

## **Chapter 10 - EIGRP**

EIGRP overview

EIGRP neighbors

EIGRP tables

Reported and feasible distance explained

Successor and feasible successor explained

EIGRP configuration

Wildcard mask explained

EIGRP and wildcard masks

Reliable Transport Protocol (RTP)

Diffusing Update Algorithm (DUAL)

EIGRP auto-summary

EIGRP manual summarization

## EIGRP overview

EIGRP (Enhanced Interior Gateway Routing Protocol) is an advanced distance vector routing protocol. It is sometimes referred to as a hybrid routing protocol because it has characteristics of both distance-vector and link-state routing protocols. EIGRP replaced Interior Gateway Routing Protocol (IGRP), an older proprietary Cisco routing protocol. EIGRP was also a proprietary protocol, but Cisco decided in 2013 to convert it to an open standard. This routing protocol is mostly used on Cisco devices and all routers in the network must support it.

The following features make EIGRP especially helpful in large and complex networks:

- support for classless routing and VLSM (Variable Length Subnet Masking).

- routes can be summarized on any router in the network.

- incremental routing updates.

- support for load-balancing.

- support for the MD5 authentication.

- support for IPv4 and IPv6.

EIGRP uses Cisco's Reliable Transport Protocol (RTP) to send messages and uses the multicast address of 224.0.0.10. The default administrative distance of EIGRP is 90, which is less than the default administrative distances of RIP and OSPF, which means that EIGRP routes will be preferred over RIP and OSPF routes. The metric is calculated using bandwidth, delay, reliability and load. By default, only bandwidth and delay are considered when calculating metric, while reliability and load are set to zero.

## EIGRP neighbors

Routers that run EIGRP must become neighbors before exchanging routing information. To dynamically discover neighbors on the directly attached networks, EIGRP routers use the multicast address of 224.0.0.10 to send Hello packets every couple of seconds. To become neighbors, the following parameters must match on both routers:

ASN (Autonomous System Number)

subnet

K values

(components of metric)

Let's examine each of these parameters in detail:

ASN (Autonomous System Number) - EIGRP uses the concept of autonomous systems. An autonomous system is simply a group of EIGRP-enabled routers that should become EIGRP neighbors and exchange routes. An AS is defined on a router by using the `router eigrp AS_NUMBER` global configuration command. In order to become neighbors, all routers must be configured with the same AS number.

subnet - interfaces on both routers must be in the same subnet.

K values - each router must be configured with the same K values used to calculate the metric. By default, the only parameters used to calculate an EIGRP metric are bandwidth and delay. If a network administrator wants to use other parameters (load and reliability) to calculate a metric, the change has to be made on both routers. Otherwise, the mismatched K-values will prevent routers to become neighbors.

Consider the following example:



We have a simple network of two routers. Both routers are running EIGRP, with the ASN of 1. Because the interfaces on both routers are in the same subnet and the K-values match, the routers will become neighbors:

```
R1(config-router)#
%DUAL-5-NBRCHANGE: IP-EIGRP 1: Neighbor 10.0.0.2 (FastEthernet0/0) is up: new
adjacency
```

The information about EIGRP neighbors are stored in the neighbor table. We can display the content of this table by running the `show ip eigrp neighbors` command:

```
R1#show ip eigrp neighbors
IP-EIGRP neighbors for process 1
H   Address          Interface           Hold Uptime      SRTT   RTO   Q   Seq
   (sec)              (ms)              (ms)   (ms)   Cnt  Num
0   10.0.0.2           Fa0/0               11    00-02:44    40    1000  0   1
```

Here is a brief description of each column

**H** - the sequential numbering of established neighbor adjacencies. The first neighbor will have a value of 0, the second neighbor a value of 1 and so on.

**Address** - the IP address of the EIGRP neighbor.

**Interface** - the interface on the local router on which the Hello packets were received.

**Hold (sec)** - the holddown timer. It specifies how long will EIGRP wait to hear from the neighbor before declaring it down.

**Uptime** - the time in hours:minutes: seconds since the local router first heard from the neighbor.

**SRTT (ms)** - Smooth round-trip time. The time it takes to send an EIGRP packet and receive an acknowledgment from the neighbor.

**RTO** - Retransmission timeout in milliseconds. This is the time that EIGRP will wait before retransmitting a packet from the retransmission queue to a neighbor.

**Q Cnt** - the number of EIGRP packets (Update, Query or Reply) in the queue that are awaiting transmission. Should be 0.

**Seq Num** - the sequence number of the last update, query, or reply packet that was received from the neighbor.

## EIGRP tables


Each EIGRP router uses three to store routing information:

**Neighbor table** - stores information about EIGRP neighbors. Before exchanging routes, routers need to establish a neighbor relationship. Information such as the IP address of the neighbor, the local interface on which the Hellos were received, the holddown timer, and Smooth round-trip time are kept in this table.

**Topology table** - stores routing information learned from neighbor routing tables. This table stores every EIGRP route inside the autonomous system. The topology table also holds the metrics for each of the listed EIGRP routes, the feasible successor and the successors.

**Routing table** - stores only the best routes to reach a remote network.

Consider the following example network:



We have a network of three routers, all running EIGRP. The routers have established neighbor relationships and the routes have been exchanged. Let's display the neighbor table on R1 using the `show ip eigrp neighbors` command:

```
R1#show ip eigrp neighbors
IP-EIGRP neighbors for process 1
H   Address          Interface          Hold Uptime      SRTT   RTO   Q   Seq
   (sec)              (sec)              (ms)              Cnt   Num
0   10.0.0.2           Fa0/0              12  00:01:09  40    1000  0   40
1   192.168.0.2       Se0/1/0            10  00:01:02  40    1000  0   16
```

You can see that R1 has established two neighbor relationships. The information such as the neighbor's IP address, the local interface, the holddown timer, are displayed.

Now let's display the topology table on R1 using the show ip eigrp topology command:

```
R1#show ip eigrp topology
IP-EIGRP Topology Table for AS 1

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

P 10.0.0.0/8, 1 successors, FD is 28160
   via Summary (28160/0), Null0
P 10.0.0.0/24, 1 successors, FD is 28160
   via Connected, FastEthernet0/0
P 172.16.0.0/16, 1 successors, FD is 20514560
   via 10.0.0.2 (20514560/20512000), FastEthernet0/0
   via 192.168.0.2 (21024000/20512000), Serial0/1/0
P 192.168.0.0/24, 1 successors, FD is 20512000
   via Connected, Serial0/1/0
```

We will explain the concept of successor and feasible successor in the next sections. For now, just remember that this table contains all routes learned by EIGRP. Notice how there are two possible routes to the 172.16.0.0/16 network, one via 10.0.0.2 and the other via 192.168.0.2.

The routing table looks like this:

```
R1#show ip route eigrp
 10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
D    10.0.0.0/8 is a summary, 00:01:44, Null0
D    172.16.0.0/16 [90/20514560] via 10.0.0.2, 00:01:49, FastEthernet0/0
```

Note that there is only one route to the 172.16.0.0/16 network - the best route, via 10.0.0.2.


## Reported and feasible distance explained

In EIGRP, a local router calculates the metric for each route, but also considers the next-hop router's metric for that same destination subnet. These metrics have their own names:

**Reported (advertised) distance (RD or AD)** - the metric advertised by the neighboring router for a specific route. This is the metric of the route used by the neighboring router to reach that specific destination network.

**Feasible distance (FD)** - the local router's metric of the best route to reach a specific network. The metric is calculated using the metric reported by the neighbor (the advertised distance) plus the metric to the neighbor reporting the route. The route with the lowest FD will be placed in the routing table.

We will try to explain these two metrics with an example:



We have a small network of two routers. Both routers are running EIGRP and the neighbor relationship has been established. R2 is directly connected to the 192.168.0.0/24 subnet and advertises that subnet using EIGRP to R1. The R2's metric to reach that network is 28160. R1 will use that metric to calculate its own metric to reach the 192.168.0.0/24 subnet. These values are stored in the R1's topology table:

```
R1#show ip eigrp topology
IP-EIGRP Topology Table for AS 1

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

P 10.0.0.0/24, 1 successors, FD is 28160
   via Connected, FastEthernet0/0
P 192.168.0.0/24, 1 successors, FD is 30720
   via 10.0.0.2 (30720/28160), FastEthernet0/0
```

Notice the numbers in parentheses (30720/28160). The first number (30720) represents the feasible distance (the metric) of R1 to reach the 192.168.0.0/24 subnet. The second number (28160) represents the advertised distance - the metric of R2 to reach that subnet.

### **Successor and feasible successor explained**

Two terms that you will often encounter in the EIGRP world are the successor and feasible successor. Here are the definitions of these two terms:

successor - the route with the best metric to reach a particular network. This route will be placed in the routing table.

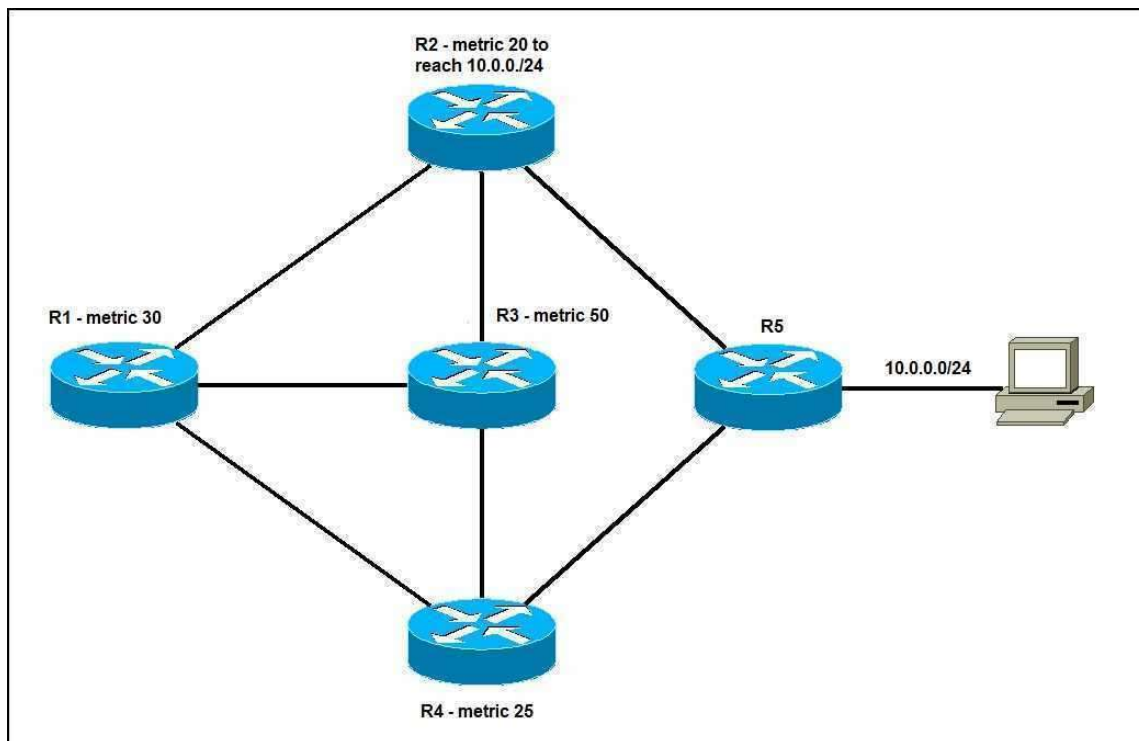
feasible successor - alternative routes to a particular network that can be used immediately if the currently best route (the successor) fails, without causing a routing loop. These routes are stored in the EIGRP topology table.

Not all alternative routes to a particular network will become feasible successors. In order for a route to become a feasible successor, the following condition must be met:

the neighbor's advertised distance (AD) for the route must be less than the successor's feasible distance (FD).

The definition above can be more easily understood with an example:





We have a network of five routers, all running EIGRP. R5 has advertised the 10.0.0.0/24 subnet. R1 has three paths to reach the 10.0.0.0/24 subnet:

- $R2 > R5$  - let's say that this is the best route (the successor route). This route will be placed in R1's routing table, with the metric of 30.
- $R3 > R2 > R5$  - for a route to become a feasible successor, the neighbor's advertised distance (AD) for the route must be less than the successor's feasible distance (FD). This is not the case here - R3 has advertised the metric of 50 to reach 10.0.0.0/24, which is greater than the feasible distance of R1 (30).
- $R4 > R5$  - this route will become a feasible successor route, since R4's advertised distance is less than the successor's feasible distance ( $25 < 30$ ). This route will be placed in R1's topology table and can be used immediately if the best route fails.

## EIGRP configuration

EIGRP can be configured using only two commands:


1. (config) router eigrp ASN - starts EIGRP on the router. In order to become EIGRP neighbors, routers must be configured with the same AS number. You can use any number between 1 and 65,535.
2. (config-router) network SUBNET [WILDCARD\_MASK] - a network command does two things:

for each interface matched by the network command, EIGRP tries to discover neighbors on that interface.

the subnet connected to the interface matched by the network command will be advertised.

By default, the network command accepts a classful network number as a parameter.

Consider the following example:



We have a simple network of two routers. Here is how we can configure EIGRP on both routers:

1. on R1, we need to enable EIGRP with the ASN of 1 and use the network command to match the interface directly connected to the 10.0.0.0/24 network:

```
R1 (config)#router eigrp 1
R1 (config-router)#network 10.0.0.0
```

2. on R2, we need to enable EIGRP with the ASN of 1, use the network command to match the interface directly connected to the 10.0.0.0/24 network, and advertise the 192.168.0.0/24 network:

```
R2 (config)#router eigrp 1
R2 (config-router)#network 10.0.0.0
R2 (config-router)#
%DUAL-5-NBRCHANGE: IP-EIGRP 1: Neighbor 10.0.0.1 (FastEthernet0/0) is up: new adjacency
R2 (config-router)#network 192.168.0.0
```

To confirm that the EIGRP neighbor relationship has been established, we can use the `show ip eigrp neighbors` command:

```
R1#show ip eigrp neighbors
IP-EIGRP neighbors for process 1
H   Address          Interface         Hold Uptime      SRTT   RTO   Q   Seq
   (sec)              (ms)              (ms)              Cnt   Num
0   10.0.0.2          Fa0/0             13 00:04:33      40    1000  0   2
```

In the picture above you can see that R1 has a single neighbor with the IP address of 10.0.0.2.

To verify that R1 has received a route to reach the 192.168.0.0/24 network, we can use the `show ip route eigrp` command:


```
R1#show ip route eigrp
   10.0.0.0/24 is subnetted, 1 subnets
D    192.168.0.0/24 [90/30720] via 10.0.0.2, 00:06:37, FastEthernet0/0
```

## Wildcard mask explained

Wildcard masks are used to specify a range of network addresses. They are usually used with routing protocols (such as EIGRP and OSPF) and access lists.

Just like a subnet mask, a wildcard mask is 32 bits long. It is a sort of inverted subnet masks, with the zero bits indicating that the corresponding bit position must match the same bit position in the IP address. The one bits indicate that the corresponding bit position does not have to match the bit position in the IP address.

The following example will help you understand the concept behind wildcard masks:



In the picture above you can see a network with three hosts and a router. The router is directly connected to three subnets. Let's say that we want to advertise only the 10.0.1.0/24 subnet in EIGRP. We can use the wildcard mask of 0.0.0.255 in the following network command to do this:

```
R1(config-router)#network 10.0.1.0 0.0.0.255
```

Why the wildcard mask of 0.0.0.255? To explain why, first we need to convert the IP address and wildcard mask to binary:

```
10.0.1.0 = 00001010.00000000.00000001.00000000
```

```
0.0.0.255 = 00000000.00000000.00000000.11111111
```

The zero bits of the wildcard mask have to match the same position in the IP address in order for the network to be included in the network command:

```
00001010.00000000.00000001.00000000
```

```
00000000.00000000.00000000.11111111
```

As you can see from the output above, the last octet doesn't have to match, because the wildcard mask bits are all ones. The first 24 bits have to match, because of the wildcard mask bits of all zeros. So, in this case, wildcard mask will match all addresses that begins with 10.0.1.x (10.0.1.0 - 10.0.1.255). In our case, we have only a single network that will be matched - 10.0.1.0/24.

What if we want to include both 10.0.0.0/24 and 10.0.1.0/24 subnets? Well, we need to use the wildcard mask of 0.0.1.255. Here is why:

10.0.0.0 = 00001010.00000000.00000000.00000000

10.0.1.0 = 00001010.00000000.00000001.00000000

0.0.1.255 = 00000000.00000000.00000001.11111111

From the output above you can see that, with the wildcard mask of 0.0.1.255, only the first 23 bits have to match. That means that all addresses in the range of 10.0.0.0 – 10.0.1.255 will be matched. So, in our case, both IP addresses have been matched.


The wildcard mask of all zeros (0.0.0.0) means that the entire IP address have to match in order for a statement to execute. For example, if we want to match only the IP address of 192.168.0.1, the command to use is `network 192.168.0.1 0.0.0.0`. A wildcard mask of all ones (255.255.255.255) means that no bits have to match. This basically means that all addresses will be matched.

## EIGRP and wildcard masks

The network command in EIGRP uses a classful network as the parameter by default, which means that all interfaces inside the classful network will participate in the EIGRP process. We can enable EIGRP only for specific networks using wildcard masks. The syntax of the command is:

```
(router-eigrp)#network IP_ADDRESS WILDCARD_MASK
```

Consider the following example:



In the picture above you can see a network with three hosts and a router. The router is directly connected to three subnets. Let's say that we want to advertise only the 10.0.0.0/24 subnet in EIGRP. We can use the wildcard mask of 0.0.0.255 to do this:

```
R1(config)#router eigrp 1
R1(config-router)#network 10.0.0.0 0.0.0.255
R1(config-router)#
```

Using the show ip protocols command, we can verify that the network 10.0.0.0/24 is included in EIGRP:

```
R1#show ip protocols

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 metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  Redistributing: eigrp 1
    Automatic network summarization is in effect
  Automatic address summarization:
    Maximum path: 4
  Routing for Networks:
    10.0.0.0/24
  Routing Information Sources:
    Gateway         Distance         Last Update
  Distance: internal 90 external 170

R1#
```

## Reliable Transport Protocol (RTP)

EIGRP sends messages without UDP or TCP; instead, a Cisco's protocol called Reliable Transport Protocol (RTP) is used for communication between EIGRP-speaking routers. As the name implies, reliability is a key feature of this protocol, and it is designed to enable quick delivery of updates and tracking of data reception.

Five different packets types are used by EIGRP:

- Update - contains route information. When these are sent in response to metric or topology changes, reliable multicasts are used. In the event that only one router needs an update, like when a new neighbor is discovered, it's sent via unicasts.
- Query - a request for specific routes and always uses the reliable multicast method. Routers send queries when they realize they've lost the path to a particular network and are searching for alternative paths.

- Reply - sent in response to a query via the unicast method. Replies can include a specific route to the queried destination or declare that there's no known route.
- Hello - used to discover EIGRP neighbors. It is sent via unreliable multicast (no acknowledgment is required).
- Acknowledgment (ACK) - sent in response to an update and is always unicast. ACKs are not sent reliably.

The acronym RTP is also used for a different, unrelated protocol - Real-time Transport Protocol (RTP), used for VoIP communication.

### **Diffusing Update Algorithm (DUAL)**

EIGRP uses Diffusing Update Algorithm (DUAL) for selecting and maintaining the best route to each remote network. DUAL is also used for the following tasks:

- discover a backup route if there's one available.
- support for variable length subnet masks (VLSMs).
- perform dynamic route recoveries.
- query neighbors for unknown alternate routes.
- send out queries for an alternate route.

EIGRP stores all routes advertised by all neighbors. The metric of these routes is used by DUAL to select efficient loop free paths. DUAL selects routes that will be inserted into a routing table. If a route fails, and there is no feasible successor, DUAL chooses a replacement route, which usually takes a second or two.

The following requirements must be met for DUAL to work properly:

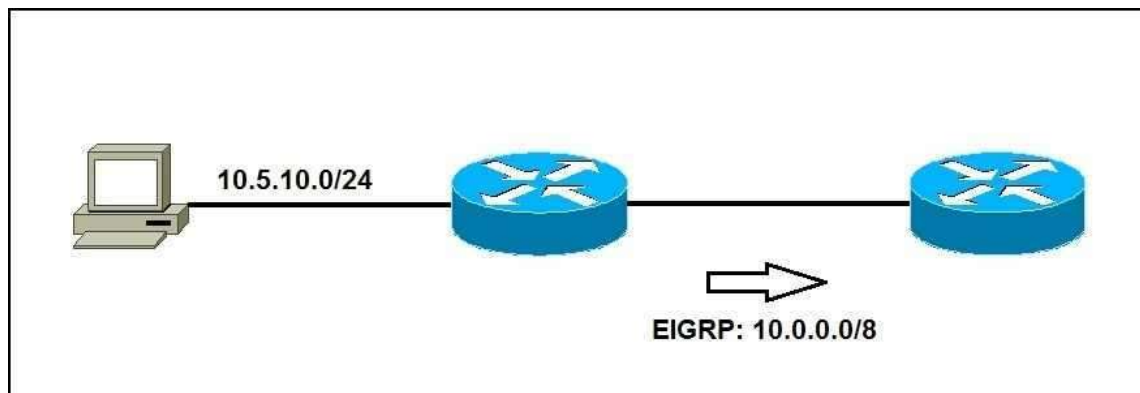
- EIGRP neighbors must be discovered.
- all transmitted EIGRP messages should be received correctly.
- all changes and messages should be processed in the order in which they're detected.



## EIGRP auto-summary

In EIGRP, the auto-summary feature is enabled by default. This means that EIGRP will automatically summarize networks at their classful boundaries. The auto-summary feature can be useful, but it can also cause problems with discontinuous networks (which are two or more subnetworks of a classful network connected together by different classful networks).

Consider the following example:



In the example above we have a network of two routers and one host. Both routers are running EIGRP and have established the neighbor relationship. The router R1 has a directly connected subnet 10.5.10.0/24 that it advertises to R2. Because the auto-summary feature is turned on, R1 will summarize the subnet 10.5.10.0/24 and send the classful route 10.0.0.0/8 to R2.

Here is the routing table on R2:

```
R2#show ip route
Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       i - IS-IS, 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

D    10.0.0.0/8 [90/30720] via 192.168.0.1, 00:01:46, FastEthernet0/1
C    192.168.0.0/24 is directly connected, FastEthernet0/1
```

Note that R2 has a route to the classful network 10.0.0.0/8. R2 will therefore send all packets destined for any IP address inside the 10.0.0.0 - 10.255.255.255 range to R1. This can cause problems with discontinuous networks. That is why the auto-summary feature is usually turned off with the no auto-summary command:

```
R1(config-router)#no auto-summary
R1(config-router)#
%DUAL-5-NBRCHANGE: IP-EIGRP 1: Neighbor 192.168.0.2
(FastEthernet0/1) resync: summary configured
```

Let's take a look at the routing table on R2:

```
R2#show ip route
Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       i - IS-IS, 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

 10.0.0.0/24 is subnetted, 1 subnets
D    10.5.10.0 [90/30720] via 192.168.0.1, 00:01:21, FastEthernet0/1
C    192.168.0.0/24 is directly connected, FastEthernet0/1
```

You can see in the picture above that R2 now has the more specific route. Note that the EIGRP neighbor relationship has to be re-established after the no auto-summary command was entered.

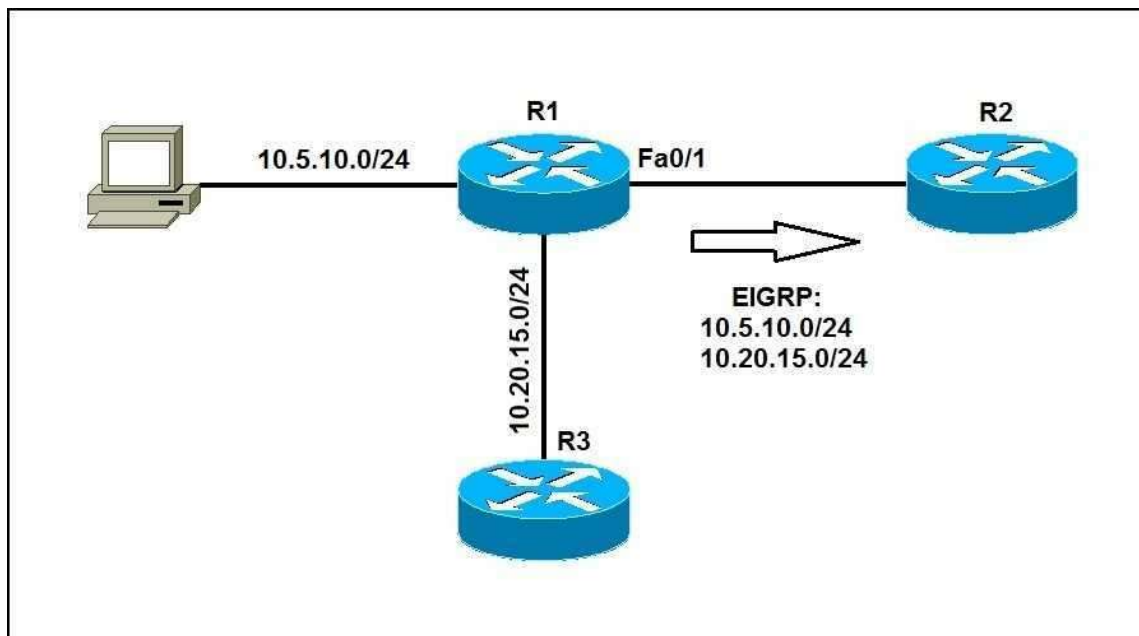
## EIGRP manual summarization

The manual summarization is a process of creating a summary route that will be used to represent multiple routes and can be used to reduce the sizes of routing tables in a network. EIGRP, unlike some other routing protocols such as OSPF, supports the manual summarization on any router within a network.

The manual summarization in EIGRP is configured on the per-interface basis. The syntax of the command is:

```
(config-if)ip summary-address eigrp ASN SUMMARY_ADDRESS
SUBNET_MASK
```

Consider the following example:



In the example above we have a network of three routers running EIGRP. R1 has two directly connected subnets: 10.5.10.0/24 and 10.20.15.0/24. Assuming that no automatic summarization is enabled, R1 will advertise these two networks to R2. Here is the routing table on R2:

```
R2#show ip route
Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       i - IS-IS, 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

 10.0.0.0/24 is subnetted, 2 subnets
D       10.5.10.0 [90/30720] via 192.168.0.1, 00:01:53, FastEthernet0/1
D       10.20.15.0 [90/284160] via 192.168.0.1, 00:00:18, FastEthernet0/1
C       192.168.0.0/24 is directly connected, FastEthernet0/1
```

Note that R2 has two routes to reach the 10.5.10.0/24 and 10.20.15.0/24 networks. To reduce the size of the R2's routing table, we can advertise a summary route that will include both subnets. We will use the following command on the Fa0/1 interface on R1:

```
R1(config-if)#ip summary-address eigrp 1 10.0.0.0 ?
  A.B.C.D IP network mask
R1(config-if)#ip summary-address eigrp 1 10.0.0.0 255.0.0.0
R1(config-if)#
%DUAL-5-NBRCHANGE: IP-EIGRP 1: Neighbor 192.168.0.2 (FastEthernet0/1) is up: new adjacency
```

R1 should now send only a single route (10.0.0.0) for both subnets. We can verify that by running the show ip route command on R2:

```
R2#show ip route
Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       i - IS-IS, 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

D    10.0.0.0/8 [90/30720] via 192.168.0.1, 00:01:28, FastEthernet0/1
C    192.168.0.0/24 is directly connected, FastEthernet0/1
```

Notice how R2 now has only a single route to reach both subnets.

In the example above, the ip summary command included two subnets directly connected to R1, but also plenty of other addresses that are not in these subnets. The range of the summarized addresses is 10.0.0.0 – 10.255.255.255, so R2 will send all packets destined for these addresses to R1. This could cause some problems if these addresses exist somewhere else in the network.