

Chapter 2: EIGRP Concept



CCNP ROUTE: Implementing IP Routing

Cisco | Networking Academy®
Mind Wide Open™



Chapter 2 Objectives

- Establishing EIGRP Neighbor Relationships
- Building the EIGRP Topology Table
- Optimizing EIGRP Behavior
- Configuring EIGRP for IPv6
- Named EIGRP Configuration

Establishing EIGRP Neighbor Relationships





Enhanced Interior Gateway Routing Protocol (EIGRP)

- Enhanced Interior Gateway Routing Protocol (EIGRP) is an advanced **distance vector** routing protocol designed by Cisco.
- The basic configuration is **simple** and **easy to understand**, so it is commonly used in **smaller networks**.
- Its advanced features, which provide **rapid convergence**, **higher scalability**, and support for **multiple routed protocols**, fulfill requirements in complex network environments.
- EIGRP supports both **IPv4** and **IPv6**.
- Although standard EIGRP configuration between IPv4 and IPv6 differs, it can be unified using newly introduced **named EIGRP configuration mode**.



Establishing EIGRP Neighbor Relationships

- EIGRP Characteristics
- EIGRP Reliable Transport
- EIGRP Operation Overview
- Configuring and Verifying Basic EIGRP for IPv4
- EIGRP Timers
- EIGRP Adjacencies in a Frame Relay Network
- EIGRP Adjacencies in a Layer 3 MPLS VPN Network
- EIGRP Adjacencies in a Layer 2 MPLS VPN Ethernet Network



EIGRP Characteristics

EIGRP provides some unique characteristics compared to other routing protocols:

- **Fast convergence**
- **Partial updates**
- **Multiple network layer support**
- **Use of multicast and unicast**
- **VLSM support**
- **Sophisticated metric**
- **Reliable transport protocol**



EIGRP Characteristics

Fast convergence

- EIGRP uses the **Diffusing Update ALgorithm (DUAL)** to achieve **rapid convergence**. A router running EIGRP stores its neighbors' routing tables so that it can quickly adapt to changes in the network.
- If no appropriate route exists in the local routing table and no appropriate backup route exists in the topology table, EIGRP **queries its neighbors** to discover an alternative route. These queries are propagated until an alternative route is found or until it is determined that no alternative route exists.



EIGRP Characteristics

Partial updates

- EIGRP sends **partial triggered updates** rather than periodic updates. These updates are sent only when the path or the metric for a route changes.
- They contain information about only that changed link **rather than the entire routing table**.
- Propagation of these partial updates is automatically bounded so that **only those routers that require the information are updated**. As a result, EIGRP consumes significantly **less bandwidth** than IGRP.
- This behavior also differs from link-state protocol operation, which sends a change update to all routers within an area.



EIGRP Characteristics

Multiple network layer support

- EIGRP supports IP Version 4 (IPv4) and IP Version 6 (IPv6) using **protocol-dependent modules** that are responsible for protocol requirements specific to the network layer. EIGRP's rapid convergence and sophisticated metric offer superior performance and stability when implemented in IPv4 and IPv6 networks.

Use of multicast and unicast

- For communication between routers, EIGRP uses multicast and unicast rather than broadcast. As a result, end stations are unaffected by routing updates or queries. The multicast address used for EIGRP for **IPv4 is 224.0.0.10**, and the multicast address for EIGRP for **IPv6 is FF00::A**.



EIGRP Characteristics

VLSM support

- EIGRP is a **classless routing protocol**, which means that it **advertises a subnet mask** for each destination network. This enables EIGRP to support discontinuous subnetworks and VLSM.

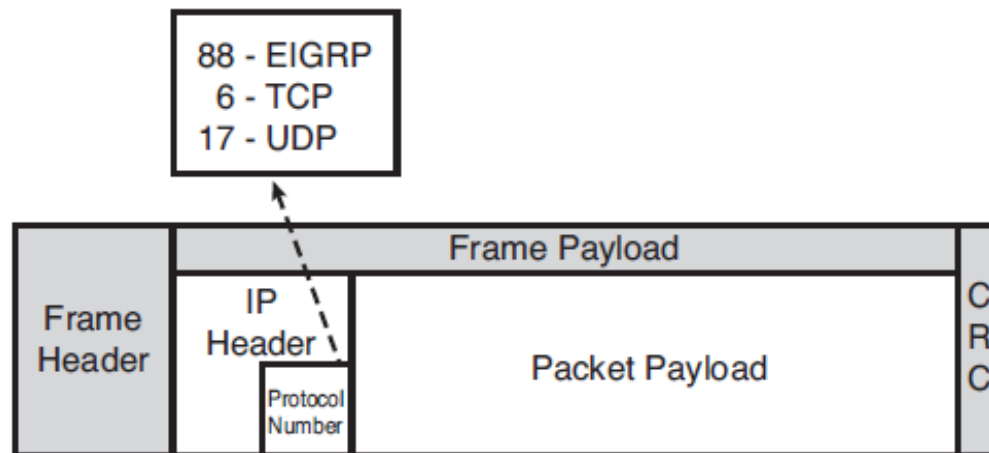
Sophisticated metric

- EIGRP represents metric values in a 32-bit format to provide enough granularity. EIGRP supports **unequal metric load balancing**, which allows administrators to distribute traffic flow more efficiently in their networks.



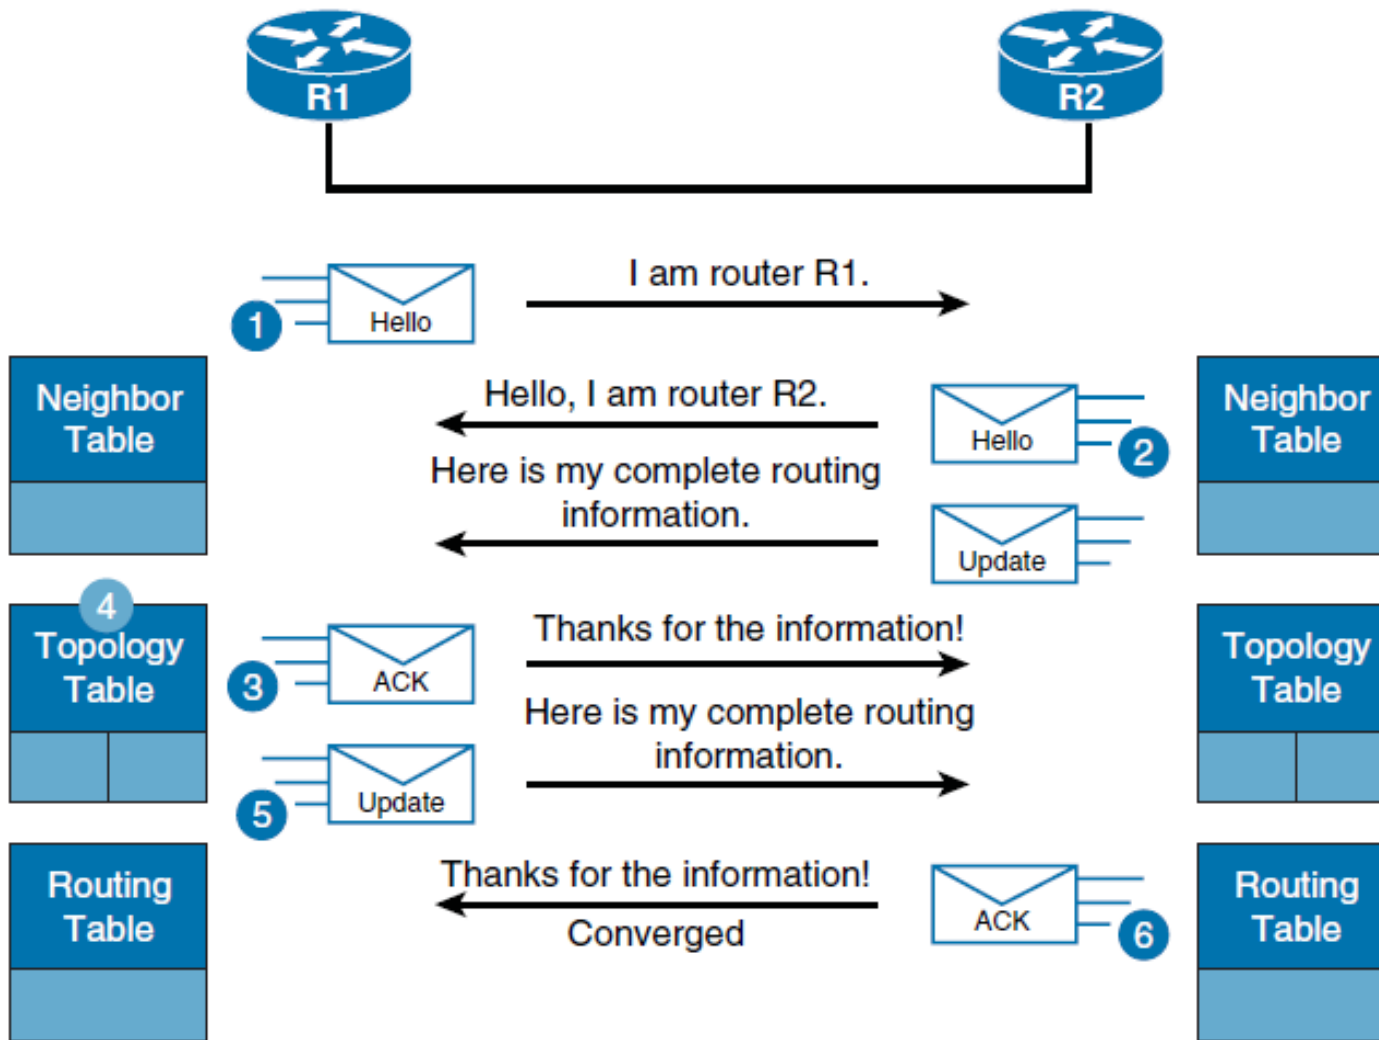
EIGRP Reliable Transport

- EIGRP runs directly above the IP layer as its own protocol, **numbered 88**.
- RTP is the component of the EIGRP responsible for **guaranteed, ordered delivery** of EIGRP packets to all neighbors.
- It supports intermixed transmission of multicast or unicast packets.





EIGRP Operation Overview – in 6 steps





EIGRP Operation Overview – in 6 steps

1. A new router (router R1 in this example) comes up on the link and sends a **Hello packet** through all its EIGRP-configured interfaces.
2. Routers that are receiving the Hello packet (R2) on one interface **reply with update packets that contain all the routes that they have in their routing tables**, except those that are learned through that interface (split horizon). R2 sends an update packet to R1, but a neighbor relationship is not established until R2 sends a hello packet to R1. The update packet from R2 has the **initialization bit set**, indicating that this is the initialization process. The update packet includes information about the routes that the neighbor (R2) is aware of, including the metric that the neighbor is advertising for each destination.




EIGRP Operation Overview – in 6 steps

3. After both routers have exchanged Hellos and the neighbor adjacency is established, R1 replies to R2 with an **ACK packet**, indicating that it received the update information.
4. R1 assimilates all the update packets in its topology table. The topology table includes all destinations that are advertised by neighboring adjacent routers. It lists each destination, all the neighbors that can reach the destination, and their associated metric.
5. R1 sends an update packet to R2.
6. Upon receiving the update packet, R2 sends an ACK packet to R1




Configuring and Verifying Basic EIGRP for IPv4



```
BR1> enable
BR1# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
BR1(config)# router eigrp 100
BR1(config-router)# network 172.16.1.0
BR1(config-router)# network 192.168.1.0
```



Configuring and Verifying Basic EIGRP for IPv4



```
BR2(config)# router eigrp 100
BR2(config-router)# network 172.16.2.0 0.0.0.3
BR2(config-router)# network 192.168.2.1 0.0.0.0
```

```
BR3(config)# router eigrp 100
BR3(config-router)# eigrp router-id 192.168.3.255
BR3(config-router)# network 0.0.0.0 255.255.255.255
```




Verifying EIGRP Neighbor Relationships

```
BR1# show ip eigrp neighbors
```

```
EIGRP-IPv4 Neighbors for AS(100)
```

H	Address	Interface	Hold (sec)	Uptime (ms)	SRTT	RTO	Q Cnt	Seq Num
0	172.16.1.1	Se0/0	13	01:29:20	17	102	0	11

```
BR1# show ip eigrp neighbors detail
```

```
EIGRP-IPv4 Neighbors for AS(100)
```

H	Address	Interface	Hold (sec)	Uptime (ms)	SRTT	RTO	Q Cnt	Seq Num
0	172.16.1.1	Se0/0	14	01:40:47	17	102	0	11

Version 7.0/3.0, Retrans: 0, Retries: 0, Prefixes: 5

Topology-ids from peer - 0



Verifying the EIGRP Interfaces

```
BR1# show ip eigrp interfaces
```

```
EIGRP-IPv4 Interfaces for AS(100)
```

Interface	Peers	Xmit Queue Un/Reliable	PeerQ Un/Reliable	Mean SRTT	Pacing Time Un/Reliable	Multicast Flow Timer	Pending Routes
Se0/0	1	0/0	0/0	17	0/16	88	0
Lo0	0	0/0	0/0	0	0/0	0	0

```
BR1# show ip eigrp interfaces detail
```

```
EIGRP-IPv4 Interfaces for AS(100)
```

Interface	Peers	Xmit Queue Un/Reliable	PeerQ Un/Reliable	Mean SRTT	Pacing Time Un/Reliable	Multicast Flow Timer	Pending Routes
Se0/0	1	0/0	0/0	17	0/16	88	0

```
Hello-interval is 5, Hold-time is 15
```

```
Split-horizon is enabled
```

```
Next xmit serial <none>
```

```
Packetized sent/expedited: 5/1
```

```
Hello's sent/expedited: 16200/2
```

```
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
```

```
<Output omitted>
```



Verifying the EIGRP Networks

```

HQ# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "eigrp 100"
  <Output omitted>
  Routing for Networks:
    0.0.0.0
  Routing Information Sources:
    Gateway          Distance      Last Update
    172.16.2.2       90           23:04:13
    172.16.3.2       90           23:04:13
    172.16.1.2       90           23:04:13
  Distance: internal 90 external 170
  
```



Configuring Passive Interfaces

```
HQ(config)# router eigrp 100
HQ(config-router)# passive-interface default
*Sep 24 03:27:31.719: %DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 172.16.3.2
  (Serial2/0) is down: interface passive
*Sep 24 03:27:31.719: %DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 172.16.1.2
  (Serial1/0) is down: interface passive
*Sep 24 03:27:31.720: %DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 172.16.2.2
  (Ethernet0/0) is down: interface passive
```

```
HQ(config-router)# no passive-interface ethernet 0/0
*Sep 24 03:31:16.376: %DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 172.16.2.2
  (Ethernet0/0) is up: new adjacency
HQ(config-router)# no passive-interface serial 1/0
*Sep 24 03:31:42.184: %DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 172.16.1.2
  (Serial1/0) is up: new adjacency
HQ(config-router)# no passive-interface serial 2/0
*Sep 24 03:31:56.265: %DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 172.16.3.2
  (Serial2/0) is up: new adjacency
```



Verifying Hello and Hold Timers

```
HQ# show ip eigrp interfaces detail
```

```
EIGRP-IPv4 Interfaces for AS(100)
```

Interface	Xmit Peers	Queue Un/Reliable	PeerQ Un/Reliable	Mean SRTT	Pacing Time Un/Reliable	Multicast Flow Timer	Pending Routes
Et0/0	1	0/0	0/0	5	0/2	50	0
Hello-interval is 5, Hold-time is 15							
Se1/0	1	0/0	0/0	13	0/15	71	0
Hello-interval is 5, Hold-time is 15							
Se2/0	1	0/0	0/0	1008	10/400	4432	0
Hello-interval is 60, Hold-time is 180							
<Output omitted>							



EIGRP Timers

- EIGRP determines default timer values based on link type. If default values are not suitable for a specific network topology, you can manipulate values to improve convergence time.
- EIGRP hello and hold timers between neighbors do not need to be identical to successfully establish EIGRP neighbor relationship; however, asymmetrical timers may lead to flapping EIGRP neighbor relationships and network instability.

```
BR3(config)# interface serial 0/0
BR3(config-if)# ip hello-interval eigrp 100 10
BR3(config-if)# ip hold-time eigrp 100 30
BR3# show ip eigrp interface detail serial 0/0
EIGRP-IPv4 Interfaces for AS(100)

```

Interface	Xmit Peers	Queue Un/Reliable	PeerQ Un/Reliable	Mean SRTT	Pacing Time Un/Reliable	Multicast Flow Timer	Pending Routes
Se0/0	1	0/0	0/0	1268	0/16	6340	0

```

Hello-interval is 10, Hold-time is 30
<Output omitted>

```




EIGRP Adjacencies in a Frame Relay Network

- When configuring EIGRP over **point-to-multipoint** subinterfaces, a single IP subnet is used.
- To emulate broadcast multiaccess network and enable EIGRP to send multicast packets over Frame Relay virtual circuits (VCs), you must add the broadcast keyword in the Frame Relay static mapping statement
 - `frame-relay map ip ip-address dlci broadcast`
- When configuring EIGRP over **point-to-point** subinterfaces, a different IP subnet is used for each subinterface.
 - These are logical interfaces that are emulating a leased line network and are a routing equivalent to point-to-point physical interfaces




EIGRP Adjacencies in a Layer 3 MPLS VPN Network



- Provider-edge (PE) routers participate in customer routing, guaranteeing optimum routing between customer sites.
- PE routers carry a separate set of routes for each customer, resulting in perfect isolation between the customers.



EIGRP Adjacencies in a Layer 2 MPLS VPN Ethernet Network



- Customer routers are located within single metropolitan area and they may be connected over the local Layer 2 MPLS VPN switch network. Customer traffic never passes through the SP backbone.
- Customer routers are located between several geographically distant areas that need to be connected over L2 MPLS VPN with point-to-point links through the SP backbone.
- Customer routers are located between several geographically distant areas that need to be connected over L2 MPLS VPN with multipoint links through the SP core. From the customer perspective SP network looks like a LAN switch.

Building the EIGRP Topology Table



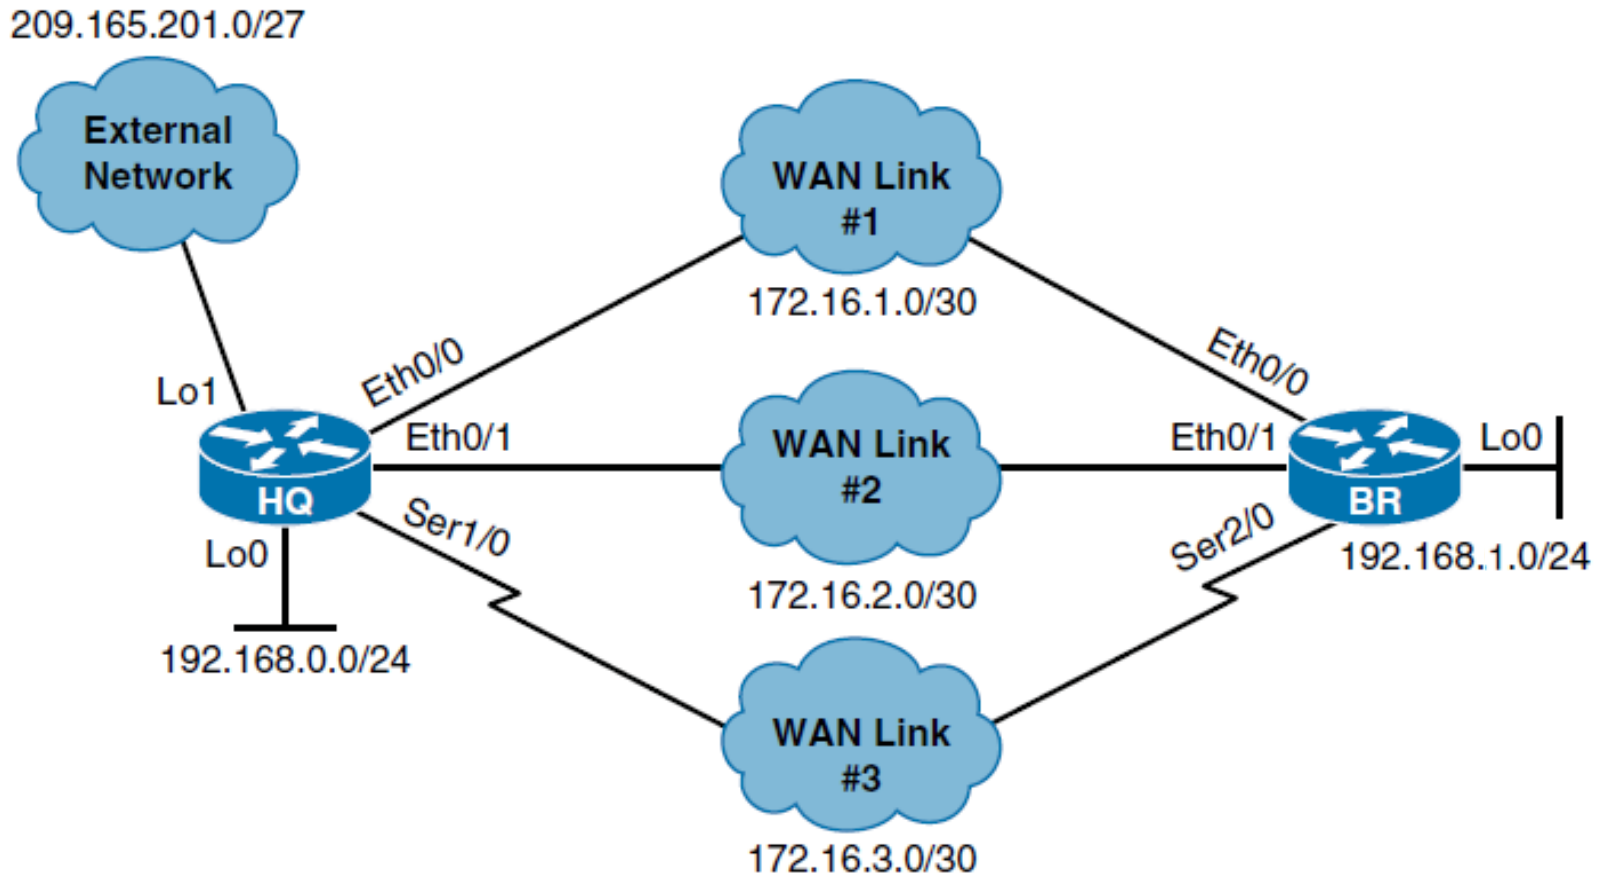


Building the EIGRP Topology Table

- EIGRP Neighbors Routing Information Exchange
- How EIGRP Chooses the Best Path through the Network
- Calculate EIGRP Metric
- Feasibility Condition prevents loops in EIGRP Networks
- Understand EIGRP Path Selection Process



EIGRP Neighbors Routing Information Exchange



EIGRP Neighbors Routing Information Exchange

```
BR# debug eigrp packet
```

```
(UPDATE, REQUEST, QUERY, REPLY, HELLO, IPXSAP, PROBE, ACK, STUB, SIAQUERY, SIAREPLY)
```

```
EIGRP Packet debugging is on
```

```
BR(config)# router eigrp 100
BR(config-router)# network 192.168.1.0 0.0.0.255
BR(config-router)# network 172.16.1.0 0.0.0.3
*Oct 8 15:20:19.227: EIGRP: Sending HELLO on Ethernet0/0
*Oct 8 15:20:19.227: AS 100, Flags 0x0:(NULL), Seq 0/0 interfaceQ 0/0 iidbQ un/
rely 0/0
*Oct 8 15:20:19.235: EIGRP: Received HELLO on Ethernet0/0 nbr 172.16.1.1
*Oct 8 15:20:19.235: AS 100, Flags 0x0:(NULL), Seq 0/0 interfaceQ 0/0
*Oct 8 15:20:19.235: %DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 172.16.1.1
(Ethernet0/0) is up: new adjacency

*Oct 8 15:20:19.261: EIGRP: Enqueueing UPDATE on Ethernet0/0 tid 0 iidbQ un/rely
0/1 serno 1-2
*Oct 8 15:20:19.266: EIGRP: Sending UPDATE on Ethernet0/0 tid 0
*Oct 8 15:20:19.266: AS 100, Flags 0x0:(NULL), Seq 2/0 interfaceQ 0/0 iidbQ un/
rely 0/0 serno 1-2
*Oct 8 15:20:19.274: EIGRP: Received ACK on Ethernet0/0 nbr 172.16.1.1
*Oct 8 15:20:19.275: AS 100, Flags 0x0:(NULL), Seq 0/2 interfaceQ 0/0 iidbQ un/
rely 0/0 peerQ un/rely 0/1
<Output omitted>
*Oct 8 15:20:19.253: EIGRP: Received UPDATE on Ethernet0/0 nbr 172.16.1.1
*Oct 8 15:20:19.253: AS 100, Flags 0x0:(NULL), Seq 2/0 interfaceQ 0/0 iidbQ un/
rely 0/0 peerQ un/rely 0/1
*Oct 8 15:20:19.360: EIGRP: Enqueueing ACK on Ethernet0/0 nbr 172.16.1.1 tid 0
*Oct 8 15:20:19.360: Ack seq 2 iidbQ un/rely 0/0 peerQ un/rely 1/0
*Oct 8 15:20:19.364: EIGRP: Sending ACK on Ethernet0/0 nbr 172.16.1.1 tid 0
*Oct 8 15:20:19.364: AS 100, Flags 0x0:(NULL), Seq 0/2 interfaceQ 0/0 iidbQ un/
rely 0/0 peerQ un/rely 1/0
Branch# no debug all
```



Verifying EIGRP Packet Traffic

```
BR# show ip eigrp traffic
EIGRP-IPv4 Traffic Statistics for AS(100)
  Hellos sent/received: 65/67
  Updates sent/received: 9/7
  Queries sent/received: 0/0
  Replies sent/received: 0/0
  Acks sent/received: 5/5
  SIA-Queries sent/received: 0/0
  SIA-Replies sent/received: 0/0
  Hello Process ID: 101
  PDM Process ID: 63
  Socket Queue: 0/10000/2/0 (current/max/highest/drops)
  Input Queue: 0/2000/2/0 (current/max/highest/drops)
```



Verifying the EIGRP Routes

```
BR# 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

D      192.168.0.0/24 [90/409600] via 172.16.1.1, 18:20:16, Ethernet0/0
D EX   209.165.201.0/27 [170/537600] via 172.16.1.1, 18:20:16, Ethernet0/0
```



Choosing the Best Path (DUAL)

```
BR# show ip eigrp topology
EIGRP-IPv4 Topology Table for AS(100)/ID(192.168.1.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 172.16.2.0/30, 1 successors, FD is 1536000
    via Connected, Ethernet0/1
P 192.168.0.0/24, 1 successors, FD is 409600
    via 172.16.1.1 (409600/128256), Ethernet0/0
    via 172.16.2.1 (1664000/128256), Ethernet0/1
P 192.168.1.0/24, 1 successors, FD is 128256
    via Connected, Loopback0
P 172.16.3.0/30, 1 successors, FD is 2169856
    via Connected, Serial2/0
P 172.16.1.0/30, 1 successors, FD is 281600
    via Connected, Ethernet0/0
P 209.165.201.0/27, 1 successors, FD is 537600
    via 172.16.1.1 (537600/426496), Ethernet0/0
    via 172.16.2.1 (1792000/426496), Ethernet0/1
```




EIGRP Topology Table

```

BR# show ip eigrp topology all-links
EIGRP-IPv4 Topology Table for AS(100)/ID(192.168.1.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 172.16.2.0/30, 1 successors, FD is 1536000, serno 4
   via Connected, Ethernet0/1
   via 172.16.3.1 (3449856/1536000), Serial2/0
   via 172.16.1.1 (1561600/1536000), Ethernet0/0
P 192.168.0.0/24, 1 successors, FD is 409600, serno 2
   via 172.16.1.1 (409600/128256), Ethernet0/0
   via 172.16.3.1 (256512000/256000000), Serial2/0
   via 172.16.2.1 (1664000/128256), Ethernet0/1
P 192.168.1.0/24, 1 successors, FD is 128256, serno 6
   via Connected, Loopback0
P 172.16.3.0/30, 1 successors, FD is 2169856, serno 5
   via Connected, Serial2/0
   via 172.16.1.1 (2195456/2169856), Ethernet0/0
   via 172.16.2.1 (3449856/2169856), Ethernet0/1
P 172.16.1.0/30, 1 successors, FD is 281600, serno 1
   via Connected, Ethernet0/0
   via 172.16.3.1 (2195456/281600), Serial2/0
   via 172.16.2.1 (1561600/281600), Ethernet0/1

```



EIGRP Routing Information

- Only the routes that are used by EIGRP, **the successor routes, get advertised.**
- Subnets of directly connected interfaces on which EIGRP has been enabled using the **network command**
- Subnets learned by **redistribution** of routes into EIGRP from other routing protocols or routing information sources
- Redistribution is a method of taking routing information from one source and advertising it into another routing protocol.
- Redistribution is used in situations when multiple routing protocols are used in the same autonomous system, or when you want to include already-defined static routes into the selected routing protocol.



EIGRP Metric

EIGRP uses a composite metric to determine the best path to the destination. The metric's value derives from a formula that can use the following parameters:

- **Bandwidth:** **Least value** of the bandwidth for all links between the local router and the destination.
- **Delay:** **Cumulative delay obtained as sum** of values of all delays for all links between the source and destination.
- **Reliability:** This value represents **the worst reliability** between source and destination (based on keepalives).
- **Load:** This value **represents the worst load** on the link between the source and the destination (based on the packet rate and the configured bandwidth of the interface).




EIGRP Metric Calculation

$$\text{Metric} = [(K1 * \text{Bandwidth} + [(K2 * \text{Bandwidth}) / (256 - \text{Load})] + K3 * \text{Delay}) * K5 / (K4 + \text{Reliability})] * 256$$

- With the default K-Values K1=1, K2=0, K3=1, K4=0, K5=0
- $\text{Metric} = (\text{Bandwidth} + \text{Delay}) * 256$
- Note that changing the K values is not recommended.



EIGRP Metric Calculation Example




- Bandwidth = $(10^7 / \text{Least bandwidth in kilobits per second})$
- Delay = $\text{microseconds} / 10$
- Metric = $(\text{Bandwidth} + \text{Delay}) * 256$

- **Bandwidth = $(10\ 000\ 000 / 10\ 000) = 1\ 000$**
- **Delay = $[4000 + 1000 + 5000] = 10\ 000$ [tens of microseconds]**
- **Metric = $(1000 + 10\ 000) * 256 = 2\ 816\ 000$**



EIGRP Successor and FS Example

How EIGRP domain stays loop free? ==> **DUAL**

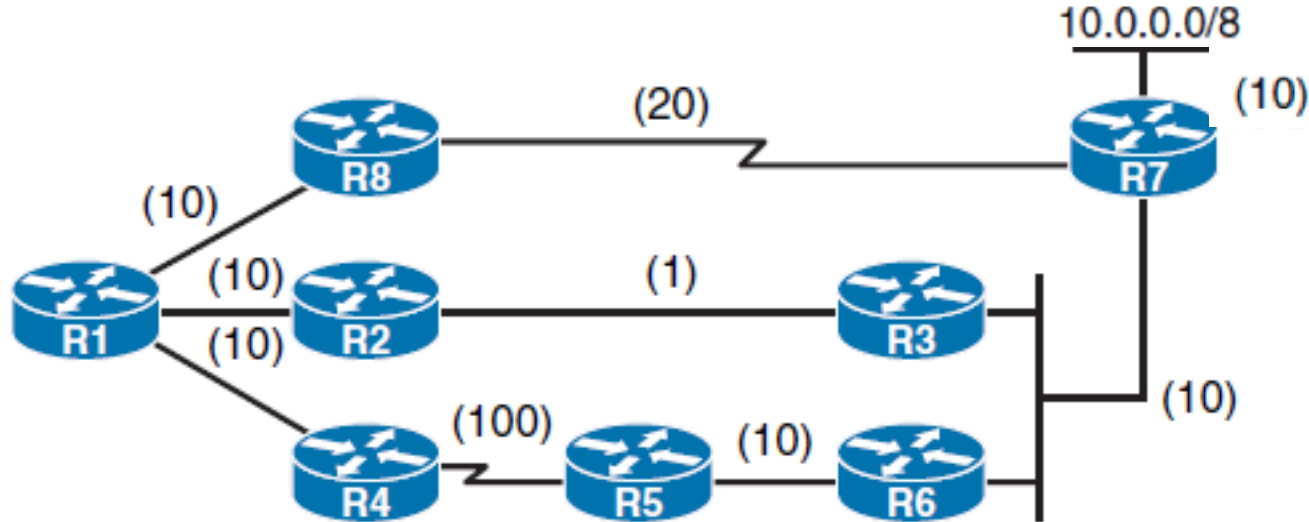


EIGRP Feasibility Condition: $FD > RD$

~~Router A: $13 > 15$~~



EIGRP Path Calculation Example




- R1 analyze the Reported Distances of neighbors to 10.0.0.0

Destination	RD	Neighbor
10.0.0.0/8	$20+10=30$	R8
10.0.0.0/8	$1+10+10=21$	R2
10.0.0.0/8	$100+10+10+10=130$	R4



EIGRP Path Calculation Example

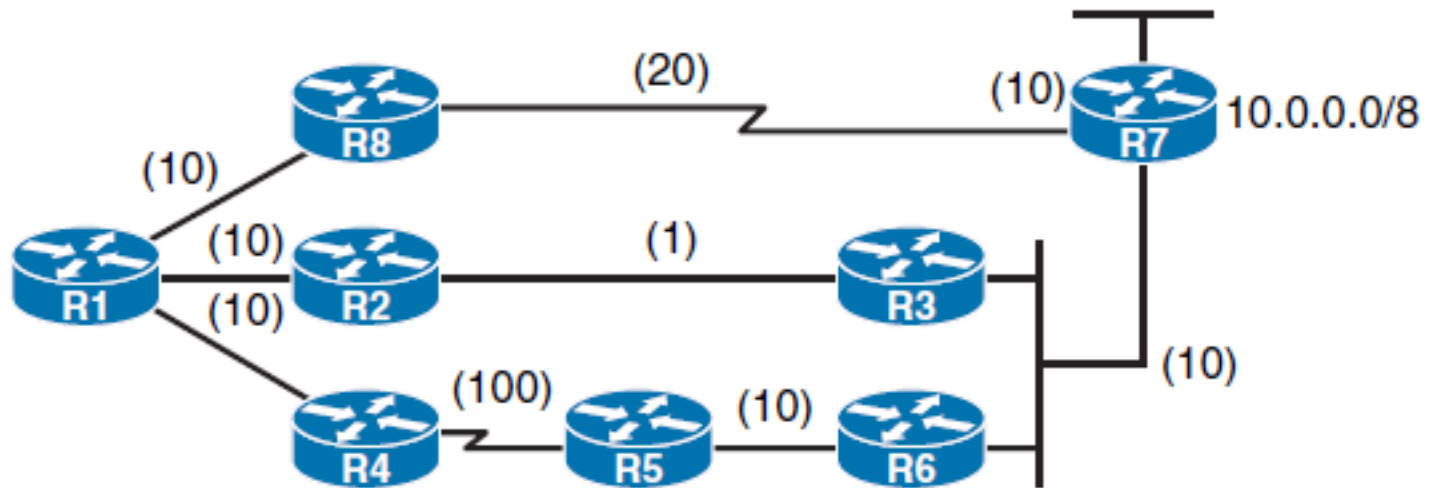


- R1's Feasible Distances to destination network 10.0.0.0/8

Destination	RD	Metric	Neighbor
10.0.0.0/8	30	$30+10=40$	R8
10.0.0.0/8	21	$21+10=31$ (FD)	R2
10.0.0.0/8	130	$130+10=140$	R4



EIGRP Path Calculation Example



- Successor
- Feasible Successor

Destination	RD	Metric	Neighbor	Status
10.0.0.0/8	30	40	R8	FS
10.0.0.0/8	21	31 (FD)	R2	S
10.0.0.0/8	130	140	R4	Non-S

Optimizing EIGRP Behavior



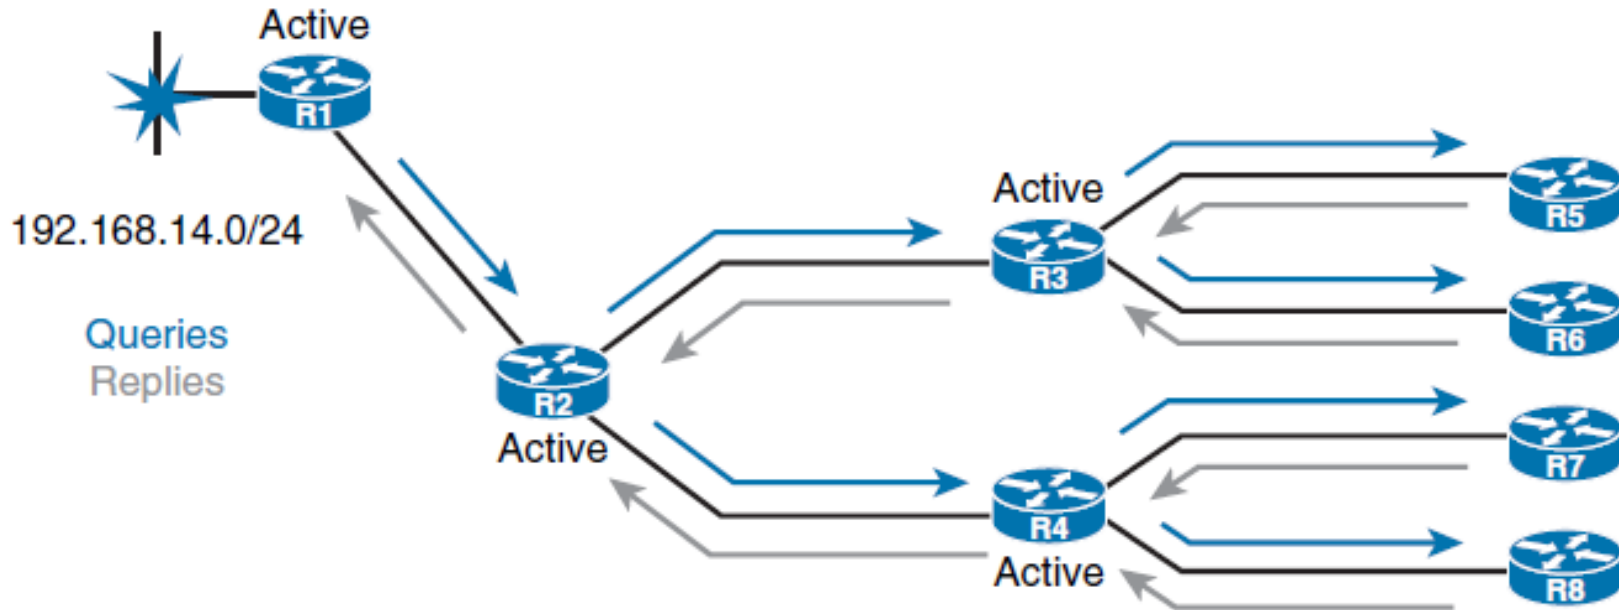


Optimizing EIGRP Behavior

- EIGRP **queries**
- Describe how **stub routing** can be used to reduce the **amount of queries** when EIGRP goes active
- EIGRP **Stuck-in-Active (SiA)** issue
- Explain how using **summary routes** lessen the impact of query scope when EIGRP goes active
- Describe load-balancing options with EIGRP



EIGRP Queries



- When a router loses a route and **does not have a feasible successor** in its topology table, **it looks for an alternative path to the destination**.
- This is known as *going active* on a route.



Active Route

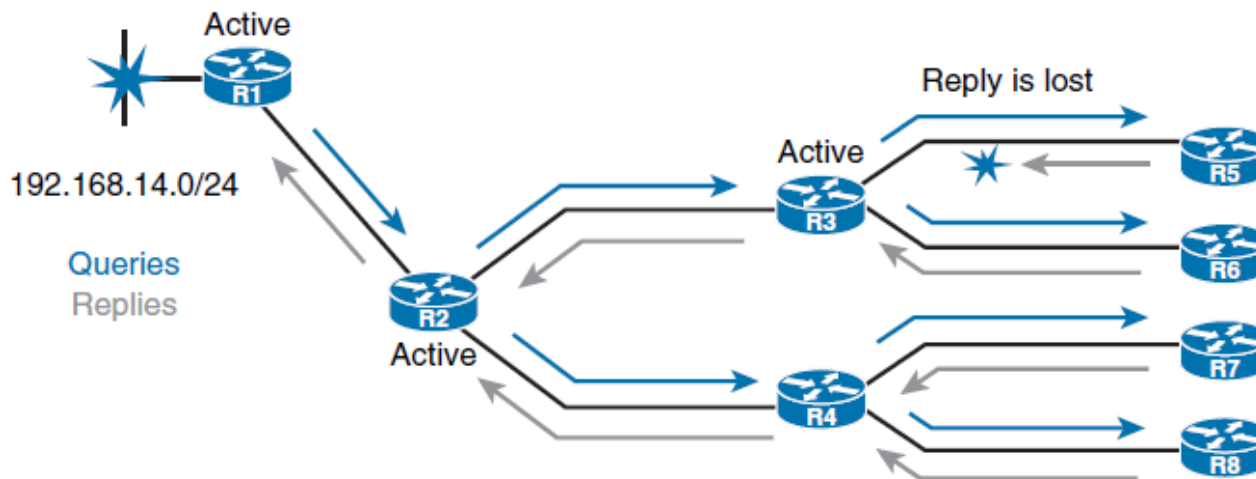
```

R1# show ip eigrp topology
EIGRP-IPv4 Topology Table for AS(1)/ID(172.16.1.2)
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 281600
   via Connected, GigabitEthernet0/1
A 192.168.14.0/24, 0 successors, FD is 409600, Q
   1 replies, active 00:00:02, query-origin: Local origin
   Remaining replies:
     via 172.16.1.1, r, GigabitEthernet0/0
P 172.16.1.0/30, 1 successors, FD is 281600
   via Connected, GigabitEthernet0/0
    
```



Stuck in Active

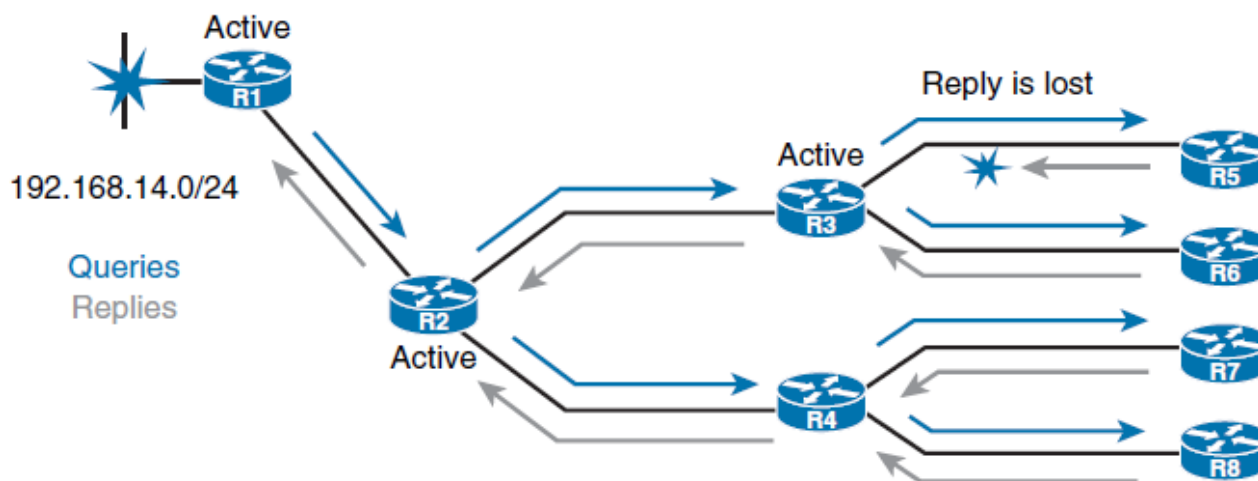
- Once a route goes active and the query sequence is initiated, it can only come out of the active state and transition to the passive state when it receives a reply for every generated query.
- If the router does not receive a reply to all the outstanding queries within 3 minutes (the default time), the route goes into the stuck-in-active (SIA) state.





Stuck in Active

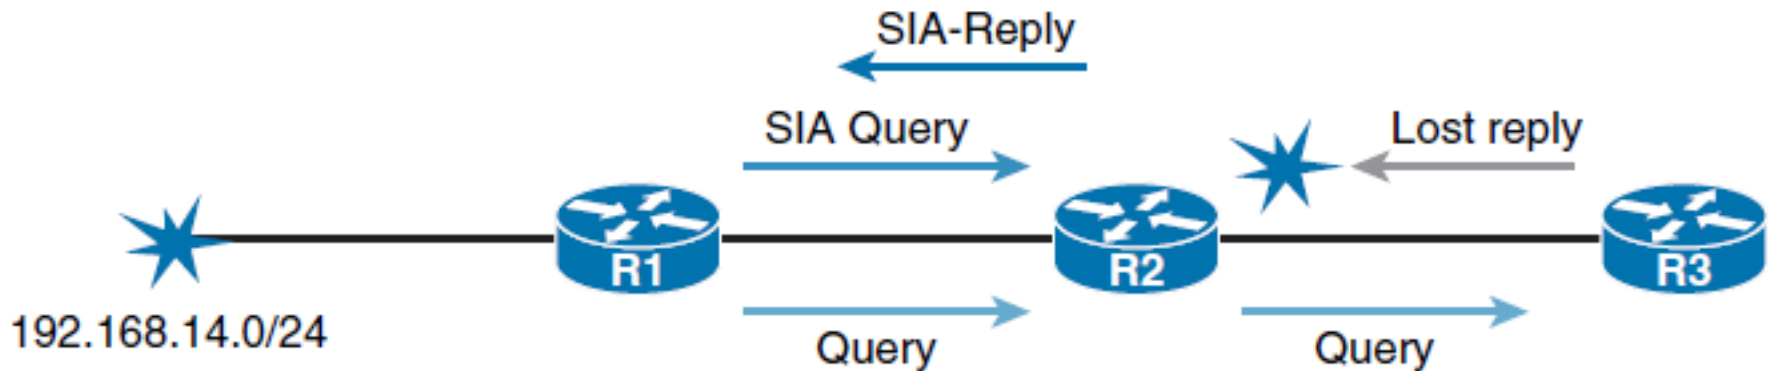
- This timer is called the **active timer**.
- Once the active timer expires, the **neighbor relationship is reset**.
- This setting causes the router to go active on **all routes that were known through the lost neighbor** and to re-advertise all the routes that it knows to the lost neighbor.





Stuck in Active – overcome

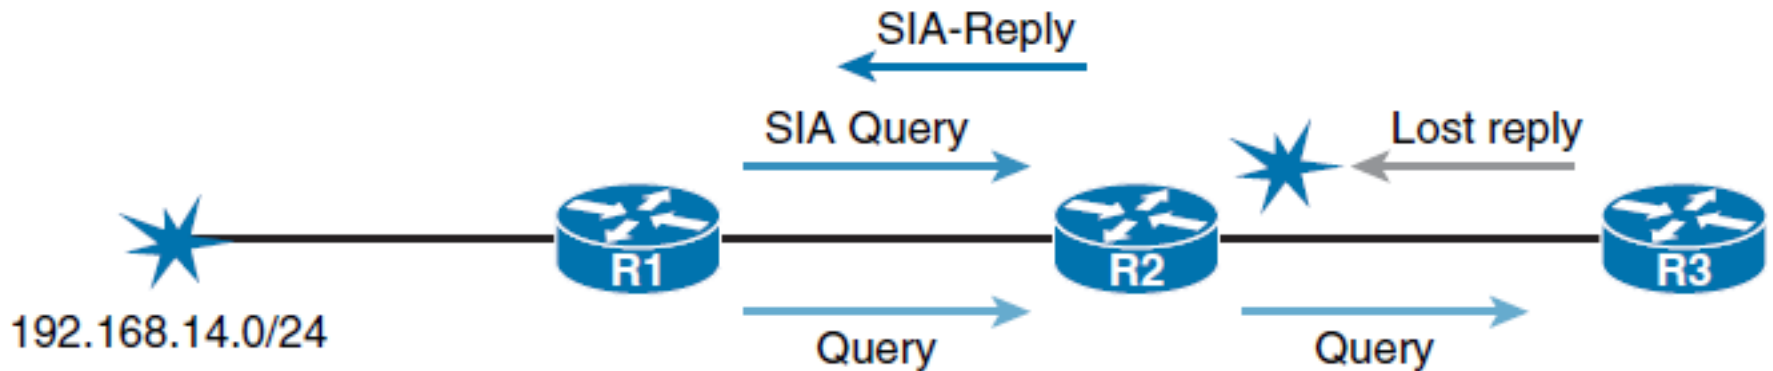
- Two new additional EIGRP packets were introduced to overcome the described limitation.
- When no reply to a query is received, **EIGRP sends an SIA query packet** when the **active timer is halfway** through (after 90 seconds).
- This enables the neighboring router to respond with a SIA reply and confirm to the upstream router that it is still searching for a replacement route.





Stuck in Active – overcome

- R1 queries downstream R2 (with an SIA query) at the midway point of the active timer about the status of the route.
- R2 responds (with an SIA reply) that it still is searching for a replacement route.
- Upon receiving this SIA reply response packet, R1 validates the status of R2 and does not terminate the neighbor relationship.
- Meanwhile, R2 will send up to three SIA queries to R3. If they go unanswered, R2 will terminate the neighbor relationship with R3. R2 will then update R1 with an SIA reply indicating that the network 192.168.14.0/24 is unreachable.
- R1 and R2 will remove the active route from their topology tables. The neighbor relationship between R1 and R2 remains intact.






Query propagation optimization

- **Two major solutions** exist to optimize the query propagation process and to limit the amount of unnecessary EIGRP load on the links.
 - **Route summarization**
 - **EIGRP stub routing**



Reducing Query Scope by Using Summary Routes


- manual summarization
- auto-summary





Configuring EIGRP Summarization

- Implementing EIGRP summarization provides several benefits. Not only does it **reduce the size of routing tables** on the routers, but it **also limits the query scope**.





HQ RT without summarization

```
HQ# show ip route
```

```
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 209.165.200.226 to network 0.0.0.0
```

```
S*    0.0.0.0/0 [1/0] via 209.165.200.226
      10.0.0.0/16 is subnetted, 7 subnets
```

```
D     10.10.0.0 [90/409600] via 172.16.1.2, 00:18:16, Ethernet0/1
D     10.11.0.0 [90/409600] via 172.16.1.2, 00:18:16, Ethernet0/1
D     10.12.0.0 [90/409600] via 172.16.1.2, 00:18:16, Ethernet0/1
D     10.13.0.0 [90/409600] via 172.16.1.2, 00:18:16, Ethernet0/1
D     10.64.0.0 [90/409600] via 172.16.2.2, 00:16:55, Ethernet0/2
D     10.65.0.0 [90/409600] via 172.16.2.2, 00:16:55, Ethernet0/2
D     10.66.0.0 [90/409600] via 172.16.2.2, 00:16:55, Ethernet0/2
```

```
      172.16.0.0/16 is variably subnetted, 4 subnets, 2 masks
```

```
C     172.16.1.0/30 is directly connected, Ethernet0/1
```



Enable auto-summarization in the BRs

```
BR1(config)# router eigrp 1
BR1(config-router)# auto-summary
*Oct 26 08:56:42.288: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 172.16.1.1
(Ethernet0/0) is resync: summary configured
*Oct 26 08:56:42.292: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 172.16.1.1
(Ethernet0/0) is resync: summary up, remove components
```

```
BR2(config)# router eigrp 1
BR2(config-router)# auto-summary
BR2(config-router)#
*Oct 26 09:30:45.251: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 172.16.2.1
(Ethernet0/0) is resync: summary configured
*Oct 26 09:30:45.255: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 172.16.2.1
(Ethernet0/0) is resync: summary up, remove components
```



HQ RT after summarization

```
HQ# show ip route
```

```
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 209.165.200.226 to network 0.0.0.0
```

```
S* 0.0.0.0/0 [1/0] via 209.165.200.226
```

```
D 10.0.0.0/8 [90/409600] via 172.16.2.2, 00:51:02, Ethernet0/2
    [90/409600] via 172.16.1.2, 00:51:02, Ethernet0/1
```

```
172.16.0.0/16 is variably subnetted, 4 subnets, 2 masks
```

```
C 172.16.1.0/30 is directly connected, Ethernet0/1
```

```
L 172.16.1.1/32 is directly connected, Ethernet0/1
```

```
C 172.16.2.0/30 is directly connected, Ethernet0/2
```

```
L 172.16.2.1/32 is directly connected, Ethernet0/2
```

```
209.165.200.0/24 is variably subnetted, 2 subnets, 2 masks
```

```
C 209.165.200.224/27 is directly connected, Ethernet0/0
```

```
L 209.165.200.225/32 is directly connected, Ethernet0/0
```



Testing Connectivity from HQ to the Summarized Network


- Because both routes have the same cost, HQ will load balance between these two routes causing unreachability issues.

```
HQ# ping 10.10.0.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.0.1, timeout is 2 seconds:
U.U.U
Success rate is 0 percent (0/5)
```

- This occurs because both remote networks belong to the same classfull, creating a discontinuous network connection.



Configuring Manual Summarization



```
BR1(config)# interface Ethernet 0/0
BR1(config-if)# ip summary-address eigrp 1 10.8.0.0/13
*Dec 3 13:22:53.406: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 172.16.1.1
(Ethernet0/0) is resync: summary configured
```

```
BR2(config)# interface Ethernet 0/0
BR2(config-if)# ip summary-address eigrp 1 10.64.0.0/14
*Dec 3 13:31:55.741: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 172.16.2.1
(Ethernet0/0) is resync: summary configured
```



HQ BR1A after Manual Summarization

```
BR1# show ip route
```

```
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, 10 subnets, 4 masks
```

```
D 10.8.0.0/13 is a summary, 00:43:25, Null0
```

```
C 10.10.0.0/16 is directly connected, Loopback0
```

```
L 10.10.0.1/32 is directly connected, Loopback0
```

```
C 10.11.0.0/16 is directly connected, Loopback1
```

```
L 10.11.0.1/32 is directly connected, Loopback1
```

```
C 10.12.0.0/16 is directly connected, Loopback2
```

```
L 10.12.0.1/32 is directly connected, Loopback2
```

```
C 10.13.0.0/16 is directly connected, Loopback3
```

```
L 10.13.0.1/32 is directly connected, Loopback3
```

```
D 10.64.0.0/14 [90/435200] via 172.16.1.1, 00:24:06, Ethernet0/0
```

```
172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
```

```
C 172.16.1.0/30 is directly connected, Ethernet0/0
```



HQ RT after Manual Summarization

```
HQ# show ip route
```

```
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 209.165.200.226 to network 0.0.0.0
```

```
S* 0.0.0.0/0 [1/0] via 209.165.200.226
```

```
10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
```

```
D 10.8.0.0/13 [90/409600] via 172.16.1.2, 01:30:05, Ethernet0/1
```

```
D 10.64.0.0/14 [90/409600] via 172.16.2.2, 01:10:46, Ethernet0/2
```

```
172.16.0.0/16 is variably subnetted, 4 subnets, 2 masks
```

```
C 172.16.1.0/30 is directly connected, Ethernet0/1
```

```
L 172.16.1.1/32 is directly connected, Ethernet0/1
```

```
C 172.16.2.0/30 is directly connected, Ethernet0/2
```

```
L 172.16.2.1/32 is directly connected, Ethernet0/2
```

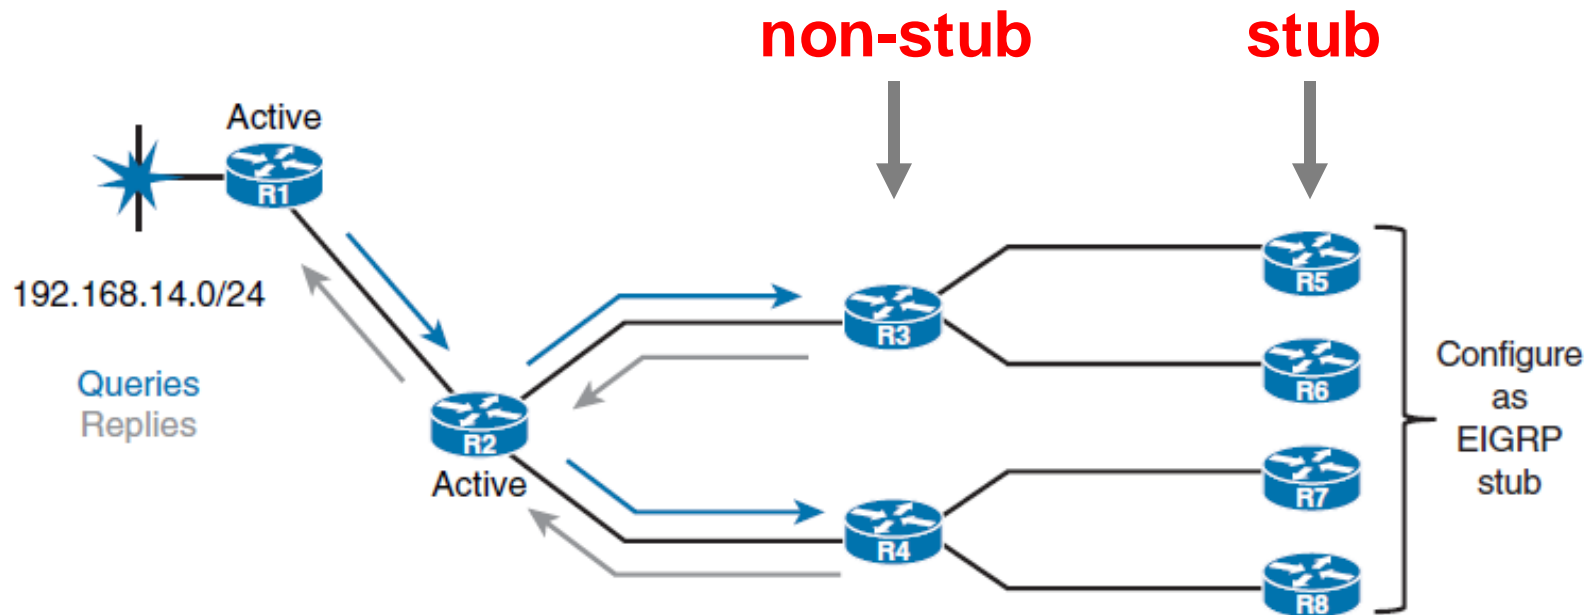
```
209.165.200.0/24 is variably subnetted, 2 subnets, 2 masks
```

```
C 209.165.200.224/27 is directly connected, Ethernet0/0
```

```
L 209.165.200.225/32 is directly connected, Ethernet0/0
```



EIGRP Stub Routers



- Routers configured as **stubs** do not forward EIGRP learned routes to other neighbors, and more importantly, **non-stub routers do not send query messages to stub routers.**
- This saves CPU cycles and bandwidth and speeds up convergence.



EIGRP Stub Options

Parameter	Description
receive-only	(Optional) Sets the router as a receive-only neighbor
leak-map <i>name</i>	(Optional) Allows dynamic prefixes based on a leak map
connected	(Optional) Advertises connected routes
static	(Optional) Advertises static routes
summary	(Optional) Advertises summary routes
redistributed	(Optional) Advertises redistributed routes from other protocols and autonomous systems

by default

- To configure a router as a stub use the **igmp stub** command in router configuration mode or address **family configuration mode**.
- To disable the EIGRP stub routing feature, use the **no** form of this command.



EIGRP Stub Options

- A router that is configured as a stub shares information about **connected** and **summary** routes with all neighboring routers **by default**.
- You can **combine all stub options** except for **receive-only** to achieve desired combination of advertised routes.
- The **connected** option permits the EIGRP stub router to advertise all connected routes for interfaces that are matched with an EIGRP network command. This option is enabled by default and is the most widely practical stub option.
- The **summary** option permits the EIGRP stub router to send summary routes. You can create summary routes manually, or you can create them automatically by enabling **auto-summary** at a major network boundary router.




EIGRP Stub Options

- The **static** option permits the EIGRP stub router to advertise static routes. You still need to redistribute static routes into EIGRP using the **redistribute static** command.
- The **redistribute** option permits the EIGRP stub router to advertise all redistributed routes, as long as redistribution is configured on the stub router using the **redistribute** command.
- The **receive-only** option restricts the stub router from sharing any of its routes with any other router within an EIGRP autonomous system. This option does not permit any other option to be specified because it prevents any type of route from being sent.



EIGRP Topology for Configuring Stub Routers





Initial Configuration (HQ)

```
HQ# show ip route
```

```
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, 2 subnets, 2 masks
C       10.1.0.0/16 is directly connected, Loopback0
L       10.1.0.1/32 is directly connected, Loopback0
172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
C       172.16.1.0/30 is directly connected, Ethernet0/0
L       172.16.1.1/32 is directly connected, Ethernet0/0
D       192.168.0.0/24 [90/409600] via 172.16.1.2, 00:12:07, Ethernet0/0
D       192.168.16.0/23 [90/307200] via 172.16.1.2, 00:12:07, Ethernet0/0
D EX   192.168.18.0/24 [170/307200] via 172.16.1.2, 00:12:07, Ethernet0/0
```



Initial Configuration (BR1A)

```
BR1A# show ip route
```

```
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
```

```
10.0.0.0/16 is subnetted, 1 subnets
```

```
D 10.1.0.0 [90/409600] via 172.16.1.1, 00:34:56, Ethernet0/0
```

```
172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
```

```
C 172.16.1.0/30 is directly connected, Ethernet0/0
```

```
L 172.16.1.2/32 is directly connected, Ethernet0/0
```

```
192.168.0.0/24 is variably subnetted, 2 subnets, 2 masks
```

```
C 192.168.0.0/24 is directly connected, Loopback0
```

```
L 192.168.0.1/32 is directly connected, Loopback0
```

```
D 192.168.16.0/23 is a summary, 03:05:24, Null0
```

```
192.168.16.0/24 is variably subnetted, 2 subnets, 2 masks
```

```
C 192.168.16.0/24 is directly connected, Ethernet0/1
```

```
L 192.168.16.1/32 is directly connected, Ethernet0/1
```

```
S 192.168.17.0/24 [1/0] via 192.168.16.2
```

```
S 192.168.18.0/24 [1/0] via 192.168.16.2
```



Configuring an EIGRP Stub Router

```

BR1A(config)# router eigrp 1
BR1A(config-router)# eigrp stub
*Oct 18 11:51:16.232: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 172.16.1.1
(Ethernet0/0) is down: peer info changed
BR1A(config-router)#
*Oct 18 11:51:20.495: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 172.16.1.1
(Ethernet0/0) is up: new adjacency

-----

*Oct 18 11:51:16.228: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 172.16.1.2
(Ethernet0/0) is down: Interface PEER-TERMINATION received
HQ#
*Oct 18 11:51:20.503: EIGRP: Adding stub (1 Peers, 1 Stubs)
*Oct 18 11:51:20.503: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 172.16.1.2
(Ethernet0/0) is up: new adjacency
*Oct 18 11:51:20.503: EIGRP: Enqueueing UPDATE on Ethernet0/0 nbr 172.16.1.2 tid 0
iidbQ un/rely 0/1 peerQ un/rely 0/0
*Oct 18 11:51:20.508: EIGRP: Received UPDATE on Ethernet0/0 nbr 172.16.1.2
<Output omitted>

```

- When the BR1A router is configured as a stub, the EIGRP adjacency needs to be reestablished.



BR1A Verified as an EIGRP Stub on HQ

```
HQ# show ip eigrp neighbors detail
```

```
EIGRP-IPv4 Neighbors for AS(1)
```

H	Address	Interface	Hold Uptime (sec)	SRTT (ms)	RTO	Q Cnt	Seq Num
0	172.16.1.2	Et0/0	11 00:39:00	7	100	0	13

```
Version 7.0/3.0, Retrans: 0, Retries: 0, Prefixes: 2
```

```
Topology-ids from peer - 0
```

```
Stub Peer Advertising (CONNECTED SUMMARY ) Routes
```

```
Suppressing queries
```

```
BFD sessions
```

```
NeighAddr      Interface
```



Change in the routing tables

- BR1A Routing Table stays exactly the same but HQ:

```
HQ# show ip route
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, 2 subnets, 2 masks
C       10.1.0.0/16 is directly connected, Loopback0
L       10.1.0.1/32 is directly connected, Loopback0
172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
C       172.16.1.0/30 is directly connected, Ethernet0/0
L       172.16.1.1/32 is directly connected, Ethernet0/0
D       192.168.0.0/24 [90/409600] via 172.16.1.2, 00:09:07, Ethernet0/0
D       192.168.16.0/23 [90/307200] via 172.16.1.2, 00:09:07, Ethernet0/0
```



Configuring an EIGRP Connected Stub

```
BR1A(config)# router eigrp 1
BR1A(config-router)# eigrp stub connected
*Oct 20 18:46:50.137: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 172.16.1.1
(Ethernet0/0) is down: peer info changed
*Oct 20 18:46:50.419: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 172.16.1.1
(Ethernet0/0) is up: new adjacency
```

```
HQ# 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   172.16.1.2                Et0/0                  14 00:10:25    12    100   0    8
Version 7.0/3.0, Retrans: 0, Retries: 0, Prefixes: 2
Topology-ids from peer - 0
Stub Peer Advertising (CONNECTED ) Routes
Suppressing queries

BFD sessions
NeighAddr                Interface
```




Change in the routing tables

```

HQ# show ip route
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, 2 subnets, 2 masks
C       10.1.0.0/16 is directly connected, Loopback0
L       10.1.0.1/32 is directly connected, Loopback0
172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
C       172.16.1.0/30 is directly connected, Ethernet0/0
L       172.16.1.1/32 is directly connected, Ethernet0/0
D       192.168.0.0/24 [90/409600] via 172.16.1.2, 00:14:52, Ethernet0/0
D       192.168.16.0/24 [90/307200] via 172.16.1.2, 00:14:52, Ethernet0/0

```



Configuring an EIGRP Receive-Only Stub

```
BR1A(config)# router eigrp 1
BR1A(config-router)# eigrp stub receive-only
*Oct 20 19:06:42.909: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 172.16.1.1
(Ethernet0/0) is down: peer info changed
BR1A(config-router)#
*Oct 20 19:06:46.356: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 172.16.1.1
(Ethernet0/0) is up: new adjacency
```

```
HQ# show ip eigrp neighbors detail
EIGRP-IPv4 Neighbors for AS(1)
```

H	Address	Interface	Hold (sec)	Uptime	SRTT (ms)	RTO	Q Cnt	Seq Num
0	172.16.1.2	Et0/0	10	00:03:03	1999	5000	0	10

```
Version 7.0/3.0, Retrans: 1, Retries: 0
Topology-ids from peer - 0
Receive-Only Peer Advertising (No) Routes
Suppressing queries

BFD sessions
NeighAddr      Interface
```




Change in the routing tables

```
HQ# show ip route
```

```
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, 2 subnets, 2 masks
```

```
C      10.1.0.0/16 is directly connected, Loopback0
```

```
L      10.1.0.1/32 is directly connected, Loopback0
```

```
172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
```

```
C      172.16.1.0/30 is directly connected, Ethernet0/0
```

```
L      172.16.1.1/32 is directly connected, Ethernet0/0
```



Obtaining Default Route

- The candidate can be a statically configured default route defined locally with the command **ip route 0.0.0.0 0.0.0.0** *next-hop | interface* .
- The candidate can also be a default route announced by the dynamic routing protocol. **EIGRP can redistribute statically defined default routes** by using the **redistribute static** configuration command.
- In addition, **any classful network** residing in the local routing table can become a default candidate when used with the **ip default-network** configuration command.
 - The command attaches an exterior flag to any classful EIGRP route, thus making it a candidate for a default route.



Redistributing Static

```

HQ# show ip route

Gateway of last resort is 209.165.200.226 to network 0.0.0.0

S*   0.0.0.0/0 [1/0] via 209.165.200.226
      10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
D     10.8.0.0/13 [90/409600] via 172.16.1.2, 23:00:26, Ethernet0/1
D     10.64.0.0/14 [90/409600] via 172.16.2.2, 22:41:07, Ethernet0/2
      172.16.0.0/16 is variably subnetted, 4 subnets, 2 masks
C     172.16.1.0/30 is directly connected, Ethernet0/1
L     172.16.1.1/32 is directly connected, Ethernet0/1
C     172.16.2.0/30 is directly connected, Ethernet0/2
L     172.16.2.1/32 is directly connected, Ethernet0/2
      209.165.200.0/24 is variably subnetted, 2 subnets, 2 masks
C     209.165.200.224/27 is directly connected, Ethernet0/0
L     209.165.200.225/32 is directly connected, Ethernet0/0

HQ# configure terminal
HQ(config)# router eigrp 1
HQ(config-router)# redistribute static
  
```



Redistributing Static

```
BR1# show ip route
```

```
Gateway of last resort is 172.16.1.1 to network 0.0.0.0
```

```
D*EX 0.0.0.0/0 [170/307200] via 172.16.1.1, 00:17:06, Ethernet0/0
```

```
10.0.0.0/8 is variably subnetted, 10 subnets, 4 masks
```

```
D 10.8.0.0/13 is a summary, 23:10:22, Null0
```

```
C 10.10.0.0/16 is directly connected, Loopback0
```

```
L 10.10.0.1/32 is directly connected, Loopback0
```

```
C 10.11.0.0/16 is directly connected, Loopback1
```

```
L 10.11.0.1/32 is directly connected, Loopback1
```

```
C 10.12.0.0/16 is directly connected, Loopback2
```

```
L 10.12.0.1/32 is directly connected, Loopback2
```

```
C 10.13.0.0/16 is directly connected, Loopback3
```

```
L 10.13.0.1/32 is directly connected, Loopback3
```

```
D 10.64.0.0/14 [90/435200] via 172.16.1.1, 22:51:03, Ethernet0/0
```

```
172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
```

```
C 172.16.1.0/30 is directly connected, Ethernet0/0
```

```
L 172.16.1.2/32 is directly connected, Ethernet0/0
```

```
D 172.16.2.0/30 [90/307200] via 172.16.1.1, 1d03h, Ethernet0/0
```

```
209.165.200.0/27 is subnetted, 1 subnets
```

```
D 209.165.200.224 [90/307200] via 172.16.1.1, 1d03h, Ethernet0/0
```



Load Balancing with EIGRP

- EIGRP can **distribute traffic over multiple links** leading to the same destination to increase the effective network bandwidth.
- It supports load balancing over **equal-metric paths** and also over **unequal-metric paths**.
- EIGRP enables load balancing between a **maximum of 4** equal-metric paths by default.
- The maximum number of parallel routes that an IP routing protocol can support can be changed using the **maximum-paths** router configuration command.




Load Balancing with EIGRP

- When a packet is **process switched**, load balancing over equal-metric paths occurs on a **per-packet basis**.
- When packets are **fast switched**, load balancing over equal-metric paths occurs on a **per-destination basis**.
- **Cisco Express Forwarding** (CEF) switching, enabled by default, supports both per-packet and per-destination load balancing.
- **Load balancing over unequal-metric links is disabled by default.**
- Only **feasible successor** paths can be included in the EIGRP load-balancing, to ensure the topology stays **loop free**.



Configuring EIGRP Load Balancing





Configuring EIGRP Load Balancing

```
HQ# show ip route
    172.16.0.0/16 is variably subnetted, 6 subnets, 2 masks
C       172.16.1.0/30 is directly connected, Ethernet0/0
L       172.16.1.1/32 is directly connected, Ethernet0/0
C       172.16.2.0/30 is directly connected, Ethernet0/1
L       172.16.2.1/32 is directly connected, Ethernet0/1
C       172.16.3.0/30 is directly connected, Serial1/0
L       172.16.3.1/32 is directly connected, Serial1/0
D       192.168.0.0/24 [90/409600] via 172.16.2.2, 00:26:18, Ethernet0/1
        [90/409600] via 172.16.1.2, 00:26:18, Ethernet0/0
```

```
HQ# show ip eigrp topology
EIGRP-IPv4 Topology Table for AS(1)/ID(172.16.3.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 172.16.2.0/30, 1 successors, FD is 281600
   via Connected, Ethernet0/1
P 192.168.0.0/24, 2 successors, FD is 409600
   via 172.16.1.2 (409600/128256), Ethernet0/0
   via 172.16.2.2 (409600/128256), Ethernet0/1
   via 172.16.3.2 (2297856/128256), Serial1/0
P 172.16.3.0/30, 1 successors, FD is 2169856
   via Connected, Serial1/0
P 172.16.1.0/30, 1 successors, FD is 281600
   via Connected, Ethernet0/0
```




Load Balancing Across Unequal-Metric Paths

- The route must be loop free. This condition is satisfied when the advertised distance is less than the total distance, or when the **route is a feasible successor**.
- The metric of the route must be lower than the metric of the best route (the successor) multiplied by the variance that is configured on the router.

```
HQ(config)# router eigrp 1
HQ(config-router)# variance 6
```



Verifying EIGRP Load-Balance

```
HQ# show ip route
```

```
172.16.0.0/16 is variably subnetted, 6 subnets, 2 masks
```

```
C      172.16.1.0/30 is directly connected, Ethernet0/0
```

```
L      172.16.1.1/32 is directly connected, Ethernet0/0
```

```
C      172.16.2.0/30 is directly connected, Ethernet0/1
```

```
L      172.16.2.1/32 is directly connected, Ethernet0/1
```

```
C      172.16.3.0/30 is directly connected, Serial1/0
```

```
L      172.16.3.1/32 is directly connected, Serial1/0
```

```
D      192.168.0.0/24 [90/2297856] via 172.16.3.2, 00:02:03, Serial1/0
          [90/409600] via 172.16.2.2, 00:02:03, Ethernet0/1
          [90/409600] via 172.16.1.2, 00:02:03, Ethernet0/0
```



Verifying EIGRP Load-Balance

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

Routing Protocol is "eigrp 1"
  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
  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: 172.16.3.1
    Topology : 0 (base)
      Active Timer: 3 min
      Distance: internal 90 external 170
      Maximum path: 4
      Maximum hopcount 100
      Maximum metric variance 6

  Automatic Summarization: disabled
  Maximum path: 4
  Routing for Networks:
    0.0.0.0
  Routing Information Sources:
    Gateway          Distance          Last Update
```

Configuring EIGRP for IPv6





Configuring EIGRP for IPv6

- Differences and Similarities of EIGRP for IPv4 and IPv6
- Configure Basic EIGRP for IPv6 Settings
- Configure and Verify EIGRP for IPv6 Summarization
- Verify basic EIGRP for IPv6 settings



Overview of EIGRP for IPv6

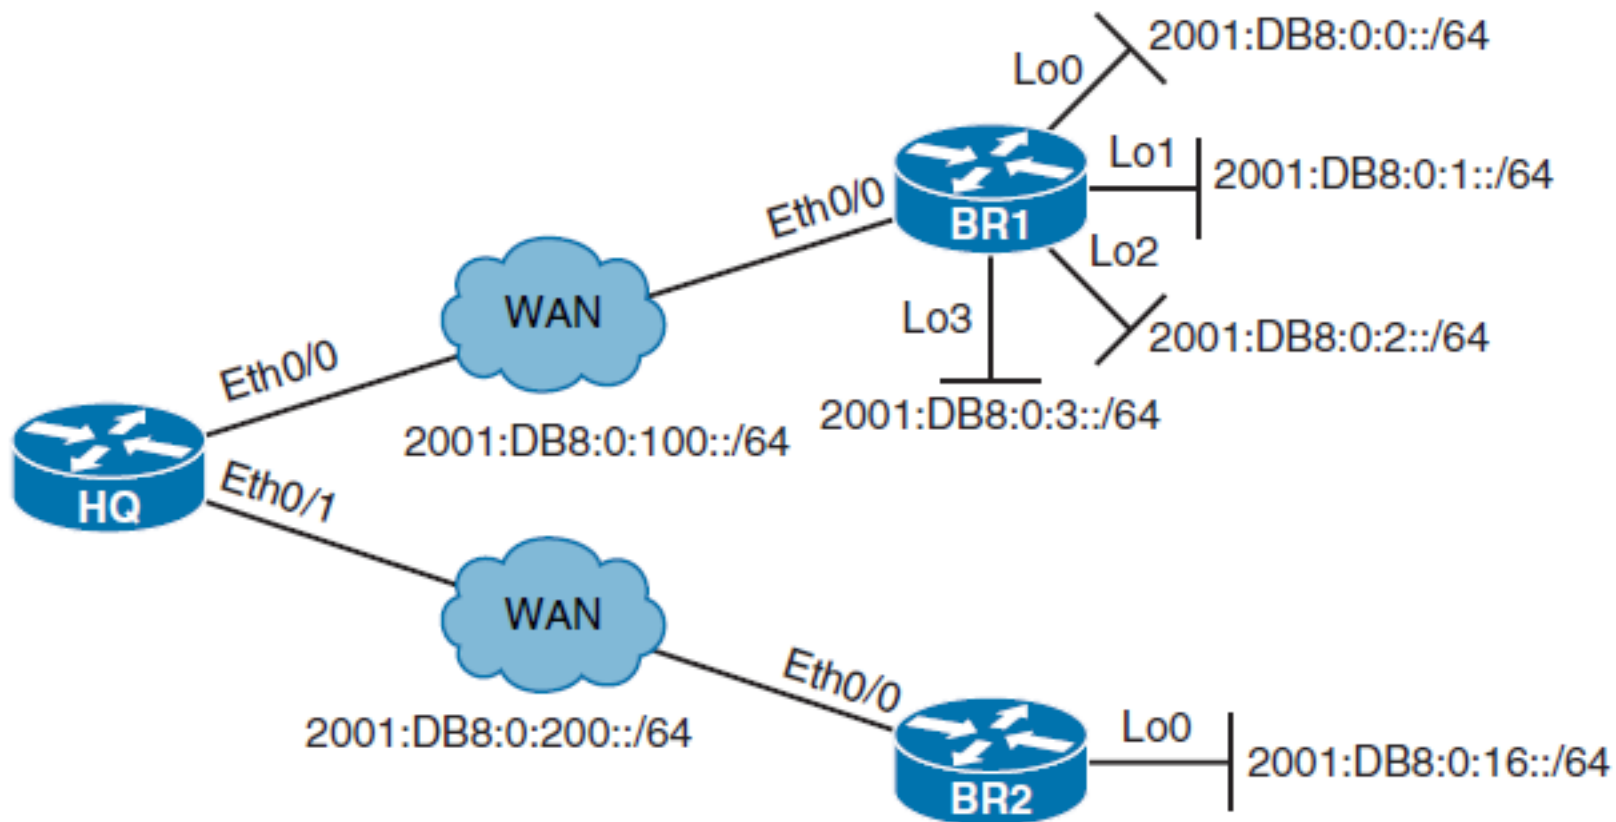
- EIGRP for IPv6 uses **IPv6 prefixes** and **lengths** rather than IPv4 subnets and masks.
- To establish EIGRP for IPv6 neighbor relationship, it uses IPv6 **link-local addresses**. EIGRP for IPv4 does not have the concept of link-local address.
- EIGRP uses **built-in authentication features** of the IPv6 protocol rather than protocol specific authentication implemented with IPv4 to guarantee message authentication.
- To transport routing information, EIGRP for IPv6 encapsulates IPv6 prefixes in the IPv6 messages, not in the IPv4 packets.
- IPv6 has **no concept of the classful network**; so, when you use EIGRP for IPv6, there is **no automatic summarization** at the class boundaries. The **only way** to summarize IPv6-advertised prefixes is through **manual summarization**.



Overview of EIGRP for IPv6

- If IPv4 address is not configured on the router, EIGRP for IPv6 **requires** an EIGRP **router ID** before it can start running. In IPv4, if you do not configure the EIGRP router ID, the router will automatically assign it using the highest loopback or the highest active interface IPv4 address.
- EIGRP for IPv6 under a specific interface intended to send and receive routing protocol messages. In EIGRP for IPv4, you configure interfaces under the routing protocol configuration mode.
- EIGRP for IPv6 uses assigned dedicated multicast address FF02::A, whereas EIGRP for IPv4 uses dedicated multicast address 224.0.0.10.

Configuring and Verifying EIGRP for IPv6





Configuring and Verifying EIGRP for IPv6

```
BR2# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
BR2(config)# ipv6 unicast-routing
```

The **ipv6 unicast-routing** command enables the router:

- To be configured for static and dynamic IPv6 routing
- To forward IPv6 packets
- To send ICMPv6 router advertisement messages



Configuring and Verifying EIGRP for IPv6

```
BR2(config)# ipv6 router eigrp 100
BR2(config-rtr)# eigrp router-id 192.168.2.1
```

- EIGRP for IPv6 has a shutdown feature.
- The routing process must be in “no shutdown” mode for EIGRP for IPv6 processing.
- No shutdown is the default on later IOSs. If necessary, you might have to issue the **no shutdown** command in EIGRP for IPv6 configuration mode.
- **Another important parameter is the EIGRP router ID.**
- Like EIGRP for IPv4, EIGRP for IPv6 uses a 32-bit router ID.
- If no IPv4 active address is configured on the router, the router will not be able to choose the EIGRP router ID.
- In this case, you must configure the router ID manually under the EIGRP routing process to make EIGRP for IPv6 operational.



Configuring and Verifying EIGRP for IPv6

- Before you enable EIGRP for IPv6 on the interface, it must have a valid IPv6 link-local address.
- EIGRP for IPv6 uses link-local addresses to form EIGRP neighbor relationships.

```
BR2(config)# interface ethernet 0/0
BR2(config-if)# ipv6 eigrp 100
*Oct 23 19:57:55.933: %DUAL-5-NBRCHANGE: EIGRP-IPv6 100: Neighbor
FE80:200::1 (Ethernet0/0) is up: new adjacency
BR2(config-if)# exit
BR2(config)# interface loopback 0
BR2(config-if)# ipv6 eigrp 100
```



Link-Local Importance for IPv6 Routing

- The link-local address is **automatically created** on an interface when the interface obtains a global IPv6 address, either manually or dynamically.
- IPv6 can also be enabled on an interface without assigning a global unicast address using the interface mode command **ipv6 enable**.
- In both cases, IPv6 link-local address will be assigned automatically to the interface using **EUI-64**.



Verifying EIGRP for IPv6 Neighbor Adjacency and Topology Table

```
BR2# show ipv6 eigrp neighbors
```

```
EIGRP-IPv6 Neighbors for AS(100)
```

H	Address	Interface	Hold Uptime (sec)	SRTT (ms)	RTO	Q Cnt	Seq Num
0	Link-local address: FE80:200::1	Et0/0	13 08:25:34	9	100	0	16

```
BR2# show ipv6 eigrp topology
```

```
EIGRP-IPv6 Topology Table for AS(100)/ID(192.168.2.1)
```

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

```
P 2001:DB8:0:2::/64, 1 successors, FD is 435200  
   via FE80:200::1 (435200/409600), Ethernet0/0  
P 2001:DB8:0:200::/64, 1 successors, FD is 281600  
   via Connected, Ethernet0/0  
P 2001:DB8::/64, 1 successors, FD is 435200  
   via FE80:200::1 (435200/409600), Ethernet0/0  
P 2001:DB8:0:1::/64, 1 successors, FD is 435200  
   via FE80:200::1 (435200/409600), Ethernet0/0  
P 2001:DB8:0:3::/64, 1 successors, FD is 435200  
   via FE80:200::1 (435200/409600), Ethernet0/0  
P 2001:DB8:0:100::/64, 1 successors, FD is 307200  
   via FE80:200::1 (307200/281600), Ethernet0/0
```



Verifying EIGRP for IPv6 Routing Table

```
BR2# show ipv6 route eigrp
IPv6 Routing Table - default - 10 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, I1 - ISIS L1, I2 - ISIS L2
       IA - ISIS interarea, IS - ISIS summary, D - EIGRP, EX - EIGRP external
       ND - 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
D 2001:DB8::/64 [90/435200]
   via FE80:200::1, Ethernet0/0
D 2001:DB8:0:1::/64 [90/435200]
   via FE80:200::1, Ethernet0/0
D 2001:DB8:0:2::/64 [90/435200]
   via FE80:200::1, Ethernet0/0
D 2001:DB8:0:3::/64 [90/435200]
   via FE80:200::1, Ethernet0/0
D 2001:DB8:0:100::/64 [90/307200]
   via FE80:200::1, Ethernet0/0
```



Configuring EIGRP for IPv6 Summary Route

Prefix	Binary Format
2001:DB8:0:0::64	2001:DB8:0:0000000000000000
2001:DB8:0:1::64	2001:DB8:0:0000000000000001
2001:DB8:0:2::64	2001:DB8:0:0000000000000010
Summary route	
2001:DB8:0:0::62	2001:DB8:0:0000000000000000

```
BR1(config)# interface Ethernet0/0
BR1(config-if)# ipv6 summary-address eigrp 100 2001:DB8:0:0::/62
*Oct 24 18:14:31.222: %DUAL-5-NBRCHANGE: EIGRP-IPv6 100: Neighbor
FE80:100::1 (Ethernet0/0) is resync: summary configured
```

Named EIGRP Configuration





Named EIGRP Configuration

- Describe how EIGRP named configuration is different from the classic EIGRP configuration
- Explain what is configured under different address family configuration modes
- Compare examples of classic and named EIGRP configuration
- Configuring and verifying EIGRP for IPv6




Named EIGRP Configuration

- Configuring EIGRP for both IPv4 and IPv6 on the same router can become a complex task because configuration takes place using different router configuration modes
- A newer configuration enables the configuration of EIGRP **for both IPv4 and IPv6 under a single configuration mode.**
- EIGRP named configuration helps eliminate configuration complexity that occurs when configuring EIGRP for both IPv4 and IPv6
- EIGRP named configuration is available in Cisco IOS Release 15.0(1)M and later releases.



Configuring Named EIGRP





Configuring Named EIGRP (Previous Config)

```

BR2# show running-config
<Output omitted>
interface Loopback0
  ip address 192.168.2.1 255.255.255.0
  ipv6 address 2001:DB8:0:2::1/64
  ipv6 enable
  ipv6 eigrp 1
!
interface Loopback1
  ip address 192.168.3.1 255.255.255.0
  ipv6 address 2001:DB8:0:3::1/64
  ipv6 enable
  ipv6 eigrp 1
!
interface Ethernet0/0
  ip address 172.16.1.6 255.255.255.252
  ipv6 address 2001:DB8:0:101::2/64
  ipv6 enable
  ipv6 eigrp 1
<Output omitted>
router eigrp 1
  network 0.0.0.0
<Output omitted>
ipv6 router eigrp 1
<Output omitted>

```



Address Families

- EIGRP named configuration mode uses the global configuration command **router eigrp *virtual-instance-name***.
- Both EIGRP for IPv4 and IPv6 can be configured within this same mode.
- **EIGRP supports multiple protocols** and can carry information about many different route types.
- **Named EIGRP configuration organizes specific route types under the same **address family**.**
- IPv4 unicast and IPv6 unicast are two of the most commonly used address families.



EIGRP for IPv4 Address Family

```
BR2(config)# router eigrp LAB
BR2(config-router)# address-family ipv4 autonomous-system 1
BR2(config-router-af)#
```

address-family ipv4 [**multicast**] [**unicast**] [**vrf vrf-name**]
autonomous-system as-number

Table 2-4 *Parameters for the EIGRP address-family ipv4 Command*

Parameter	Description
ipv4	Selects the IPV4 protocol address family.
multicast	(Optional) Specifies the multicast address family. This keyword is available only in EIGRP named IPv4 configurations.
unicast	(Optional) Specifies the unicast address family. This is the default.
vrf <i>vrf-name</i>	(Optional) Specifies the name of the VRF.
autonomous-system <i>autonomous-system-number</i>	Specifies the autonomous system number.



EIGRP for IPv4 Address Family

- The **address-family ipv4** command enables the IPv4 address family and starts EIGRP for the defined autonomous system.
- In IPv4 address family configuration mode, you can enable EIGRP for specific interfaces by using the **network** command, and you can define some other general parameters such as **router-id** or **eigrp stub**.
- Unless specified otherwise, address family is by default defined as unicast address family used the exchange unicast routes.

```
BR2(config-router-af)# network 0.0.0.0
*Dec 27 14:15:53.944: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 172.16.1.5
(Ethernet0/0) is up: new adjacency
```



EIGRP for IPv6 Address Family

```
BR2(config)# router eigrp LAB
BR2(config-router)# address-family ipv6 autonomous-system 1
BR2(config-router-af)#
*Dec 30 09:37:23.652: %DUAL-5-NBRCHANGE: EIGRP-IPv6 1: Neighbor
FE80::A8BB:CCFF:FE00:3310 (Ethernet0/0) is up: new adjacency
```

address-family ipv6 [unicast] [vrf vrf-name] autonomous-system as-number

Parameter

Description

Ipv6

Selects the IPv6 protocol address family.

unicast

(Optional) Specifies the unicast address family. This is the default.

vrf vrf-name

(Optional) Specifies the name of the VRF.

autonomous-system

Specifies the autonomous system number.

autonomous-system- number



EIGRP for IPv6 Address Family

- IPv6 EIGRP neighbor relationship gets established as soon as you define the IPv6 address family.
- All IPv6-enabled interfaces are automatically included in the EIGRP for IPv6 process.
- The IPv6 address family configuration will show up in the running configuration as a **unicast address family by default**.
- You can configure or remove individual interfaces from the EIGRP for IPv6 process by using the **af-interface** *interface-type interface number* command in address family configuration mode



EIGRP for IPv6 Address Family

- **af-interface** { **default** | *interface-type interface number* }

Parameter	Description
default	Specifies the default address family interface configuration mode. Commands applied under this mode affect all interfaces used by this address family instance.
<i>interface-type interface number</i>	Interface type and number of the interface that the address family submode commands will affect.

```
R2(config)#router eigrp LAB
R2(config-router)# !
R2(config-router)# address-family ipv4 unicast autonomous-system 12345
R2(config-router-af)#af-interface default
R2(config-router-af-interface)#authentication mode hmac-sha-256 0 SECURerouter
R2(config-router-af-interface)#
```

By configuring this on the af-interface default, I am configuring it by default for all interfaces in the address-family. This is easier and faster than the past where commands needed to be entered manually on all involved router interfaces.



Disabling EIGRP for IPv6 on an Interface

```
BR2(config)# router eigrp LAB
BR2(config-router)# address-family ipv6 autonomous-system 1
BR2(config-router-af)# af-interface ethernet 0/1
BR2(config-router-af-interface)# shutdown
```



Named EIGRP Final Config

```
BR2# show running configuration | section router eigrp
router eigrp LAB
!
address-family ipv4 unicast autonomous-system 1
!
topology base
exit-af-topology
network 0.0.0.0
exit-address-family
!
address-family ipv6 unicast autonomous-system 1
!
topology base
exit-af-topology
exit-address-family
```



EIGRP Summarization in Named Configuration

```
BR2(config)# router eigrp LAB
BR2(config-router)# address-family ipv4 autonomous-system 1
BR2(config-router-af)# af-interface ethernet 0/0
BR2(config-router-af-interface)# summary-address 192.168.2.0/23
BR2(config-router-af-interface)#
*Dec 30 13:36:07.935: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor
172.16.1.5 (Ethernet0/0) is resync: summary configured
```



Configuring IPv6 Passive Interfaces

```

BR2(config)# router eigrp LAB
BR2(config-router)# address-family ipv6 autonomous-system 1
BR2(config-router-af)# af-interface default
BR2(config-router-af-interface)# passive-interface
*Dec 31 08:42:40.864: %DUAL-5-NBRCHANGE: EIGRP-IPv6 1: Neighbor
FE80::A8BB:CCFF: FE00:F010 (Ethernet0/0) is down: interface passive
BR2(config-router-af-interface)# exit
BR2(config-router-af)# af-interface ethernet0/0
BR2(config-router-af-interface)# no passive-interface
*Dec 31 08:42:57.111: %DUAL-5-NBRCHANGE: EIGRP-IPv6 1: Neighbor
FE80::A8BB:CCFF: FE00:F010 (Ethernet0/0) is up: new adjacency

```



Named EIGRP Configuration Modes

Three different configuration modes:

■ Address family configuration mode

- General EIGRP configuration commands for selected address family are entered under address family configuration mode. Here you can configure the **router ID** and define **network statements** and also configure router as an **EIGRP stub**.
- Address family configuration mode gives you access to two additional configuration modes: address family interface configuration mode and address family topology configuration mode.

■ Address family interface configuration mode

- You should use address family interface configuration mode for all those commands that you have **previously configured directly under interfaces**. Most common options are setting **summarization** with the **summary-address** command or marking interfaces as passive using **passive-interface** command. You can also modify default hello and hold-time timers.

■ Address family topology configuration mode

- Address family topology configuration mode gathers all configuration options that **directly impact the EIGRP topology table**. Here you can set **load-balancing** parameters such as **variance** and **maximum-paths**, or you can redistribute static routes using the **redistribute** command.



Address Family Configuration Mode

```
BR1(config)# router eigrp LAB
```

```
BR1(config-router)# address-family ipv6 unicast autonomous-system 1
```

```
BR1(config-router-af)# ?
```

Address Family configuration commands:

af-interface	Enter Address Family interface configuration
default	Set a command to its defaults
eigrp	EIGRP Address Family specific commands
exit-address-family	Exit Address Family configuration mode
help	Description of the interactive help system
maximum-prefix	Maximum number of prefixes acceptable in aggregate
metric	Modify metrics and parameters for address advertisement
neighbor	Specify an IPv6 neighbor router
no	Negate a command or set its defaults
shutdown	Shutdown address family
timers	Adjust peering based timers
topology	Topology configuration mode

```
BR1(config-router-af)#
```




Address Family Interface Configuration Mode

```

BR1(config)# router eigrp LAB
BR1(config-router)# address-family ipv6 unicast autonomous-system 1
BR1(config-router-af)# af-interface ethernet 0/0
BR1(config-router-af-interface)# ?
Address Family Interfaces configuration commands:
  authentication      authentication subcommands
  bandwidth-percent  Set percentage of bandwidth percentage limit
  bfd                 Enable Bidirectional Forwarding Detection
  dampening-change    Percent interface metric must change to cause update
  dampening-interval Time in seconds to check interface metrics
  default             Set a command to its defaults
  exit-af-interface   Exit from Address Family Interface configuration mode
  hello-interval      Configures hello interval
  hold-time           Configures hold time
  next-hop-self       Configures EIGRP next-hop-self
  no                  Negate a command or set its defaults
  passive-interface   Suppress address updates on an interface
  shutdown            Disable Address-Family on interface
  split-horizon       Perform split horizon
  summary-address     Perform address summarization

BR1(config-router-af-interface)#

```



Address Family Topology Configuration Mode

```

BR1(config)# router eigrp LAB
BR1(config-router)# address-family ipv6 unicast autonomous-system 1
BR1(config-router-af)# topology base
BR1(config-router-af-topology)# ?
Address Family Topology configuration commands:
  default          Set a command to its defaults
  default-information Control distribution of default information
  default-metric   Set metric of redistributed routes
  distance         Define an administrative distance
  distribute-list  Filter entries in eigrp updates
  eigrp           EIGRP specific commands
  exit-af-topology Exit from Address Family Topology configuration mode
  maximum-paths   Forward packets over multiple paths
  metric          Modify metrics and parameters for advertisement
  no             Negate a command or set its defaults
  redistribute     Redistribute IPv6 prefixes from another routing protocol
  summary-metric  Specify summary to apply metric/filtering
  timers         Adjust topology specific timers
  traffic-share   How to compute traffic share over alternate paths
  variance       Control load balancing variance

BR1(config-router-af-topology)#

```



Chapter 2 Summary

- EIGRP is an advanced distance vector protocol.
- EIGRP uses RTP for reliable, guaranteed delivery of packets.
- Hello and hold timers can be adapted to influence network convergence.
- IGRP adapts well to various technologies such as Frame Relay, Layer 3 MPLS VPN, and Layer 2 MPLS VPN.
- EIGRP uses hello, update, query, reply, and acknowledgment packets.
- EIGRP uses a composite metric that is by default based on bandwidth and delay.
- Reported distance is the metric value reported by the neighboring router.
- Feasible distance is the lowest distance to a destination from the perspective of the local router.
- Alternative path must satisfy the feasibility condition to become a feasible successor.
- The reported distance of an alternate path must be less than the feasible distance.
- When a route is lost and no feasible successor is available, queries are sent to all neighboring routers on all interfaces.
- EIGRP stub configuration improves network stability and reduces resource utilization.
- Summarization decreases the size of the IP routing table and optimizes exchange of routing information.



Chapter 2 Summary

- EIGRP performs equal-cost load balancing.
- To support unequal-cost load balancing, a **variance** parameter must be configured.
- EIGRP for IPv6 uses IPv6 link-local addresses to form neighbor relationships.
- EIGRP for IPv6 supports only manual prefix summarization.
- To configure EIGRP for IPv6, you must define the routing process and configure interfaces participating in EIGRP routing.
- EIGRP for IPv6 verification commands have similar syntax to EIGRP for IPv4 commands.
- Classic EIGRP configuration is divided over different configuration modes.
- Named EIGRP configuration gathers EIGRP configuration in one place.
- Named EIGRP configuration unifies configuration commands for different address families.
- Named EIGRP configuration is hierarchically organized using three address family configuration modes.
- The same verification commands for classic EIGRP are used to verify named EIGRP configuration.



Chapter 2 Labs

- **CCNPv7_ROUTE_Lab2-1_EIGRP-Load-Balancing**
- **CCNPv7_ROUTE_Lab2-2_EIGRP-Stub-Routing**
- **CCNPv7_ROUTE_Lab2-3_EIGRP-IPv6**
- **CCNPv7_ROUTE_Lab2-4_EIGRP-Named-Configuration**

Cisco | Networking Academy[®]

Mind Wide Open[™]



Acknowledgment

- Some of images and texts are from Implementing Cisco IP Routing (ROUTE) Foundation Learning Guide by Diane Teare, Bob Vachon and Rick Graziani (1587204568)
- Copyright © 2015 – 2016 Cisco Systems, Inc.
- Special Thanks to *Bruno Silva*