stub static command:

```
R2(config)#router ei CCIE
R2(config-router)#address-family ipv4 autonomous-system 100
R2(config-router-af)#topology base
R2(config-router-af-topology)#redistribute static
R2(config-router-af-topology)#exit
R2(config-router-af)#eigrp stub static
```

Let's verify the routing table of R1:
R1 has only R2's static route to R3's Loopback0:

```
R1#show ip route eigrp | beg Gate
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 5 subnets, 3 masks
D EX     10.3.3.0/24 [170/23796062] via 10.0.12.2, 00:00:40, Serial1/1
R1#
```

When using the eigrp stub static command, R2 removed the connected and summary options as shown by the show running-config | section eigrp command:

```
R2#show running-config | section eigrp
router eigrp CCIE
 !
 address-family ipv4 unicast autonomous-system 100
  !
  af-interface Serial1/1
   summary-address 192.168.0.0 255.255.248.0
  exit-af-interface
  !
  topology base
   redistribute static
  exit-af-topology
  network 10.0.0.0
  network 192.168.0.0 0.0.255.255
  eigrp stub static
 exit-address-family
R2#
```

We can confirm with the show ip eigrp neighbors command:

```
R1#show ip eigrp neighbors detail
EIGRP-IPv4 VR(CCIE) Address-Family Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                   (sec)         (ms)       Cnt
Num
0   10.0.12.2               Se1/1                    12 00:01:30   99    594  0  12
   Version 11.0/2.0, Retrans: 0, Retries: 0, Prefixes: 1
```

```
   Topology-ids from peer - 0
   Stub Peer Advertising (STATIC ) Routes
   Suppressing queries
R1#
```

Therefore we need to configure R2 EIGRP stub routing to include the connected, summary and static options using the eigrp stub connected summary static command:

```
R2(config)#router ei CCIE
R2(config-router)#address-family ipv4 autonomous-system 100
R2(config-router-af)#eigrp stub connected summary static
```

Now R2 is sending its connected, summarized and static routes to R1:

```
R1#show ip route eigrp | beg Gate
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 7 subnets, 3 masks
D        10.0.23.0/30 [90/23796062] via 10.0.12.2, 00:00:19, Serial1/1
D        10.2.2.0/24 [90/13556702] via 10.0.12.2, 00:00:19, Serial1/1
D EX     10.3.3.0/24 [170/23796062] via 10.0.12.2, 00:00:19, Serial1/1
D     192.168.0.0/21 [90/13556702] via 10.0.12.2, 00:00:19, Serial1/1
R1#
```

Let's verify that the R2 EIGRP stub routing is modified using the show ip eigrp neighbors detail command:

```
R1#show ip eigrp neighbors detail
EIGRP-IPv4 VR(CCIE) Address-Family Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                   (sec)         (ms)       Cnt
Num
0   10.0.12.2               Se1/1                    13 00:01:24 1326   5000  0  15
   Version 11.0/2.0, Retrans: 0, Retries: 0, Prefixes: 4
   Topology-ids from peer - 0
   Stub Peer Advertising (CONNECTED STATIC SUMMARY ) Routes
   Suppressing queries
R1#
```

```
R2#show running-config | section eigrp
router eigrp CCIE
 !
 address-family ipv4 unicast autonomous-system 100
  !
  af-interface Serial1/1
   summary-address 192.168.0.0 255.255.248.0
  exit-af-interface
  !
  topology base
   redistribute static
  exit-af-topology
  network 10.0.0.0
  network 192.168.0.0 0.0.255.255
  eigrp stub connected static summary
 exit-address-family
R2#
```

Verify the connectivity after including the connected and summary keywords:

```
R1#ping 10.3.3.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 88/104/128 ms
R1#
```

```
R3#ping 10.1.1.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 132/184/212 ms
R3#
```

Now let's verify another option which is receive-only.
The receive-only option prevents the stub router from sending any of its routes with any other router in the EIGRP AS. This option does not permit any other option to be included.

```
R2(config)#router eig CCIE
R2(config-router)#address-family ipv4 autonomous-system 100
R2(config-router-af)#eigrp stub receive-only
```

Let's verify the routing table of R1 once again:

R1 does not receive any EIGRP routes from R2.

```
R1#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is not set

R1#
```

Now R1 sees R2 as a receive-only stub router using the show ip eigrp neighbor detail command:

```
R1#show ip eigrp neighbor detail
EIGRP-IPv4 VR(CCIE) Address-Family Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                   (sec)         (ms)       Cnt
Num
0   10.0.12.2               Se1/1                    11 00:01:03   87   522  0  17
   Version 11.0/2.0, Retrans: 0, Retries: 0
   Topology-ids from peer - 0
   Receive-Only Peer Advertising (No) Routes
   Suppressing queries
R1#
```

Lab 12: Maximum-prefix limit and Warning-Only Mode

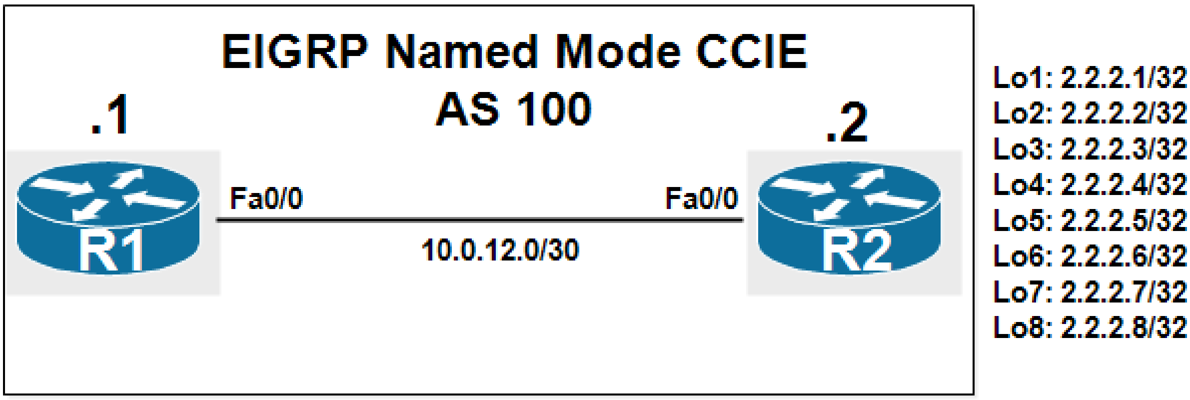

Basic configuration of routers:

```
R1(config)#interface FastEthernet0/0
R1(config-if)# ip address 10.0.12.1 255.255.255.252
R1(config-if)# no shutdown
```

```
R2(config)#interface Loopback1
R2(config-if)# ip address 2.2.2.1 255.255.255.255
R2(config-if)#interface Loopback2
R2(config-if)# ip address 2.2.2.2 255.255.255.255
R2(config-if)#interface FastEthernet0/0
R2(config-if)# ip address 10.0.12.2 255.255.255.252
R2(config-if)# no shutdown
```

Configuration of EIGRP named mode and advertises the connected networks:

```
R1(config-if)#router eigrp CCIE
R1(config-router)#address-family ipv4 unicast autonomous-system 100
R1(config-router-af)#network 10.0.12.1 0.0.0.0
```

```
R2(config-if)#router eigrp CCIE
R2(config-router)#address-family ipv4 unicast autonomous-system 100
R2(config-router-af)#network 10.0.12.2 0.0.0.0
R1(config-router-af)#network 2.0.0.0
```

Verify the EIGRP neighbor relationship:

```
R1#show ip eigrp neighbors
EIGRP-IPv4 VR(CCIE) Address-Family Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                   (sec)         (ms)       Cnt
Num
0   10.0.12.2               Fa0/0                    13 00:00:27    1  3000  0   1
R1#
```

Verify the routing table of R1, you should find two EIGRP routes toward the two loopback networks:

```
R1#show ip route eigrp | beg Gate
Gateway of last resort is not set

      2.0.0.0/32 is subnetted, 2 subnets
D        2.2.2.1 [90/103040] via 10.0.12.2, 00:00:08, FastEthernet0/0
D        2.2.2.2 [90/103040] via 10.0.12.2, 00:00:08, FastEthernet0/0
R1#
```

Configure R1 to limit the number of learned EIGRP routes from R2 to 8 and if the percentage 50% of maximum prefixes is received, R1 triggers a warning message:

```
R1(config-router-af)#router ei CCIE
R1(config-router)#address-family ipv4 unicast autonomous-system 100
R1(config-router-af)#neighbor 10.0.12.2 maximum-prefix 8 50 warning-only
```

The number of EIGRP route are not reached yet, let's advertise two prefixes from R2:

```
R2(config-if)#interface Loopback3
R2(config-if)# ip address 2.2.2.3 255.255.255.255
R2(config-if)#interface Loopback4
R2(config-if)# ip address 2.2.2.4 255.255.255.255
```

We can see that the IOS on R1 displays a warning message telling that the threshold of number prefixes 4 is reached (which the percentage of the maximum prefix:

```
R1#
*Jul  5 12:35:19.479: %DUAL-4-PFXLIMITTHR: EIGRP-IPv4 100: Neighbor threshold
prefix level(4) reached.
R1#
```

Verify the routing table of R1:

```
R1#show ip route eigrp | beg Gate
Gateway of last resort is not set

      2.0.0.0/32 is subnetted, 4 subnets
D        2.2.2.1 [90/103040] via 10.0.12.2, 00:16:13, FastEthernet0/0
D        2.2.2.2 [90/103040] via 10.0.12.2, 00:16:13, FastEthernet0/0
D        2.2.2.3 [90/103040] via 10.0.12.2, 00:05:07, FastEthernet0/0
D        2.2.2.4 [90/103040] via 10.0.12.2, 00:02:29, FastEthernet0/0
R1#
```

Let's advertise 5 prefixes to reach the maximum number of prefixes 8:

```
R2(config)#int lo5
R2(config-if)#ip address 2.2.2.5 255.255.255.255
R2(config-if)#int lo6
R2(config-if)#ip add 2.2.2.6 255.255.255.255
R2(config-if)#int lo7
R2(config-if)#ip add 2.2.2.7 255.255.255.255
R2(config-if)#int lo8
R2(config-if)#ip add 2.2.2.8 255.255.255.255
R2(config-if)#int lo9
R2(config-if)#ip add 2.2.2.9 255.255.255.255
```

Now R1 displays another warning message telling that the maximum number of prefixes 8 is reached:

```
R1#
*Jul  5 12:37:32.859: %DUAL-3-PFXLIMIT: EIGRP-IPv4 100: Neighbor prefix limit
reached(8).
R1#
```

Verify the routing table of R1:

```
R1#show ip route eigrp | beg Gate
Gateway of last resort is not set

      2.0.0.0/32 is subnetted, 9 subnets
D        2.2.2.1 [90/103040] via 10.0.12.2, 00:23:18, FastEthernet0/0
D        2.2.2.2 [90/103040] via 10.0.12.2, 00:23:18, FastEthernet0/0
D        2.2.2.3 [90/103040] via 10.0.12.2, 00:12:12, FastEthernet0/0
D        2.2.2.4 [90/103040] via 10.0.12.2, 00:09:34, FastEthernet0/0
D        2.2.2.5 [90/103040] via 10.0.12.2, 00:03:53, FastEthernet0/0
D        2.2.2.6 [90/103040] via 10.0.12.2, 00:02:04, FastEthernet0/0
D        2.2.2.7 [90/103040] via 10.0.12.2, 00:01:53, FastEthernet0/0
D        2.2.2.8 [90/103040] via 10.0.12.2, 00:01:41, FastEthernet0/0
D        2.2.2.9 [90/103040] via 10.0.12.2, 00:01:28, FastEthernet0/0
R1#
```

Let's disable the Loopback interfaces from 3 to 9:

```
R2(config)#int range loopback 3-9
R2(config-if-range)#shutdown
```

Now configure R1 so that when it receives the maximum number of prefixes, a warning message will be displayed and the adjacency with R2 will be killed:

```
R1(config)#router ei CCIE
R1(config-router)#address-family ipv4 unicast autonomous-system 100
R1(config-router-af)#no neighbor 10.0.12.2 maximum-prefix 4 50 warning-only
R1(config-router-af)#neigh 10.0.12.2 maximum-prefix 8 50
```

Let's advertise the loopback 3 and 4:

```
R2(config)#int range loopback 3-4
R2(config-if-range)#no shutdown
```

A warning message is displayed because the threshold value 4 is reached:

```
R1#
*Jul  5 12:46:06.243: %DUAL-4-PFXLIMITTHR: EIGRP-IPv4 100: Neighbor threshold
prefix level(4) reached.
R1#
```

Now let's advertise the loopback interfaces from 5 to 9:

```
R2(config-if-range)#int range loopback 5-9
R2(config-if-range)#no shutdown
```

R1 displays a warning message because the maximum number of prefixes is reached (8) and another message is showed meaning that the adjacency is lost with R2:

```
R1#
*Jul  5 12:51:02.291: %DUAL-3-PFXLIMIT: EIGRP-IPv4 100: Neighbor prefix limit
reached(8).
```

```
*Jul  5 12:51:02.307: %DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 10.0.12.2
(FastEthernet0/0) is down: prefix-limit exceeded
R1#
```

The show ip eig nei command confirms the result, no adjacency with R2:

```
R1#show ip eigrp neighbors
EIGRP-IPv4 VR(CCIE) Address-Family Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                   (sec)         (ms)       Cnt
Num
-   10.0.12.2               Fa0/0                     -     --      -     -   -   -
R1#
```

R1 does not have any EIGRP routes:

```
R1#show ip route | beg Gate
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        10.0.12.0/30 is directly connected, FastEthernet0/0
L        10.0.12.1/32 is directly connected, FastEthernet0/0
R1#
```

Lab 13: Two EIGRP AS number and two equal paths to the same destination

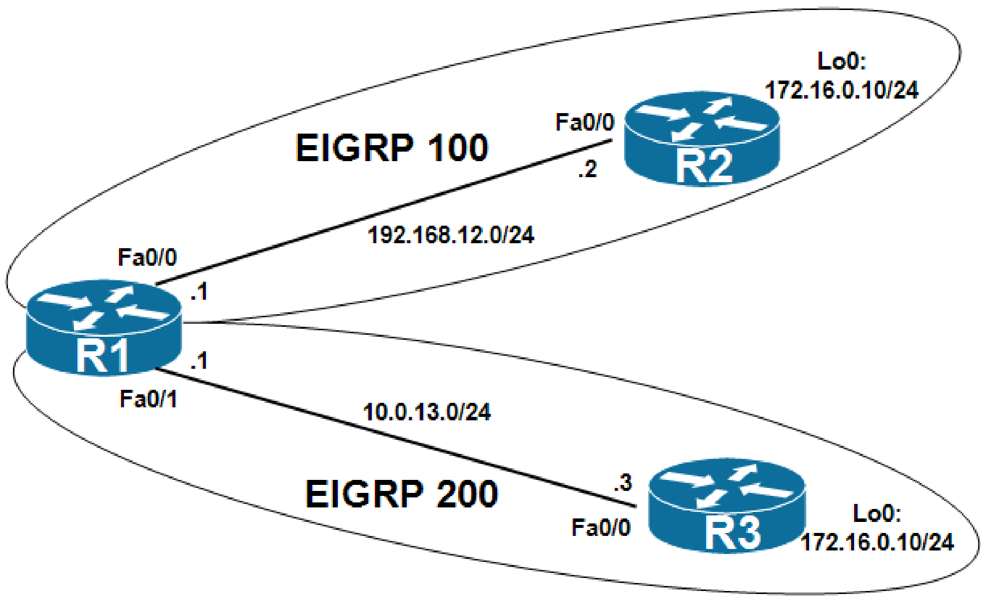

Configuration of IP addressing:

R1:

```
interface FastEthernet0/0
 ip address 192.168.12.1 255.255.255.0
 no shutdown
!
interface FastEthernet0/1
 ip address 10.0.13.1 255.255.255.0
 no shutdown
```

R2:

```
interface Loopback0
 ip address 172.16.0.10 255.255.255.0
!
interface FastEthernet0/0
 ip address 192.168.12.2 255.255.255.0
 no shutdown
```

R3:

```
interface Loopback0
 ip address 172.16.0.10 255.255.255.0
!
interface FastEthernet0/0
 ip address 10.0.13.3 255.255.255.0
 no shutdown
```

Configuration of EIGRP AS 100 between R1 and R2:

R1:

```
router eigrp 100
 network 192.168.12.0
```

```
 no auto-summary
```

R2:

```
router eigrp 100
 network 172.16.0.0
 network 192.168.12.0
 no auto-summary
```

Configuration of EIGRP AS 200 between R1 and R3:

R1:

```
router eigrp 200
 network 10.0.0.0
 no auto-summary
```

R3:

```
router eigrp 200
 network 10.0.0.0
 network 172.16.0.0
 no auto-summary
```

All routers are using the same bandwidth in their fastethernet interfaces: 100000 Kbit

Verification of the neighbor relationship:
R1 is adjacent with R2 in EIGRP AS 100 domain and with R3 in EIGRP AS 200 domain:

```
R1#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                   (sec)         (ms)       Cnt
Num
0   192.168.12.2            Fa0/0                    11 00:12:53  201  1206  0  4
EIGRP-IPv4 Neighbors for AS(200)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                   (sec)         (ms)       Cnt
Num
0   10.0.13.3               Fa0/1                    13 00:07:51  188  1128  0  2
R1#
```

R1 learns two EIGRP routes toward the same destination 172.16.0.0/24,one from R2 via EIGRP AS 100 domain and another route from R3 via EIGRP AS 200 domain:
We can see that the topology table of EIGRP AS 100 displays one successor R2 to 172.16.0.0/24 and an FD: 156160.
And the topology table of EIGRP AS 200 displays 0 successors in other words there are not a successors and the FD is Infinity.

```
R1#show ip eigrp topo
EIGRP-IPv4 Topology Table for AS(100)/ID(192.168.12.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 192.168.12.0/24, 1 successors, FD is 28160
        via Connected, FastEthernet0/0
P 172.16.0.0/24, 1 successors, FD is 156160
        via 192.168.12.2 (156160/128256), FastEthernet0/0
```

```
EIGRP-IPv4 Topology Table for AS(200)/ID(192.168.12.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 172.16.0.0/24, 0 successors, FD is Infinity
        via 10.0.13.3 (156160/128256), FastEthernet0/1
P 10.0.13.0/24, 1 successors, FD is 28160
        via Connected, FastEthernet0/1
R1#
```

As a result R1 installs the route learned from R2 via EIGRP 100 because R1 prefers the route that was learned through the EIGRP process with the lower Autonomous System (AS) number, EIGRP AS 100:

```
R1#show ip route eigrp  | beg Gate
Gateway of last resort is not set

      172.16.0.0/24 is subnetted, 1 subnets
D        172.16.0.0 [90/156160] via 192.168.12.2, 00:25:06, FastEthernet0/0
R1#
```

Let' remove EIGRP 200 between R1 and R2.
And let's configure EIGRP AS 99.

```
R1(config)#no router eig 200
R3(config)#no router eig 200
```

R1:

```
router eigrp 99
 network 10.0.0.0
 no auto-summary
```

R3:

```
router eigrp 99
 network 10.0.0.0
 network 172.16.0.0
 no auto-summary
```

Let's take a look at the topology table of R1:
Now the topology table of EIGRP AS 100 displays 0 successors in other words there are not a successors and the FD is Infinity.
And the topology table of EIGRP AS 99 displays 1 successor R3 to 172.16.0.0/24 and an FD: 156160.

```
R1#show ip eigrp topology
EIGRP-IPv4 Topology Table for AS(100)/ID(192.168.12.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 192.168.12.0/24, 1 successors, FD is 28160
        via Connected, FastEthernet0/0
P 172.16.0.0/24, 0 successors, FD is Infinity
        via 192.168.12.2 (156160/128256), FastEthernet0/0

EIGRP-IPv4 Topology Table for AS(99)/ID(192.168.12.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status
```

```
P 172.16.0.0/24, 1 successors, FD is 156160
        via 10.0.13.3 (156160/128256), FastEthernet0/1
P 10.0.13.0/24, 1 successors, FD is 28160
        via Connected, FastEthernet0/1
R1#
```

Let's verify the routing table of R1, R1 installs the EIGRP route learned through the lowest AS number 99 from R3:

```
R1#show ip route eigrp  | beg Gate
Gateway of last resort is not set

      172.16.0.0/24 is subnetted, 1 subnets
D        172.16.0.0 [90/156160] via 10.0.13.3, 00:02:31, FastEthernet0/1
R1#
```

Therefore when a router learns two EIGRP routes with two equal paths metric , the route learned through the lower Autonomous System (AS) number is installed in the routing table.

Let's decrease the bandwidth of the preferred path R3:

```
R1(config)#int fa0/1
R1(config-if)#band
R1(config-if)#bandwidth ?
  <1-10000000>  Bandwidth in kilobits
  inherit       Specify how bandwidth is inherited
  receive       Specify receive-side bandwidth

R1(config-if)#bandwidth 1000
```

Now the topology table of EIGRP AS 100 displays 1 successor R2 to 172.16.0.0/24 and an FD: 156160.
And the topology table of EIGRP AS 99 displays 0 successors in other words there are not a successors and the FD is Infinity.

```
R1(config-if)#do show ip ei topo
EIGRP-IPv4 Topology Table for AS(100)/ID(192.168.12.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 192.168.12.0/24, 1 successors, FD is 28160
        via Connected, FastEthernet0/0
P 172.16.0.0/24, 1 successors, FD is 156160
        via 192.168.12.2 (156160/128256), FastEthernet0/0

EIGRP-IPv4 Topology Table for AS(99)/ID(192.168.12.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 172.16.0.0/24, 0 successors, FD is Infinity
        via 10.0.13.3 (2690560/128256), FastEthernet0/1
P 10.0.13.0/24, 1 successors, FD is 2562560
        via Connected, FastEthernet0/1
R1#
```

As a result R1 prefers the route through R2 regardless of the Autonomous Systeme AS number:

```
R1(config-if)#do show ip route eigrp  | beg Gate
Gateway of last resort is not set

      172.16.0.0/24 is subnetted, 1 subnets
D        172.16.0.0 [90/156160] via 192.168.12.2, 00:01:51, FastEthernet0/0
R1#
```

Lab 14: Feasible Successor issues

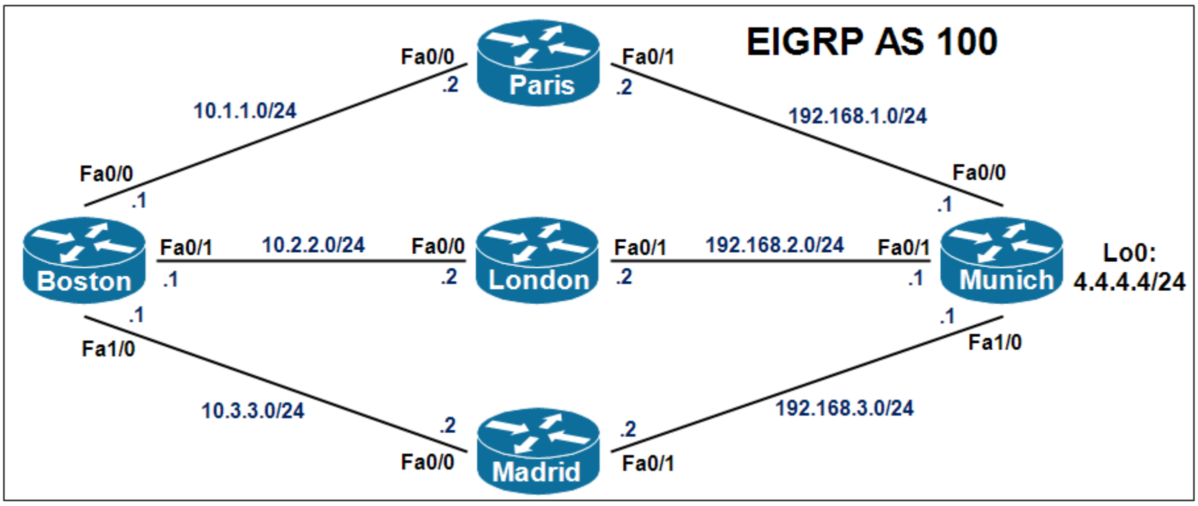

Configuration of all routers:

Boston:

```
interface FastEthernet0/0
 ip address 10.1.1.1 255.255.255.0
 no shutdown
!
interface FastEthernet0/1
 ip address 10.2.2.1 255.255.255.0
 delay 1
 no shutdown
!
interface FastEthernet1/0
 bandwidth 1000
 ip address 10.3.3.1 255.255.255.0
 delay 100
 no shutdown
!
router eigrp 100
 network 10.0.0.0
 no auto-summary
```

London:

```
interface FastEthernet0/0
 ip address 10.2.2.2 255.255.255.0
 no shutdown
!
interface FastEthernet0/1
 ip address 192.168.2.2 255.255.255.0
 no shutdown
!
router eigrp 100
 network 10.0.0.0
 network 192.168.2.0
 no auto-summary
```

Madrid:

```
interface FastEthernet0/0
 ip address 10.3.3.2 255.255.255.0
 no shutdown
!
interface FastEthernet0/1
 ip address 192.168.3.2 255.255.255.0
 no shutdown
!
router eigrp 100
 network 10.0.0.0
 network 192.168.3.0
 no auto-summary
```

Munich:

```
interface Loopback0
 ip address 4.4.4.4 255.255.255.0
!
interface FastEthernet0/0
 ip address 192.168.1.1 255.255.255.0
 no shutdown
!
interface FastEthernet0/1
 ip address 192.168.2.1 255.255.255.0
 no shutdown
!
interface FastEthernet1/0
 bandwidth 10000
 ip address 192.168.3.1 255.255.255.0
 delay 100
 no shutdown
!
router eigrp 100
 network 4.0.0.0
 network 192.168.0.0 0.0.255.255
 no auto-summary
```

Paris:

```
interface FastEthernet0/0
 ip address 10.1.1.2 255.255.255.0
 no shutdown
!
interface FastEthernet0/1
 ip address 192.168.1.2 255.255.255.0
 no shutdown
!
router eigrp 100
 offset-list 0 in 1000 FastEthernet0/1
 network 10.0.0.0
 network 192.168.1.0
 no auto-summary
```

We know that successor is the route with the best metric to reach a destination. That route is stored in the routing table. A feasible successor is a backup path to reach that same destination that can be used immediately if the successor route fails. These backup routes are stored in the topology table.

Let's verify the topology table of Boston with the show ip eigrp topology all-links command and we will focus over the destination subnet 4.4.4.0/24:

```
Boston#show ip eigrp topology all-links
IP-EIGRP Topology Table for AS(100)/ID(10.3.3.1)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 4.4.4.0/24, 1 successors, FD is 409856, serno 25
        via 10.2.2.2 (409856/409600), FastEthernet0/1
        via 10.1.1.2 (436200/410600), FastEthernet0/0
        via 10.3.3.2 (2739200/409600), FastEthernet1/0
P 10.3.3.0/24, 1 successors, FD is 2585600, serno 3
        via Connected, FastEthernet1/0
        via 10.1.1.2 (359400/333800), FastEthernet0/0
        via 10.2.2.2 (333056/332800), FastEthernet0/1, serno 19
P 10.2.2.0/24, 1 successors, FD is 256256, serno 15
        via Connected, FastEthernet0/1
P 10.1.1.0/24, 1 successors, FD is 281600, serno 12
        via Connected, FastEthernet0/0
P 192.168.1.0/24, 1 successors, FD is 307200, serno 24
        via 10.1.1.2 (307200/281600), FastEthernet0/0, serno 18
        via 10.2.2.2 (307456/307200), FastEthernet0/1, serno 23
        via 10.3.3.2 (2636800/307200), FastEthernet1/0, serno 21
P 192.168.2.0/24, 1 successors, FD is 281856, serno 11
        via 10.2.2.2 (281856/281600), FastEthernet0/1

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

        via 10.3.3.2 (2636800/307200), FastEthernet1/0, serno 22
P 192.168.3.0/24, 1 successors, FD is 307456, serno 26
        via 10.2.2.2 (307456/307200), FastEthernet0/1
        via 10.1.1.2 (333800/308200), FastEthernet0/0
        via 10.3.3.2 (2611200/281600), FastEthernet1/0
Boston#
```

From Boston to 4.4.4.0/24:
-The Path through London : FD=409856
-The Path through Madrid : FD=2739200
-The Path through Paris: FD= 436200

Therefore London is the Successor:

```
Boston#show ip route | inc 4.4.4.0
D 4.4.4.0 [90/409856] via 10.2.2.2, 00:22:36, FastEthernet0/1
Boston#
```

AD of Paris: 410600
AD of Madrid: 409600

409600 is LESS than 409856
410600 is GREATER than 409856

Therefore Madrid is the Feasible Successor as confirmed by the show ip eigrp topology command:

```
Boston#show ip eigrp topology
IP-EIGRP Topology Table for AS(100)/ID(10.3.3.1)
```

```
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 4.4.4.0/24, 1 successors, FD is 409856
        via 10.2.2.2 (409856/409600), FastEthernet0/1
        via 10.3.3.2 (2739200/409600), FastEthernet1/0
P 10.3.3.0/24, 1 successors, FD is 2585600
        via Connected, FastEthernet1/0
        via 10.1.1.2 (359400/333800), FastEthernet0/0
        via 10.2.2.2 (333056/332800), FastEthernet0/1, serno 19
P 10.2.2.0/24, 1 successors, FD is 256256
        via Connected, FastEthernet0/1
P 10.1.1.0/24, 1 successors, FD is 281600
        via Connected, FastEthernet0/0
P 192.168.1.0/24, 1 successors, FD is 307200
        via 10.1.1.2 (307200/281600), FastEthernet0/0, serno 18
P 192.168.2.0/24, 1 successors, FD is 281856
        via 10.2.2.2 (281856/281600), FastEthernet0/1
P 192.168.3.0/24, 1 successors, FD is 307456
        via 10.2.2.2 (307456/307200), FastEthernet0/1
        via 10.3.3.2 (2611200/281600), FastEthernet1/0
Boston#
```

Now assume the Successor London is down and fails.
Which path Boston will install in its routing table to reach 4.4.4.0/24?

By definition :The feasible successor routes are not used in the forwarding of traffic, but can be injected into the routing table immediately in place of a failed successor route, to avoid the querying process mentioned earlier.
Let's verify:

```
Boston#debug eigrp packets query
EIGRP Packets debugging is on
    (QUERY)
Boston#

Boston(config)#int fa0/1
Boston(config-if)#shutdown
Boston(config-if)#end
```

```
*Mar  1 00:06:14.135: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 10.2.2.2
(FastEthernet0/1) is down: interface down
*Mar  1 00:06:14.147: EIGRP: Enqueueing QUERY on FastEthernet0/0 iidbQ un/rely 0/1
serno 27-30
*Mar  1 00:06:14.147: EIGRP: Enqueueing QUERY on FastEthernet1/0 iidbQ un/rely 0/1
serno 27-30
*Mar  1 00:06:14.155: EIGRP: Enqueueing QUERY on FastEthernet0/0 nbr 10.1.1.2
iidbQ un/rely 0/0 peerQ un/rely 0/0 serno 27-30
*Mar  1 00:06:14.159: EIGRP: Sending QUERY on FastEthernet0/0
*Mar  1 00:06:14.159:   AS 100, Flags 0x0, Seq 40/0 idbQ 0/0 iidbQ un/rely 0/0
serno 27-30
*Mar  1 00:06:14.167: EIGRP: Enqueueing QUERY on FastEthernet1/0 nbr 10.3.3.2
iidbQ un/rely 0/0 peerQ un/rely 0/0 serno 27-30
*Mar  1 00:06:14.171: EIGRP: Sending QUERY on FastEthernet1/0
*Mar  1 00:06:14.175:   AS 100, Flags 0x0, Seq 41/0 idbQ 0/0 iidbQ un/rely 0/0
serno 27-30
*Mar  1 00:06:14.255: EIGRP: Received QUERY on FastEthernet1/0 nbr 10.3.3.2
*Mar  1 00:06:14.259:   AS 100, Flags 0x0, Seq 29/0 idbQ 0/0 iidbQ un/rely 0/0
peerQ un/rely 0/0
```

Boston installs the path through Paris instead of the FS Madrid.
because the path paris has a better FD 436200 than the FD of Madrid 2739200:

```
Boston#show ip route eigrp | inc 4.4.4.0
D 4.4.4.0 [90/436200] via 10.1.1.2, 00:00:38, FastEthernet0/0
Boston#
```

```
Boston#show ip eigrp topology
IP-EIGRP Topology Table for AS(100)/ID(10.3.3.1)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 4.4.4.0/24, 1 successors, FD is 436200
        via 10.1.1.2 (436200/410600), FastEthernet0/0
        via 10.3.3.2 (2739200/409600), FastEthernet1/0
P 10.3.3.0/24, 1 successors, FD is 2585600
        via Connected, FastEthernet1/0
        via 10.1.1.2 (359400/333800), FastEthernet0/0
P 10.2.2.0/24, 1 successors, FD is 359400
        via 10.1.1.2 (359400/333800), FastEthernet0/0
        via 10.3.3.2 (2662400/332800), FastEthernet1/0, serno 31
P 10.1.1.0/24, 1 successors, FD is 281600
        via Connected, FastEthernet0/0
P 192.168.1.0/24, 1 successors, FD is 307200
        via 10.1.1.2 (307200/281600), FastEthernet0/0, serno 18
P 192.168.2.0/24, 1 successors, FD is 333800
        via 10.1.1.2 (333800/308200), FastEthernet0/0
        via 10.3.3.2 (2636800/307200), FastEthernet1/0, serno 22
P 192.168.3.0/24, 1 successors, FD is 333800
        via 10.1.1.2 (333800/308200), FastEthernet0/0

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

        via 10.3.3.2 (2611200/281600), FastEthernet1/0
Boston#
```

we can see from the debug eigrp packets query that Boston is sending the query packets to find a new route , but by definition When EIGRP loses its successor route and does not have a FS, it sends out a query message to all of its neighbors asking if they know a path. In this case even if Boston has an FS, it will send a query packet out its interfaces, so the FS in this case is not useful to avoid a query messages.

Lab 15: EIGRP Summary and Leak Map

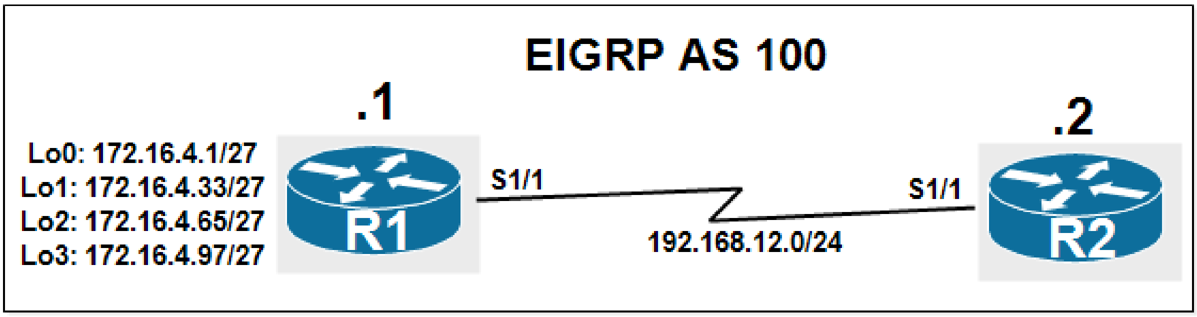

Route leaking can be used when you have summarized a subnet but at the same time, you want also to advertise a more specific subnet
Basic Configuration of R1 and R2:

R1:

```
interface Loopback0
 ip address 172.16.4.1 255.255.255.224
!
interface Loopback1
 ip address 172.16.4.33 255.255.255.224
!
interface Loopback2
 ip address 172.16.4.65 255.255.255.224
!
interface Loopback3
 ip address 172.16.4.97 255.255.255.224
!
interface Serial1/1
 ip address 192.168.12.1 255.255.255.0
 no shutdown
```

R2:

```
interface Serial1/1
 ip address 192.168.12.2 255.255.255.0
 no shudown
```

Configure EIGRP AS 100 and advertises all networks presents in the topology:

R1:

```
router eigrp 100
 network 172.16.0.0
 network 192.168.12.0
 no auto-summary
```

R2:

```
router eigrp 100
 network 192.168.12.0
 no auto-summary
```

R2 receives 4 EIGRP routes (the four subnets connected to the loopback interfaces of R1):

```
R2#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is not set

      172.16.0.0/27 is subnetted, 4 subnets
D        172.16.4.0 [90/2297856] via 192.168.12.1, 00:00:36, Serial1/1
D        172.16.4.32 [90/2297856] via 192.168.12.1, 00:00:36, Serial1/1
D        172.16.4.64 [90/2297856] via 192.168.12.1, 00:00:36, Serial1/1
D        172.16.4.96 [90/2297856] via 192.168.12.1, 00:00:36, Serial1/1
R2#
```

Let's summarize the four loopback addresses to the most specific summary possible:

R1:

```
int s1/1
 ip summary-address eigrp 100 172.16.4.0 255.255.255.128
```

Let's verify the routing table of R2 :
Only one summarized subnet is installed on R2.

```
R2#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is not set

      172.16.0.0/25 is subnetted, 1 subnets
D        172.16.4.0 [90/2297856] via 192.168.12.1, 00:00:14, Serial1/1
R2#
```

Let' configure the route leaking on R1, We identify the specific subnet 172.16.4.32/27 by a standard ACL named LEAK-MAP then we create a route-map called LEAKMAP which matches the ACL with the permit action:

On R1:

```
ip access-list standard LEAK-MAP
 permit 172.16.4.32 0.0.0.31
!
route-map LEAKMAP permit 10
 match ip address LEAK-MAP
!
int s1/1
 ip summary-address eigrp 100 172.16.4.0 255.255.255.128 leak-map LEAKMAP
```

Let's verify the routing table of R2 once again:

We can see that R2 receives two EIGRP Routes:

-Summarized (172.16.4.0/25)
-Specific (172.16.4.32/27)

```
R2#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is not set

      172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
D        172.16.4.0/25 [90/2297856] via 192.168.12.1, 00:01:29, Serial1/1
D        172.16.4.32/27 [90/2297856] via 192.168.12.1, 00:02:21, Serial1/1
R2#
```

Lab 18: show ip eigrp topology all-links subcommand

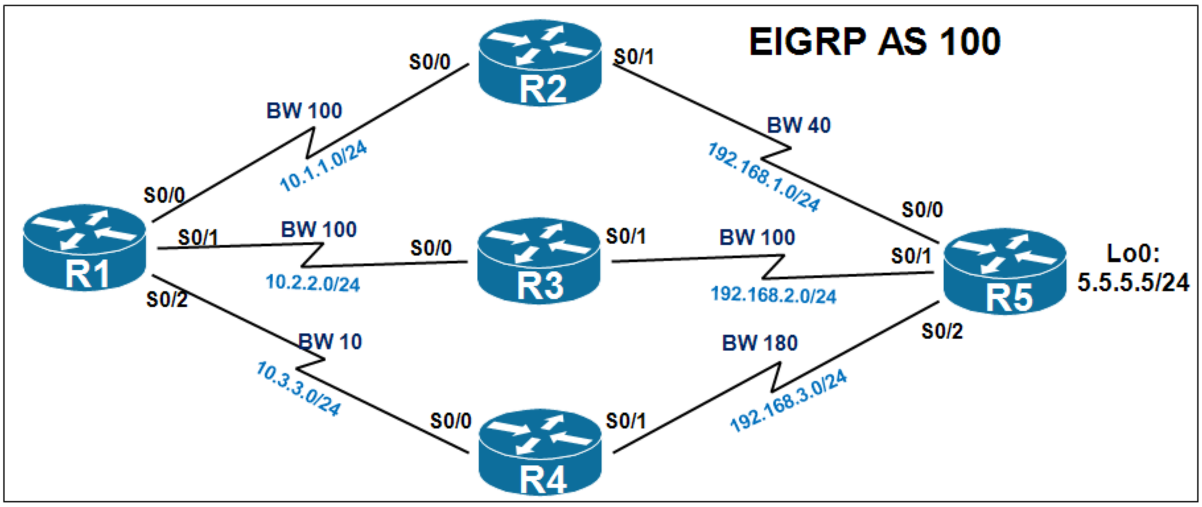

Configuration of all routers:

On R1:

```
R1#show run int   serial 0/0
Building configuration...

Current configuration : 97 bytes
!
interface Serial0/0
 bandwidth 100
 ip address 10.1.1.1 255.255.255.0
 clock rate 2000000
end
```

```
R1#show run int   serial 0/1
Building configuration...

Current configuration : 97 bytes
!
interface Serial0/1
 bandwidth 100
 ip address 10.2.2.1 255.255.255.0
 clock rate 2000000
end
```

```
R1#show run int   serial 0/2
Building configuration...

Current configuration : 96 bytes
!
interface Serial0/2
 bandwidth 10
 ip address 10.3.3.1 255.255.255.0
 clock rate 2000000
```

```
end
```

On R2:

```
R2#show run int   serial 0/0
Building configuration...

Current configuration : 128 bytes
!
interface Serial0/0
 bandwidth 100
 ip address 10.1.1.2 255.255.255.0
 clock rate 2000000
end
```

```
R2#show run int   serial 0/1
Building configuration...

Current configuration : 99 bytes
!
interface Serial0/1
 bandwidth 40
 ip address 192.168.1.2 255.255.255.0
 clock rate 2000000
end
```

On R3:

```
R3#show run int   serial 0/0
Building configuration...

Current configuration : 97 bytes
!
interface Serial0/0
 bandwidth 100
 ip address 10.2.2.2 255.255.255.0
 clock rate 2000000
end
```

```
R3#show run int   serial 0/1
Building configuration...

Current configuration : 100 bytes
!
interface Serial0/1
 bandwidth 100
 ip address 192.168.2.2 255.255.255.0
 clock rate 2000000
end
```

On R4:

```
R4#show run int   serial 0/0
Building configuration...
```

```
Current configuration : 96 bytes
!
interface Serial0/0
 bandwidth 10
 ip address 10.3.3.2 255.255.255.0
 clock rate 2000000
end
```

```
R4#show run int   serial 0/1
Building configuration...

Current configuration : 100 bytes
!
interface Serial0/1
 bandwidth 180
 ip address 192.168.3.2 255.255.255.0
 clock rate 2000000
end
```

On R5:

```
R5#show run int loopback 0
Building configuration...

Current configuration : 61 bytes
!
interface Loopback0
 ip address 5.5.5.5 255.255.255.0
end
```

```
R5#show run int   serial 0/0
Building configuration...

Current configuration : 99 bytes
!
interface Serial0/0
 bandwidth 40
 ip address 192.168.1.1 255.255.255.0
 clock rate 2000000
end
```

```
R5#show run int   serial 0/1
Building configuration...

Current configuration : 100 bytes
!
interface Serial0/1
 bandwidth 100
 ip address 192.168.2.1 255.255.255.0
 clock rate 2000000
end
```

```
R5#show run int   serial 0/2
```

```
Building configuration...

Current configuration : 100 bytes
!
interface Serial0/2
 bandwidth 180
 ip address 192.168.3.1 255.255.255.0
 clock rate 2000000
end
```

R3 is the Successor because it has a lowest metric (FD).

```
R1#show ip route eigrp
     5.0.0.0/24 is subnetted, 1 subnets
D       5.5.5.0 [90/26752000] via 10.2.2.2, 00:02:38, Serial0/1
D    192.168.1.0/24 [90/65024000] via 10.1.1.2, 00:10:15, Serial0/0
D    192.168.2.0/24 [90/26624000] via 10.2.2.2, 00:10:15, Serial0/1
D    192.168.3.0/24 [90/27136000] via 10.2.2.2, 00:10:15, Serial0/1
R1#
```

Let's check the topology table of R1:
We can see that R3 is the Feasible Successor because the AD ' R3 14862080 is less than the FD ' R1 26752000:

```
R1#show ip eigrp topology
IP-EIGRP Topology Table for AS(100)/ID(10.3.3.1)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 5.5.5.0/24, 1 successors, FD is 26752000
        via 10.2.2.2 (26752000/26240000), Serial0/1
        via 10.3.3.2 (257152000/14862080), Serial0/2
P 10.3.3.0/24, 1 successors, FD is 256512000
        via Connected, Serial0/2
P 10.2.2.0/24, 1 successors, FD is 26112000
        via Connected, Serial0/1
P 10.1.1.0/24, 1 successors, FD is 26112000
        via Connected, Serial0/0
P 192.168.1.0/24, 1 successors, FD is 65024000
        via 10.1.1.2 (65024000/64512000), Serial0/0
P 192.168.2.0/24, 1 successors, FD is 26624000
        via 10.2.2.2 (26624000/26112000), Serial0/1
P 192.168.3.0/24, 1 successors, FD is 27136000
        via 10.2.2.2 (27136000/26624000), Serial0/1
        via 10.3.3.2 (257024000/14734080), Serial0/2
R1#
```

The show ip eigrp topology all-links command displays all learned routes that are neither a Successor nor Feasible Successor .

Let's verify on R1:

Neighbor R2 is not in the topology table using the "all-links" option

```
R1#show ip eigrp topology all-links
IP-EIGRP Topology Table for AS(100)/ID(10.3.3.1)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
```

```
       r - reply Status, s - sia Status

P 5.5.5.0/24, 1 successors, FD is 26752000, serno 15
        via 10.2.2.2 (26752000/26240000), Serial0/1
        via 10.3.3.2 (257152000/14862080), Serial0/2
P 10.3.3.0/24, 1 successors, FD is 256512000, serno 3
        via Connected, Serial0/2
P 10.2.2.0/24, 1 successors, FD is 26112000, serno 2
        via Connected, Serial0/1
        via 10.3.3.2 (258048000/27136000), Serial0/2
P 10.1.1.0/24, 1 successors, FD is 26112000, serno 1
        via Connected, Serial0/0
        via 10.3.3.2 (258560000/27648000), Serial0/2
P 192.168.1.0/24, 1 successors, FD is 65024000, serno 5
        via 10.1.1.2 (65024000/64512000), Serial0/0
        via 10.2.2.2 (65536000/65024000), Serial0/1
        via 10.3.3.2 (257536000/65024000), Serial0/2
P 192.168.2.0/24, 1 successors, FD is 26624000, serno 4
        via 10.2.2.2 (26624000/26112000), Serial0/1
        via 10.3.3.2 (257536000/26624000), Serial0/2
P 192.168.3.0/24, 1 successors, FD is 27136000, serno 10

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

        via 10.2.2.2 (27136000/26624000), Serial0/1
        via 10.3.3.2 (257024000/14734080), Serial0/2
R1#
```

Why? R1 is simply not learning a route to 5.5.5.0/24 from R2

According to the choice of bandwidth, R2's best route to 5.5.5.0/24 is through R1,R2 won't advertise a route to 5.5.5.0/24 to R1 due to split horizon rule.

Looks carefully the bandwidth values used in the topology the Path R2--R3--R5 is more faster than the path through directly R5.If you just did the math the metric through R1 is better (higher bandwidth) than the metric (slower bandwidth 40) through directly R5 .

In other words R1 didn't learn that route from R2 because the split horizon rule.To test it either disable split horizon or change the bandwidth so that the best path of R2 for 4.4.4.0/24 will be direct to R5 instead of through R1.

```
R2#show ip route | inc 5.5.5.0
D       5.5.5.0 [90/27264000] via 10.1.1.1, 00:05:47, Serial0/0
R2#
```

Disable the split-horizon rule on R2's s0/0:

```
R2(config)#int s0/0
R2(config-if)#no ip split-horizon eigrp 100
```

Because now the split-horizon is disabled, R1 can learn the route from R2 as a result the route is added in the topology table as shown by the output of the show ip eigrp topology all-links command:

```
R1#show ip eigrp topology all-links
IP-EIGRP Topology Table for AS(100)/ID(10.3.3.1)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
```

```
       r - reply Status, s - sia Status

P 5.5.5.0/24, 1 successors, FD is 26752000, serno 15
        via 10.2.2.2 (26752000/26240000), Serial0/1
        via 10.1.1.2 (27776000/27264000), Serial0/0
        via 10.3.3.2 (257152000/14862080), Serial0/2
P 10.3.3.0/24, 1 successors, FD is 256512000, serno 3
        via Connected, Serial0/2
        via 10.1.1.2 (257536000/257024000), Serial0/0
P 10.2.2.0/24, 1 successors, FD is 26112000, serno 2
        via Connected, Serial0/1
        via 10.1.1.2 (27136000/26624000), Serial0/0
        via 10.3.3.2 (258048000/27136000), Serial0/2
P 10.1.1.0/24, 1 successors, FD is 26112000, serno 1
        via Connected, Serial0/0
        via 10.1.1.2 (26624000/26112000), Serial0/0
        via 10.3.3.2 (258560000/27648000), Serial0/2
P 192.168.1.0/24, 1 successors, FD is 65024000, serno 5
        via 10.1.1.2 (65024000/64512000), Serial0/0
        via 10.2.2.2 (65536000/65024000), Serial0/1

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

        via 10.3.3.2 (257536000/65024000), Serial0/2
P 192.168.2.0/24, 1 successors, FD is 26624000, serno 4
        via 10.2.2.2 (26624000/26112000), Serial0/1
        via 10.1.1.2 (27648000/27136000), Serial0/0
        via 10.3.3.2 (257536000/26624000), Serial0/2
P 192.168.3.0/24, 1 successors, FD is 27136000, serno 10
        via 10.2.2.2 (27136000/26624000), Serial0/1
        via 10.1.1.2 (28160000/27648000), Serial0/0
        via 10.3.3.2 (257024000/14734080), Serial0/2
R1#
```

Lab 19: Successor and Feasible Successor

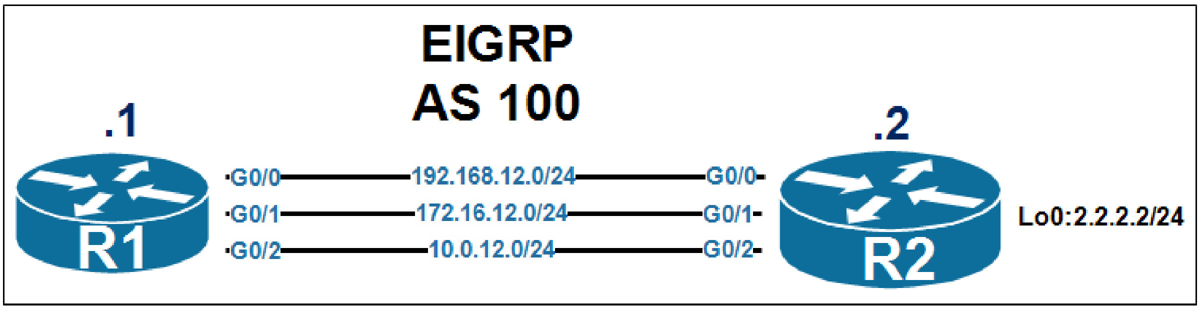

Configuration of all routers:

R1

```
interface GigabitEthernet0/0
 ip address 192.168.12.1 255.255.255.0
 duplex auto
 speed auto
!
interface GigabitEthernet0/1
 ip address 172.16.12.1 255.255.255.0
 duplex auto
 speed auto
!
interface GigabitEthernet0/2
 ip address 10.0.12.1 255.255.255.0
 duplex auto
 speed auto
!
router eigrp 100
 network 10.0.0.0
 network 172.16.0.0
 network 192.168.12.0
```

R2

```
interface Loopback0
 ip address 2.2.2.2 255.255.255.0
!
interface GigabitEthernet0/0
 ip address 192.168.12.2 255.255.255.0
 duplex auto
 speed auto
!
interface GigabitEthernet0/1
 ip address 172.16.12.2 255.255.255.0
 duplex auto
 speed auto
!
interface GigabitEthernet0/2
 ip address 10.0.12.2 255.255.255.0
 duplex auto
 speed auto
!
router eigrp 100
 network 2.0.0.0
```

```
 network 10.0.0.0
 network 172.16.0.0
 network 192.168.12.0
```

R2 advertises a route toward the subnet 2.2.2.0 through the three links. R1 receives three routes or Successors from R2 with the same FD (metric) = 130816 therfore it installs a load-balancing toward 2.2.2.0/24 through 3 links:

```
R1#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is not set

      2.0.0.0/24 is subnetted, 1 subnets
D        2.2.2.0 [90/130816] via 192.168.12.2, 00:00:09, GigabitEthernet0/0
                 [90/130816] via 172.16.12.2, 00:00:09, GigabitEthernet0/1
                 [90/130816] via 10.0.12.2, 00:00:09, GigabitEthernet0/2
R1#
```

The show ip eigrp topology command displays only the Successor and the Feasible Successor in the topology table ,in this case R1 has 3 successors as highlighted:

```
R1#sho ip eigrp topology
EIGRP-IPv4 Topology Table for AS(100)/ID(192.168.12.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 192.168.12.0/24, 1 successors, FD is 2816
        via Connected, GigabitEthernet0/0
P 2.2.2.0/24, 3 successors, FD is 130816
        via 10.0.12.2 (130816/128256), GigabitEthernet0/2
        via 172.16.12.2 (130816/128256), GigabitEthernet0/1
        via 192.168.12.2 (130816/128256), GigabitEthernet0/0
P 172.16.12.0/24, 1 successors, FD is 2816
        via Connected, GigabitEthernet0/1
P 10.0.12.0/24, 1 successors, FD is 2816
        via Connected, GigabitEthernet0/2
R1#
```

The show ip eigrp topology all-links command displays all learned routes,we have three learned routes :

```
R1#show ip eigrp topology all-links
EIGRP-IPv4 Topology Table for AS(100)/ID(192.168.12.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 192.168.12.0/24, 1 successors, FD is 2816, serno 21
        via Connected, GigabitEthernet0/0
        via 172.16.12.2 (3072/2816), GigabitEthernet0/1
        via 10.0.12.2 (3072/2816), GigabitEthernet0/2
P 2.2.2.0/24, 3 successors, FD is 130816, serno 24
```

```
        via 10.0.12.2 (130816/128256), GigabitEthernet0/2
        via 172.16.12.2 (130816/128256), GigabitEthernet0/1
        via 192.168.12.2 (130816/128256), GigabitEthernet0/0
P 172.16.12.0/24, 1 successors, FD is 2816, serno 2
        via Connected, GigabitEthernet0/1
        via 192.168.12.2 (3072/2816), GigabitEthernet0/0
        via 10.0.12.2 (3072/2816), GigabitEthernet0/2
P 10.0.12.0/24, 1 successors, FD is 2816, serno 12
        via Connected, GigabitEthernet0/2
        via 192.168.12.2 (3072/2816), GigabitEthernet0/0
        via 172.16.12.2 (3072/2816), GigabitEthernet0/1
R1#
```

Let's change the bandwitdh as follow, by changing the bandwidth in the interfaces g0/0 and g0/2 will force R1 to compute the metrics through the links g0/0--g0/0 and g0/2--g0/2:

```
R1(config-if)#in g0/0
R1(config-if)#bandwidth 100
R1(config-if)#in g0/2
R1(config-if)#bandwidth 100
```

Let's see the topology table of R1:

The Feasible Distance FD:
-The route through the link g0/0--g0/0: FD = 25728256
-The route through the link g0/2--g0/2: FD = 25728256
-The route through the link g0/1--g0/1: FD = 130816

The FD through the link g0/1--g0/1 is better than the FD of the links g0/0--g0/0 and g0/2--g0/2. The route learned through g0/1--g0/1 is the Successor.

The Advertised Distance:
-The route through the link g0/2--g0/2 : AD = 128256
-The route through the link g0/0--g0/0 : AD = 128256

The AD of the two links meets the feasibility condition : AD 128256 is less than the FD's R1 130816,therfore the routes learned through the two links are considered as the Feasible Successors.

The show ip eigrp topology all-link command shown one Successor through the link g0/1--g0/1 and two Feasible Successors through g0/0--g0/0 and g0/2--g0/2:

```
R1#show ip eigrp topology all-links
EIGRP-IPv4 Topology Table for AS(100)/ID(192.168.12.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 192.168.12.0/24, 1 successors, FD is 25600256, serno 25
        via Connected, GigabitEthernet0/0
        via 172.16.12.2 (25600512/25600256), GigabitEthernet0/1
        via 10.0.12.2 (25600512/25600256), GigabitEthernet0/2
P 2.2.2.0/24, 1 successors, FD is 130816, serno 28
        via 172.16.12.2 (130816/128256), GigabitEthernet0/1
        via 10.0.12.2 (25728256/128256), GigabitEthernet0/2
        via 192.168.12.2 (25728256/128256), GigabitEthernet0/0
P 172.16.12.0/24, 1 successors, FD is 2816, serno 2
        via Connected, GigabitEthernet0/1
        via 192.168.12.2 (25600512/2816), GigabitEthernet0/0
        via 10.0.12.2 (25600512/2816), GigabitEthernet0/2
```

```
P 10.0.12.0/24, 1 successors, FD is 25600256, serno 26
        via Connected, GigabitEthernet0/2
        via 192.168.12.2 (25600512/25600256), GigabitEthernet0/0
        via 172.16.12.2 (25600512/25600256), GigabitEthernet0/1

R1#
```

After the change of the bandwidth R1 preferred now the route through the link g0/1--g0/1:

```
R1#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is not set

      2.0.0.0/24 is subnetted, 1 subnets
D        2.2.2.0 [90/130816] via 172.16.12.2, 00:03:58, GigabitEthernet0/1
R1#
```

Let's disables the link g0/1--g0/1:

```
R1(config)#int g0/1
R1(config-if)#shutdown
R1(config-if)#end
```

After that the Successor failed R1 installs the two Feasible Successor and does a Load-Balancing because they have the same FD:

```
R1#show ip route eigrp
Codes: L - local, 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, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is not set

      2.0.0.0/24 is subnetted, 1 subnets
D        2.2.2.0 [90/25728256] via 192.168.12.2, 00:00:07, GigabitEthernet0/0
                 [90/25728256] via 10.0.12.2, 00:00:07, GigabitEthernet0/2
R1#
```

Now the topology table displays two Successors :

```
R1#show ip eigrp topology all-links
EIGRP-IPv4 Topology Table for AS(100)/ID(192.168.12.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status
```

```
P 192.168.12.0/24, 1 successors, FD is 25600256, serno 30
        via Connected, GigabitEthernet0/0
        via 10.0.12.2 (25600512/25600256), GigabitEthernet0/2
P 2.2.2.0/24, 2 successors, FD is 130816, serno 32
        via 10.0.12.2 (25728256/128256), GigabitEthernet0/2
        via 192.168.12.2 (25728256/128256), GigabitEthernet0/0
P 10.0.12.0/24, 1 successors, FD is 25600256, serno 26
        via Connected, GigabitEthernet0/2
        via 192.168.12.2 (25600512/25600256), GigabitEthernet0/0
R1#
```

Before continuing let's enables the interface g0/1 of R1:

```
R1(config)#int g0/1
R1(config-if)#no shutdown
```

Before increasing the AD's R2 through the Link G0/2--G0/2,let's verify the topology table of R2: The FD of R2 to 2.2.2.0 is 128256 which is the AD advertised by R2 to R1:

```
R2#show ip ei topology
EIGRP-IPv4 Topology Table for AS(1)/ID(192.168.12.2)

EIGRP-IPv4 Topology Table for AS(100)/ID(2.2.2.2)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 192.168.12.0/24, 1 successors, FD is 25600256
        via Connected, GigabitEthernet0/0
P 2.2.2.0/24, 1 successors, FD is 128256
        via Connected, Loopback0
P 172.16.12.0/24, 1 successors, FD is 2816
        via Connected, GigabitEthernet0/1
P 10.0.12.0/24, 1 successors, FD is 25600256
        via Connected, GigabitEthernet0/2
R2#
```

Now we will force the link g0/2--g0/2 to be a non-FS , by using an offset-list we will increase the AD advertised by R2 through the link G0/2 and it will be greater than the FD of the successor through g0/1--g0/1 which is 130816:

```
R2(config)#access-list 1 permit 2.2.2.0 0.0.0.255
R2(config)#router eigrp 100
R2(config-router)#offset-list 1 out 100000 g0/2
```

R2 calculates a new AD for 2.2.2.0 as follow: 128256+100000=228256
The AD of the route through g0/2--g0/2 has an AD 228256 greater than the FD of the successor 130816, it does not meet the feasibility condition therefore it is not considered as the Feasible Successor.

By the show ip eigrp topology table command displays only the Successor and the Feasible Successor.

Let's verify with this command, we can see that the output shown one Successor through g0/1--g0/1 and only one Feasible Successor g0/0--g0/0, the route g0/2--g0/2 is not displayed because it is neither a Successor nor Feasible Successor:

```
R1#show ip eigrp topology
EIGRP-IPv4 Topology Table for AS(100)/ID(192.168.12.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
```

```
       r - reply Status, s - sia Status

P 192.168.12.0/24, 1 successors, FD is 25600256
        via Connected, GigabitEthernet0/0
P 2.2.2.0/24, 1 successors, FD is 130816
        via 172.16.12.2 (130816/128256), GigabitEthernet0/1
        via 192.168.12.2 (25728256/128256), GigabitEthernet0/0
P 172.16.12.0/24, 1 successors, FD is 2816
        via Connected, GigabitEthernet0/1
P 10.0.12.0/24, 1 successors, FD is 25600256
        via Connected, GigabitEthernet0/2
R1#
```

By definition the show ip eigrp topology all-links displays all learned routes neither a successor nor feasible successor.

let's verify :

As expected the output below includes the non-FS route g0/2--g0/2.

Notice the topology table of R1 which lists the Successor g0/1, the Feasible Successor g0/0 and the path g0/2 which is neither a successor nor feasible successor:

```
R1#show ip eigrp topology all
R1#show ip eigrp topology all-links
EIGRP-IPv4 Topology Table for AS(100)/ID(192.168.12.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 192.168.12.0/24, 1 successors, FD is 25600256, serno 90
        via Connected, GigabitEthernet0/0
        via 10.0.12.2 (25600512/25600256), GigabitEthernet0/2
        via 172.16.12.2 (25600512/25600256), GigabitEthernet0/1
P 2.2.2.0/24, 1 successors, FD is 130816, serno 91
        via 172.16.12.2 (130816/128256), GigabitEthernet0/1
        via 192.168.12.2 (25728256/128256), GigabitEthernet0/0
        via 10.0.12.2 (25828256/228256), GigabitEthernet0/2
P 172.16.12.0/24, 1 successors, FD is 2816, serno 79
        via Connected, GigabitEthernet0/1
        via 192.168.12.2 (25600512/2816), GigabitEthernet0/0
        via 10.0.12.2 (25600512/2816), GigabitEthernet0/2
P 10.0.12.0/24, 1 successors, FD is 25600256, serno 26
        via Connected, GigabitEthernet0/2
        via 192.168.12.2 (25600512/25600256), GigabitEthernet0/0
        via 172.16.12.2 (25600512/25600256), GigabitEthernet0/1
R1#
```

Here the detailed topology table of the subnet 2.2.2.0/24 on R1 which shows all learned routes through all the three links so the show ip eigrp topology all-links command includes non-FS routes.

```
R1#show ip eigrp topology 2.2.2.0/24
EIGRP-IPv4 Topology Entry for AS(100)/ID(192.168.12.1) for 2.2.2.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 130816
  Descriptor Blocks:
  172.16.12.2 (GigabitEthernet0/1), from 172.16.12.2, Send flag is 0x0
      Composite metric is (130816/128256), route is Internal
      Vector metric:
        Minimum bandwidth is 1000000 Kbit
```

```
        Total delay is 5010 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
        Originating router is 2.2.2.2
  192.168.12.2 (GigabitEthernet0/0), from 192.168.12.2, Send flag is 0x0
      Composite metric is (25728256/128256), route is Internal
      Vector metric:
        Minimum bandwidth is 100 Kbit
        Total delay is 5010 microseconds
        Reliability is 255/255
        Load is 2/255
        Minimum MTU is 1500
        Hop count is 1
        Originating router is 2.2.2.2
  10.0.12.2 (GigabitEthernet0/2), from 10.0.12.2, Send flag is 0x0
      Composite metric is (25828256/228256), route is Internal
      Vector metric:
        Minimum bandwidth is 100 Kbit
        Total delay is 8916 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
        Originating router is 2.2.2.2
R1#
```

Lab 20: Next-hop field between two AS eigrp IPv6

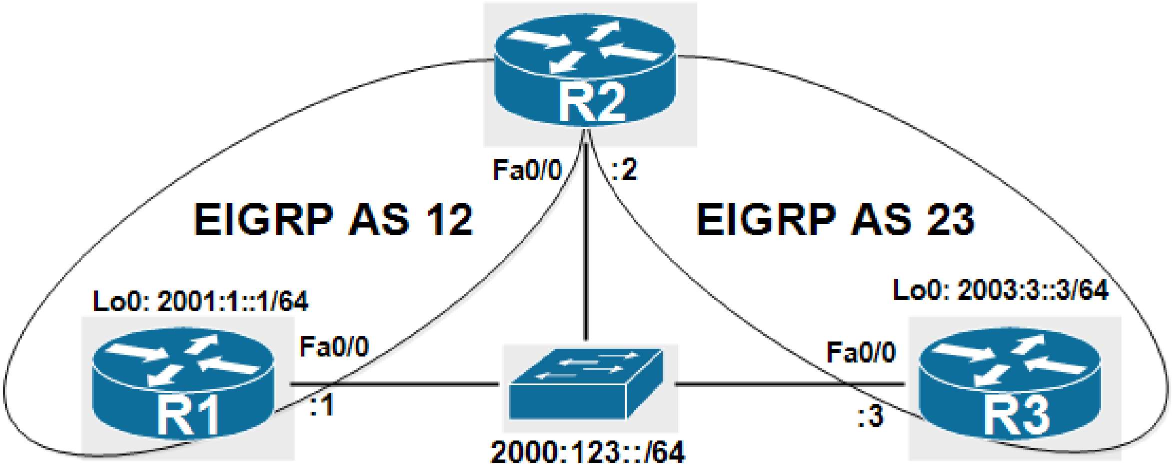

-R2 redistributes between EIGRP AS 12 and EIGRP AS 23.

R1 and R3 receives the prefixes 2003:3::/64 and 2001:1::/64 respectively as an external EIGRP Route:

```
R1#show ipv6 route | s 2003:3::/64
EX  2003:3::/64 [170/1333760]
     via FE80::C801:26FF:FEE4:8, FastEthernet0/0
R1#
```

```
R3#show ipv6 route | s 2001:1::/64
EX  2001:1::/64 [170/1333760]
     via FE80::C801:26FF:FEE4:8, FastEthernet0/0
R3#
```

To reach these prefixes ,R1 and R3 are going through R2 because the next-hop Field of the EIGRP updates (packets) are set to 0.0.0.0 which means that to reach these prefixes ,the routers R1 and R3 will use the router which advertises these external prefixes R2:

```
R1#traceroute 2003:3::3
Type escape sequence to abort.
Tracing the route to 2003:3::3

  1 2000:123::2 108 msec 76 msec 68 msec
  2 2000:123::3 152 msec 100 msec 156 msec
R1#
```

```
R3#traceroute 2001:1::1
Type escape sequence to abort.
Tracing the route to 2001:1::1

  1 2000:123::2 116 msec 88 msec 4 msec
  2 2000:123::1 204 msec 168 msec 124 msec
R3#
```

To avoid this extra-hop and permit R1 to reach 2003:3::/64 through R3 and permit R3 to reach 2001:1::/64 through R1 directly, we need to change the next-hop field by using the no ipv6 next-hop-self command as follow:

```
R2(config)#int fa0/0
R2(config-if)#no ipv6 next-hop-self eigrp 12
R2(config-if)#no ipv6 next-hop-self eigrp 23
```

Now the next-hop of the external routes are R3 and R1 for R1 and R3 respectively:

```
R1#show ipv6 route | s 2003:3::/64
EX  2003:3::/64 [170/1333760]
     via FE80::C802:2FFF:FE10:8, FastEthernet0/0
R1#
```

```
R3#show ipv6 route | s 2001:1::/64
EX  2001:1::/64 [170/1333760]
     via FE80::C800:26FF:FEE4:8, FastEthernet0/0
R3#
```

R1 goes through R3 to reach the external prefix:

```
R1#traceroute 2003:3::3
Type escape sequence to abort.
Tracing the route to 2003:3::3

  1 2000:123::3 112 msec 92 msec 60 msec
R1#
```

R2 goes through R1 to reach the external prefix:

```
R3#traceroute 2001:1::1
Type escape sequence to abort.
Tracing the route to 2001:1::1

  1 2000:123::1 100 msec 88 msec 76 msec
R3#
```

Lab 21: Filter EIGRP external routes

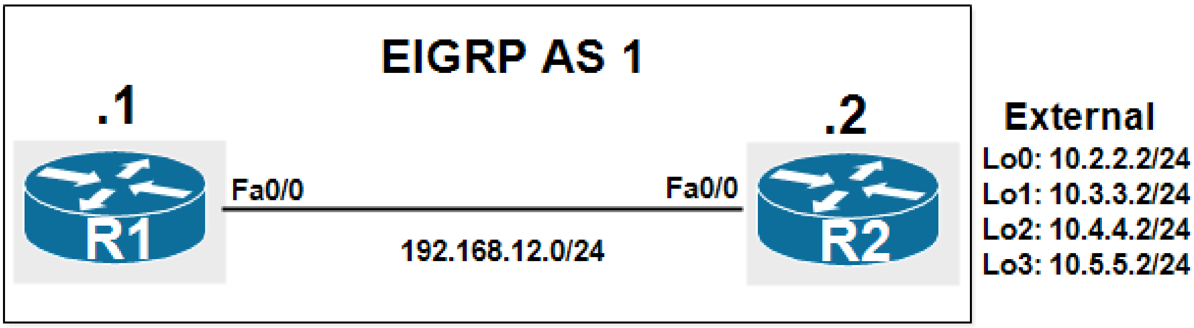

Configuration of all routers:

R1:

```
interface FastEthernet0/0
 ip address 192.168.12.1 255.255.255.0
 no shutdown
!
router eigrp 1
 network 192.168.12.0
 no auto-summary
 eigrp router-id 1.1.1.1
```

R2:

```
interface Loopback0
 ip address 10.2.2.2 255.255.255.0
!
interface Loopback1
 ip address 10.3.3.2 255.255.255.0
!
interface Loopback2
 ip address 10.4.4.2 255.255.255.0
!
interface Loopback3
 ip address 10.5.5.2 255.255.255.0
!
interface FastEthernet0/0
 ip address 192.168.12.2 255.255.255.0
 no shutdown
!
router eigrp 1
 redistribute connected metric 1500 10 255 1 1500
 network 192.168.12.0
 no auto-summary
 eigrp router-id 2.2.2.2
```

R1 and R2 are running EIGRP.
R2 redistributes the four loopback addresses into eigrp.
R1 installs successfully four external eigrp routes in its routing table.

```
R1#show 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 not set

C 192.168.12.0/24 is directly connected, FastEthernet0/0
10.0.0.0/24 is subnetted, 4 subnets
D EX 10.5.5.0 [170/1734656] via 192.168.12.2, 00:00:27, FastEthernet0/0
D EX 10.4.4.0 [170/1734656] via 192.168.12.2, 00:00:27, FastEthernet0/0
D EX 10.3.3.0 [170/1734656] via 192.168.12.2, 00:00:27, FastEthernet0/0
D EX 10.2.2.0 [170/1734656] via 192.168.12.2, 00:00:27, FastEthernet0/0

R1#
```

Requirement:prevent R1 to install these external routes with only one command

First solution:

```
R1(config)#router eigrp 1
R1(config-router)#distance eigrp 90 255
```

The distance eigrp 90 255 command sets the Administrative Distance of all EIGRP internal routes to 90 and all EIGRP external routes to 255.
From Cisco website :If the administrative distance is 255, the router does not believe the source of that route and does not install the route in the routing table.
So When an Administrative Distance is set to 255 then the router considers these routes unreachable and thus they are ignored (in other word they are not installed in the routing table).
Notice the routing table of R1 ;the four networks loopback are not installed:

```
R1(config-router)#do show 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 not set


C 192.168.12.0/24 is directly connected, FastEthernet0/0


R1(config-router)#
```

Second solution:

let's remove the distance command:

```
R1(config)#router eigrp 1
R1(config-router)#no distance eigrp 90 255
```

Notice, R1 and R2 have a different router-id,1.1.1.1 and 2.2.2.2 respectively,unlike for OSPF's router-id,with EIGRP ,duplicate Router-id do not prevent routers from becoming neighbors.
But the only problem that we can meet with duplicate router-id is when injecting external routes into EIGRP.
Let's configure R1 with the router-id 2.2.2.2 ,the same RID of R2:

```
R1(config-router)#eigrp router-id 2.2.2.2
```

Notice the routing table R1 ,the four networks loopback are not installed:

```
R1(config-router)#do show 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 not set

C 192.168.12.0/24 is directly connected, FastEthernet0/0

R1(config-router)#
```

R1 and R2 are adjacent even if they have the same RID as shown by the following output:

```
R1(config)#do show ip eigrp nei
IP-EIGRP neighbors for process 1
H               Address                 Interface               Hold
Uptime          SRTT        RTO      Q         Seq
                                                                                  (
sec)                        (ms)                     Cnt      Num
0          192.168.12.2               Fa0/0                    11
00:30:20             33             200        0           25
```

Lab 22: EIGRP Stub Feature

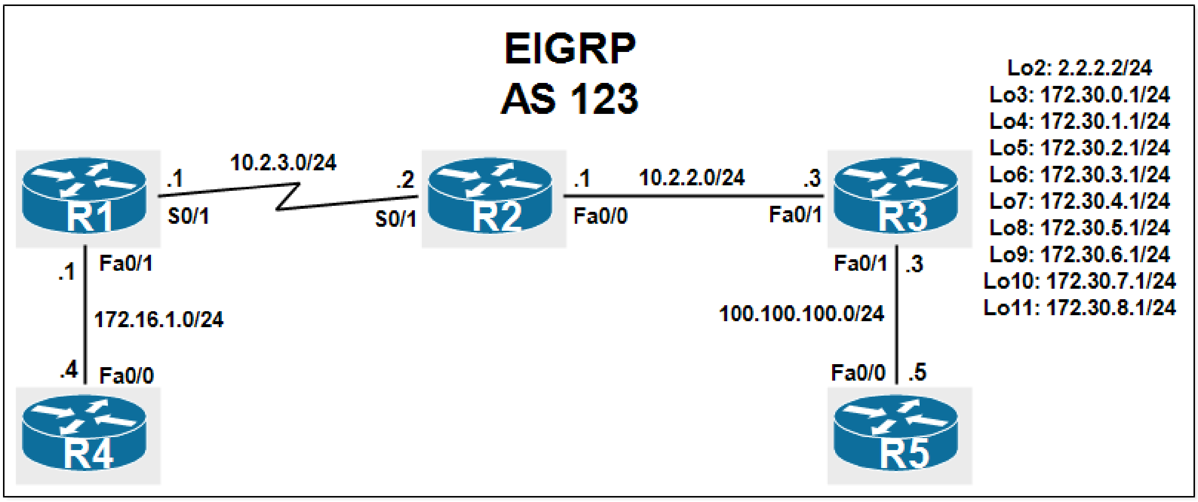

Configuration of all routers:

R1:

```
interface FastEthernet0/1
 ip address 172.16.1.1 255.255.255.0
 no shutdown
!
interface Serial0/1
 ip address 10.2.3.1 255.255.255.0
 no shutdown
!
router eigrp 123
 network 10.2.3.1 0.0.0.0
 network 172.16.1.1 0.0.0.0
 no auto-summary
```

R2:

```
interface FastEthernet0/1
 ip address 10.2.2.1 255.255.255.0
 no shutdown
!
interface Serial0/1
 ip address 10.2.3.2 255.255.255.0
 no shutdown
!
router eigrp 123
 network 10.2.2.1 0.0.0.0
 network 10.2.3.2 0.0.0.0
 no auto-summary
```

R3:

```
interface Loopback2
 ip address 2.2.2.2 255.255.255.255
!
interface Loopback3
 ip address 172.30.0.1 255.255.255.0
!
interface Loopback4
```

```
 ip address 172.30.1.1 255.255.255.0
!
interface Loopback5
 ip address 172.30.2.1 255.255.255.0
!
interface Loopback6
 ip address 172.30.3.1 255.255.255.0
!
interface Loopback7
 ip address 172.30.4.1 255.255.255.0
!
interface Loopback8
 ip address 172.30.5.1 255.255.255.0
!
interface Loopback9
 ip address 172.30.6.1 255.255.255.0
!
interface Loopback10
 ip address 172.30.7.1 255.255.255.0
!
interface Loopback11
 ip address 172.30.8.1 255.255.255.0
!
interface FastEthernet0/0
 ip address 10.2.2.3 255.255.255.0
 no shutdown
!
interface FastEthernet0/1
 ip address 100.100.100.3 255.255.255.0
 no shutdown
!
router eigrp 123
 network 2.2.2.2 0.0.0.0
 network 3.3.3.3 0.0.0.0
 network 10.2.2.3 0.0.0.0
 network 100.100.100.3 0.0.0.0
 network 172.30.0.0
 no auto-summary
```

R4:

```
interface FastEthernet0/0
 ip address 172.16.1.4 255.255.255.0
 no shutdown
!
router eigrp 123
 network 172.16.1.4 0.0.0.0
 no auto-summary
```

R5:

```
interface FastEthernet0/0
 ip address 100.100.100.5 255.255.255.0
 no shutdown
!
router eigrp 123
 network 100.100.100.5 0.0.0.0

 no auto-summary
```

R3 creates and advertises a summary route for the loopback interfaces Lo3 through Lo10 out fa0/0 and a static route toward 3.3.3.3 as follow:

```
R3(config)#int fa0/0
R3(config-if)#ip summary-address eigrp 123 172.30.0.0 255.255.248.0
```

```
R3(config)#router eigrp 123
R3(config-router)# network 3.3.3.3 0.0.0.0
R3(config)#ip route 3.3.3.3 255.255.255.255 Null0
```

R1,R2 and R4 installs successfully the summary route in the routing tables as shown by the following outputs:

```
R1#show ip route
     2.0.0.0/32 is subnetted, 1 subnets
D       2.2.2.2 [90/2323456] via 10.2.3.2, 00:01:19, Serial0/1
     100.0.0.0/24 is subnetted, 1 subnets
D       100.100.100.0 [90/2221056] via 10.2.3.2, 00:01:19, Serial0/1
     3.0.0.0/32 is subnetted, 1 subnets
D       3.3.3.3 [90/2195456] via 10.2.3.2, 00:00:50, Serial0/1
     172.16.0.0/24 is subnetted, 1 subnets
C       172.16.1.0 is directly connected, FastEthernet0/1
     172.30.0.0/16 is variably subnetted, 2 subnets, 2 masks
D       172.30.0.0/21 [90/2323456] via 10.2.3.2, 00:01:21, Serial0/1
D       172.30.8.0/24 [90/2323456] via 10.2.3.2, 00:01:21, Serial0/1
     10.0.0.0/24 is subnetted, 2 subnets
D       10.2.2.0 [90/2195456] via 10.2.3.2, 00:01:21, Serial0/1
C       10.2.3.0 is directly connected, Serial0/1
R1#
```

```
R4#show ip route
     2.0.0.0/32 is subnetted, 1 subnets
D       2.2.2.2 [90/2349056] via 172.16.1.1, 00:07:11, FastEthernet0/0
     100.0.0.0/24 is subnetted, 1 subnets
D       100.100.100.0 [90/2246656] via 172.16.1.1, 00:07:11, FastEthernet0/0
     3.0.0.0/32 is subnetted, 1 subnets
D       3.3.3.3 [90/2221056] via 172.16.1.1, 00:07:07, FastEthernet0/0
     172.16.0.0/24 is subnetted, 1 subnets
C       172.16.1.0 is directly connected, FastEthernet0/0
     172.30.0.0/16 is variably subnetted, 2 subnets, 2 masks
D       172.30.0.0/21 [90/2349056] via 172.16.1.1, 00:07:13,  FastEthernet0/0
D       172.30.8.0/24 [90/2349056] via 172.16.1.1, 00:07:13, FastEthernet0/0
     10.0.0.0/24 is subnetted, 2 subnets
D       10.2.2.0 [90/2221056] via 172.16.1.1, 00:07:13, FastEthernet0/0
D       10.2.3.0 [90/2195456] via 172.16.1.1, 00:07:13, FastEthernet0/0
R4#
```

```
R2#show ip route
     2.0.0.0/32 is subnetted, 1 subnets
D       2.2.2.2 [90/409600] via 10.2.2.3, 00:12:35, FastEthernet0/1
     100.0.0.0/24 is subnetted, 1 subnets
D       100.100.100.0 [90/307200] via 10.2.2.3, 00:12:35, FastEthernet0/1
     3.0.0.0/32 is subnetted, 1 subnets
D       3.3.3.3 [90/281600] via 10.2.2.3, 00:09:04, FastEthernet0/1
     172.16.0.0/24 is subnetted, 1 subnets
D       172.16.1.0 [90/2195456] via 10.2.3.1, 00:09:33, Serial0/1
     172.30.0.0/16 is variably subnetted, 2 subnets, 2 masks
D       172.30.0.0/21 [90/409600] via 10.2.2.3, 00:12:38, FastEthernet0/1
```

```
D       172.30.8.0/24 [90/409600] via 10.2.2.3, 00:12:36, FastEthernet0/1
     10.0.0.0/24 is subnetted, 2 subnets
C       10.2.2.0 is directly connected, FastEthernet0/1
C       10.2.3.0 is directly connected, Serial0/1
R2#
```

R5 receives all routes:

```
R5#show 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 not set

     2.0.0.0/32 is subnetted, 1 subnets
D       2.2.2.2 [90/409600] via 100.100.100.3, 00:10:00, FastEthernet0/0
     100.0.0.0/24 is subnetted, 1 subnets
C       100.100.100.0 is directly connected, FastEthernet0/0
     3.0.0.0/32 is subnetted, 1 subnets
D       3.3.3.3 [90/281600] via 100.100.100.3, 00:10:00, FastEthernet0/0
     172.16.0.0/24 is subnetted, 1 subnets
D       172.16.1.0 [90/2246656] via 100.100.100.3, 00:10:00, FastEthernet0/0
     172.30.0.0/24 is subnetted, 9 subnets
D       172.30.2.0 [90/409600] via 100.100.100.3, 00:10:01, FastEthernet0/0
D       172.30.3.0 [90/409600] via 100.100.100.3, 00:10:01, FastEthernet0/0
D       172.30.0.0 [90/409600] via 100.100.100.3, 00:10:01, FastEthernet0/0
D       172.30.1.0 [90/409600] via 100.100.100.3, 00:10:01, FastEthernet0/0
D       172.30.6.0 [90/409600] via 100.100.100.3, 00:10:01, FastEthernet0/0
D       172.30.7.0 [90/409600] via 100.100.100.3, 00:10:01, FastEthernet0/0
D       172.30.4.0 [90/409600] via 100.100.100.3, 00:10:01, FastEthernet0/0
D       172.30.5.0 [90/409600] via 100.100.100.3, 00:10:01, FastEthernet0/0
D       172.30.8.0 [90/409600] via 100.100.100.3, 00:10:01, FastEthernet0/0
     10.0.0.0/24 is subnetted, 2 subnets
D       10.2.2.0 [90/307200] via 100.100.100.3, 00:10:01, FastEthernet0/0
D       10.2.3.0 [90/2221056] via 100.100.100.3, 00:10:01, FastEthernet0/0
R5#
```

let's configure R3 with eigrp stub command:

```
R3(config)#router eigrp 123
R3(config-router)#eigrp stub
```

let's verify the routing table of R1 ,notice that the static route is missing on R1 but the summary route is still installed ,because when issued simply as eigrp stub, R3 uses default parameters, which are the connected and summary options.

```
R1#show ip route
     2.0.0.0/32 is subnetted, 1 subnets
D       2.2.2.2 [90/2323456] via 10.2.3.2, 00:00:31, Serial0/1
     100.0.0.0/24 is subnetted, 1 subnets
D       100.100.100.0 [90/2221056] via 10.2.3.2, 00:00:31, Serial0/1
     172.16.0.0/24 is subnetted, 1 subnets
C       172.16.1.0 is directly connected, FastEthernet0/1
     172.30.0.0/16 is variably subnetted, 2 subnets, 2 masks
```

```
D        172.30.0.0/21 [90/2323456] via 10.2.3.2, 00:00:31, Serial0/1
D        172.30.8.0/24 [90/2323456] via 10.2.3.2, 00:00:31, Serial0/1
     10.0.0.0/24 is subnetted, 2 subnets
D        10.2.2.0 [90/2195456] via 10.2.3.2, 00:15:35, Serial0/1
C        10.2.3.0 is directly connected, Serial0/1
R1#
```

Notice that we can ping 100.100.100.3 the ip add of fa0/1's R3 but not 100.100.100.5 the ip add of fa0/0's R5:

```
R1#ping 100.100.100.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 100.100.100.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 20/25/36 ms
R1#
```

```
R1#ping 100.100.100.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 100.100.100.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
R1#
```

Because simply R5 is located behind the stub thus it does not have all informations about the routes located beyond the sub as shown by the show ip route on R5:

```
R5#show ip route
     2.0.0.0/32 is subnetted, 1 subnets
D        2.2.2.2 [90/409600] via 100.100.100.3, 00:02:01, FastEthernet0/0
     100.0.0.0/24 is subnetted, 1 subnets
C        100.100.100.0 is directly connected, FastEthernet0/0
     172.30.0.0/24 is subnetted, 9 subnets
D        172.30.2.0 [90/409600] via 100.100.100.3, 00:02:01, FastEthernet0/0
D        172.30.3.0 [90/409600] via 100.100.100.3, 00:02:01, FastEthernet0/0
D        172.30.0.0 [90/409600] via 100.100.100.3, 00:02:01, FastEthernet0/0
D        172.30.1.0 [90/409600] via 100.100.100.3, 00:02:03, FastEthernet0/0
D        172.30.6.0 [90/409600] via 100.100.100.3, 00:02:03, FastEthernet0/0
D        172.30.7.0 [90/409600] via 100.100.100.3, 00:02:03, FastEthernet0/0
D        172.30.4.0 [90/409600] via 100.100.100.3, 00:02:03, FastEthernet0/0
D        172.30.5.0 [90/409600] via 100.100.100.3, 00:02:03, FastEthernet0/0
D        172.30.8.0 [90/409600] via 100.100.100.3, 00:02:03, FastEthernet0/0
     10.0.0.0/24 is subnetted, 1 subnets
D        10.2.2.0 [90/307200] via 100.100.100.3, 00:02:03, FastEthernet0/0
R5#
```

Let's try the eigrp stub static connected command:

```
R3(config)#router eigrp 123
R3(config-router)#eigrp stub static connected
```

We will take a look at the routing table of R1,notice that the static route appears because after adding the keyword static R3 includes the static route in its update, notice that R3 advertises all the connected routes (the loopback interfaces) instead of the summary route configured above:

```
R1#show ip route

      2.0.0.0/32 is subnetted, 1 subnets
D        2.2.2.2 [90/2323456] via 10.2.3.2, 00:00:48, Serial0/1
      100.0.0.0/24 is subnetted, 1 subnets
D        100.100.100.0 [90/2221056] via 10.2.3.2, 00:00:48, Serial0/1
      3.0.0.0/32 is subnetted, 1 subnets
D        3.3.3.3 [90/2195456] via 10.2.3.2, 00:00:48, Serial0/1
      172.16.0.0/24 is subnetted, 1 subnets
C        172.16.1.0 is directly connected, FastEthernet0/1
      172.30.0.0/24 is subnetted, 9 subnets
D        172.30.2.0 [90/2323456] via 10.2.3.2, 00:00:49, Serial0/1
D        172.30.3.0 [90/2323456] via 10.2.3.2, 00:00:49, Serial0/1
D        172.30.0.0 [90/2323456] via 10.2.3.2, 00:00:49, Serial0/1
D        172.30.1.0 [90/2323456] via 10.2.3.2, 00:00:49, Serial0/1
D        172.30.6.0 [90/2323456] via 10.2.3.2, 00:00:49, Serial0/1
D        172.30.7.0 [90/2323456] via 10.2.3.2, 00:00:49, Serial0/1
D        172.30.4.0 [90/2323456] via 10.2.3.2, 00:00:49, Serial0/1
D        172.30.5.0 [90/2323456] via 10.2.3.2, 00:00:49, Serial0/1
D        172.30.8.0 [90/2323456] via 10.2.3.2, 00:00:49, Serial0/1
      10.0.0.0/24 is subnetted, 2 subnets
D        10.2.2.0 [90/2195456] via 10.2.3.2, 00:25:49, Serial0/1
C        10.2.3.0 is directly connected, Serial0/1
R1#
```

Let's try eigrp stub receive-only command:

```
R3(config)#router eigrp 123
R3(config-router)#eigrp stub receive-only
```

R1 now has only routes connected and the eigrp route toward 10.2.2.0 because with the keyword receive-only ,R3 does not advertise any routes:

```
R1#show ip route

      172.16.0.0/24 is subnetted, 1 subnets
C        172.16.1.0 is directly connected, FastEthernet0/1
      10.0.0.0/24 is subnetted, 2 subnets
D        10.2.2.0 [90/2195456] via 10.2.3.2, 00:33:36, Serial0/1
C        10.2.3.0 is directly connected, Serial0/1
R1#
```