

## CCNP BSCI 642-801 Exam Notes

### Enhanced IGRP (EIGRP)

#### 1. EIGRP Basics

1.1 **Enhanced IGRP** (EIGRP) is a **balanced hybrid** routing protocol. It has some features similar to distance vector routing protocols and some features similar to link-state routing protocols. It has the following characteristics:

##### Routing Protocol Type

- Cisco-proprietary protocol.
- Internal routing protocol.
- Classless routing protocol, i.e. support VLSM.
- EIGRP implements **protocol-dependent modules** for IP, IPX and AppleTalk. They support routing for the specific Layer 3 protocols.

##### Metric

- Composite metric is used, which is computed using the IGRP metric formula (details can be found in Chapter 3), and is then multiplied by 256 to achieve a finer metric granularity, i.e. Default Metric = (Bw + Delay) x 256, where:
  - Bw = [10 Gbps / bandwidth of the slowest link along the path].
  - Delay = Sum of delays of all the links along the path (in 10s of microseconds).
- The maximum allowable hop count is 255 (100 by default).
- **Diffusing Update Algorithm (DUAL)** is used to calculate the shortest paths and alternative routes. It exchanges more topology information than distance vector routing protocols, but less than link-state routing protocols.
- Load-balancing using equal or unequal cost paths is supported.

##### Convergence

- Short convergence time (several seconds). EIGRP maintains backup routes (known as feasible successors) in its topology database. When a route fails, EIGRP can use its feasible successor immediately. Therefore, the convergence time can be very short.

##### Neighbor Relationship

- A **neighbor** is a directly connected router running EIGRP and with the same AS number.
- A router sends hello messages to its neighbors periodically when the **hello timer** expires. The messages are used for neighbor discovery and maintenance.

- A neighbor is assumed dead if no hello message from that neighbor is received before the **hold timer** expires (default = 3 times of the hello timer value). All routes learned from that neighbor will also be deleted.

### Routing updates

- **Reliable Transport Protocol (RTP)**, a Cisco proprietary protocol, is used for transmitting EIGRP messages between routers. It can provide reliable packet delivery when necessary, by using sequence number, acknowledgement, and retransmission with unicast (a packet will be retransmitted up to 16 times if no acknowledgement is received, before declaring the neighbor router is dead).
- Both unicast and multicast traffic are used for inter-router communication. The multicast address used by EIGRP is 224.0.0.10.
- Partial routing updates are sent between neighbors when there are link-state changes. Full routing updates (the entire topology table) are not sent regularly. They are sent only when two neighbors initiate communication.
- Authentication of the source of routing updates is supported.

### 1.2 EIGRP consists of **four components** :

- Protocol-dependent modules
- Diffusing Update Algorithm (DUAL)
- Neighbor discovery/recovery
- Reliable Transport Protocol (RTP)

### 1.3 EIGRP classifies the next hop routers of different paths from the router to a destination network as follows:

- **Successor**
  - The next hop router of the lowest cost path to the destination network.
  - It is placed in the routing table.
  - **Feasible distance (FD)** refers to the metric of the lowest cost path from the router to the destination network.
- **Feasible successors (FS)**
  - The next hop routers of the paths that meet the feasible condition.
  - **Feasible condition (FC)** refers to the condition that the metric of the path from the next hop router to the destination network is lower than the router's feasible distance. It implies that the next hop router is downstream from the router, and the path from the next hop router to the destination network does not pass through the router. Therefore, **routing loops** can be prevented with the FC checking.
- **Possibilities**
  - The next hop routers of the paths that cannot meet the feasible condition.

1.4 EIGRP classifies routes into the following three types:

- **Internal routes**

- Routes that are learned within this EIGRP AS.
- The default administrative distance of EIGRP internal routes is **90**.
- In the output of the command "#show ip route", EIGRP internal routes are represented by the code '**D**'. For example:

```
D          192.168.1.0/24 [90/2674123] via 10.1.1.6, 0:01:00, Ethernet0
```

- **External routes**

- Routes that are redistributed into this EIGRP AS from another routing process.
- IP EIGRP automatically redistributes routes with IP IGRP of the same AS.
- IPX EIGRP automatically redistributes routes with IPX RIP and NLSP.
- AppleTalk EIGRP automatically redistributes routes with AppleTalk RTMP.
- The default administrative distance of EIGRP external routes is **170**.
- In the output of the command "#show ip route", EIGRP external routes are represented by the code '**D EX**'. For example:

```
D EX      172.31.1.0/24 [170/3110962] via 10.1.3.5, 0:01:31, Ethernet1
```

- **Summary routes**

- Summary routes advertised by a router performing summarization on more-specific EIGRP routes.
- EIGRP can perform route summarization at any interface on any router (rather than on an area border router only for OSPF).
- Auto-summarization (i.e. automatic route summarization at classful network boundary) is enabled by default. To manually configure summarization, auto-summarization should be disabled first.
- A router performing route summarization will automatically create a static route to the null interface (null0) for the summarized address. This ensures that if the router receives a packet destined for an address within the summarized address range, and for which it has no route (e.g. 172.31.111.0 in the example below), it will forward the packet to the null interface (i.e. drop the packet), rather than using some other routes such as the default route for routing the packet.
- An example of the output of the command "#show ip route" on a router performing route summarization is as follows:

```
172.31.0.0/16 is subnetted, 2 subnets
C          172.31.100.0/24 is directly connected, Ethernet0
C          172.31.101.0/24 is directly connected, Ethernet1
D          172.31.0.0/16 is a summary, 0:01:19, Null0
```

- The metric of a summary route is the best metric among the summarized routes.
- The default administrative distance of EIGRP summary routes is **5**.

- In the output of the command "#show ip route" on a router received an EIGRP summary route, the route is represented by the code 'D', same as an EIGRP internal route.

## 2. EIGRP Operation

2.1 EIGRP makes use of the following three tables in its operation:

- **Neighbor table.**
- **Topology table.**
- **Routing table.**

Each Layer 3 protocol supported by EIGRP (i.e. IP, IPX, and AppleTalk) has its own set of tables. Therefore, there are 9 active EIGRP tables on a router if EIGRP is configured for all the three Layer 3 protocols.

2.2 **Neighbor table** has the following characteristics:

- A neighbor table contains an entry for each neighbor.
- Each entry in the table (IP) contains the following information of a neighbor:
  - IP address of the neighbor.
  - Interface of the router connected to the neighbor.
  - The neighbor's uptime, or how long the neighbor was added to the neighbor table.
  - Hold-time - how long the router waits before considering the neighbor is down.
  - Smooth round-trip time (SRTT) - the time period for a packet to be sent to the neighbor and a reply to be received. It is used to calculate the RTO.
  - Retransmission timeout (RTO) - the time period that the router will wait for an acknowledgement from the neighbor before re-transmitting a packet.
  - The sequence number of the last EIGRP packet received from the neighbor.
  - Number of EIGRP packets that the router is waiting to send to the neighbor.
  - etc.

2.3 **Topology table** has the following characteristics:

- It contains an entry for every destination network learned through EIGRP.
- Each table entry contains the following information of a destination:
  - IP address prefix and netmask of the destination network.
  - Feasible distance.
  - Successors.
  - For each path to the destination network:
    - IP address of the next hop router (i.e. either a successor, feasible successor, or possibility).

- **Advertised Distance (AD)**, also known as **Reported Distance (RD)**, i.e. the metric of the path from the next hop router to the destination network (metric advertised by the next hop router).
- Outgoing interface connected to the next hop router.
- Status of the route to the destination network - **Passive** (i.e. healthy) or **Active** (i.e. the router is querying its neighbors to find a path to the network).
- Whether the router has sent an update / query / reply packet about the route to its neighbors, or is waiting for a reply.

2.4 **Routing table** has the following characteristics:

- It is built from the topology table after DUAL has been run.
- The table contains the following information for each destination:
  - IP address prefix and netmask of the destination network.
  - Successors for the destination network. Up to six successors (four by default) can be placed in the routing table for each destination network. This number is configurable.

2.5 There are five type of EIGRP packets:

- **Hello** packet
  - It is used for discovering neighbors and maintaining neighbor relationship.
  - It is sent to neighbors periodically (**hello interval**). The default hello interval is:
    - Non-broadcast multi-access (NBMA) network with bandwidth  $\leq$  1.5 Mbps (e.g. T1 frame relay, ISDN BRI, etc.) - **60 seconds**.
    - Other networks (e.g. Ethernet, point-to-point serial links, high speed frame relay, etc.) - **5 seconds**.
  - A neighbor is assumed dead if no hello message from that neighbor is received before the **hold timer** expires (default = 3 times of the hello timer value). All routes learned from that neighbor will also be deleted.
  - Neighbor routers can use different hello and hold-time intervals. A router informs the neighbors its hold-time interval through hello packets.
  - It is sent as **multicast**.
  - Acknowledgement is not required (i.e. **unreliable** delivery).
- **Update** packet
  - It is used for sending routing information to neighbors as follows:
    - **Full routing updates** (i.e. the whole topology table) during initialization of the routing process.
    - **Incremental routing updates** (i.e. routing updates about the paths that have been changed) when there is a change in the network topology or metric.
  - It is sent to the relevant routers as **unicast** or **multicast**.

- Acknowledgement is required (i.e. **reliable** delivery).
- **Query** packet
  - It is used for querying neighbors if they have feasible successors for a destination network.
  - It is sent during a diffusing computation (to be explained later in this Section).
  - It is sent as **multicast**.
  - Acknowledgement is required (i.e. **reliable** delivery).
- **Reply** packet
  - It is used for replying a query with the requested routing information (e.g. information of the best path to the destination network, or destination unreachable).
  - It is sent as **unicast**.
  - Acknowledgement is required (i.e. **reliable** delivery).
- **Acknowledgement (ACK)** packet
  - It is used for acknowledging the receipt of an update / query / reply packet.
  - It is sent as **unicast**.
  - Acknowledgement is not required (i.e. **unreliable** delivery).

2.6 **Diffusing Update Algorithm (DUAL)** is used to determine the successors and feasible successors for a destination network. A router will run DUAL when:

- There is a change in the status or metric of a directly connected link.
- An update is received from a neighbor (through update / query / reply packet).

2.7 **DUAL** works as follows:

#### 1. **Local Computation**

- The router recalculates the distance to the destination network for each feasible successor.
- If the lowest distance is smaller than the FD:
  - The corresponding feasible successor will become the new successor.
  - The FD will be revised to the lowest distance value.
  - Routing updates will be sent to the neighbors.
- When the router is performing a local computation, the route remains in the **passive state**.
- If a feasible successor for the route cannot be found in the topology table, the router will run the diffusing algorithm.

#### 2. **Diffusing Computation**

- The state of the route is changed from passive to **active**.
  - The router sends queries to its neighbors to ask for routes to the destination network.

- When a neighbor receives the query, it will perform a local computation. If one or more feasible successors are found, the neighbor will send a reply back to the querying router with the best route. Otherwise, it will perform a diffusing computation. Therefore, the query will propagate or diffuse until a reply is received. If no feasible successor can be found eventually, the neighbor returns an unreachable message to the querying router.
- The router waits for a reply from every neighbor it queried until the **active timer** (default = 3 minutes) expires. The active timer is designed to prevent a route from being permanently active.
- The router updates its topology table with the replies received. The metric for each feasible successor is calculated based on the cost advertised by the neighbor and the cost of the link to the neighbor.
- The feasible successor with the lowest metric is selected as the successor, and the feasible distance is updated accordingly.
- The state of the route is changed from active to **passive**.

### 3. Scalability & Performance

- 3.1 A route will be marked as **stuck-in-active** (SIA) if one or more neighbors do not reply to a query for the route before the active timer expires. The unresponsive neighbors will be removed from the neighbor table. The routes with these neighbors as the next hop router will also be deleted.
- 3.2 SIAs should not occur in a stable and well-designed network. If a neighbor router or a link has problem, it should be detected by the expiry of the hold timer, long before the expiry of the active timer. Therefore, such problems will not cause SIAs normally.
- 3.3 However, SIAs may occur because of looping of queries, or because of heavily congested / low-bandwidth links and/or overloaded routers in a large network. Therefore, SIAs may lead to flushing of valid neighbors and routes from the neighbor table and topology table, which affect the stability of the network.
- 3.4 Route queries can propagate throughout the entire network if feasible successors are not found. It causes scalability limitation on the network size and may even lead to SIAs. The problem is more complicated if the network contains redundant paths, such that the queries may loop back. To prevent these problems, the following methods can be used to limit the scope of query range, i.e. **query scoping**:
- **Route summarization** If a subnet is hidden in route summarization, the router

will reply that the subnet is unreachable and the querying router will purge the route from its database. It is the most effective method for limiting the scope of query range.

- **EIGRP Stub Feature.** In a hub-and-spoke network, by configuring the remote routers as stubs, the hub router can answer queries immediately without propagating the queries to the remote sites.
- **Autonomous System Boundary.** Queries do not propagate into other ASs unless redistribution with the ASs has been configured.
- **Route Filtering.** Filters for routing updates (details can be found in Chapter 7) can be used to mark query replies as unreachable.

3.5 The **EIGRP Stub Routing** feature can improve network stability (by preventing SIAs) and reduce resources required on stub routers. It has the following characteristics:

- It is typically used in a **hub-and-spoke** network (similar to ODR).
- It is suitable for **stub routers** (stub networks' routers which connected to the distribution router or the hub) with **limited resources**.
- Since the stub routers have no EIGRP neighbor router other than the distribution router (hub), they must forward all non-local traffic to the distribution router. Therefore, the routers do not need to maintain full topology tables and routing tables, but only a **default route** to the distribution router **for non-local traffic**.
- With the stub routing feature enabled, the stub routers direct all non-local traffic to the distribution router. Besides, the distribution router will not query the stub routers for any routes (i.e. **query scoping**).

3.6 EIGRP controls the **bandwidth usage** of a link for transmitting EIGRP traffic. This feature is known as **padding**. It works as follows:

- By default, **at most 50%** of the configured bandwidth of an interface or subinterface is used for transmitting EIGRP traffic. It prevents the routing traffic from saturating the link. Otherwise, data packets and EIGRP packets may be dropped.
- For a **LAN interface**, the default bandwidth value reflects the correct bandwidth. No additional manual configuration is required normally.
- For a **serial interface**, the default bandwidth value is always 1544 or 1.544 Mbps. Therefore, the bandwidth value may need to be configured manually using the command "(config-if)#bandwidth".
- For a **point-to-point subinterface**, the default bandwidth value is also 1.544 Mbps. Therefore the following manual configuration is required:
  - Configure the bandwidth value as the **committed information rate (CIR)** of the subinterface (e.g. frame relay DLCI).
  - If the total bandwidth (for all subinterfaces of the physical interface) that can be

used by EIGRP traffic exceeds the **access speed** of the physical interface, reduce the percentage of the configured bandwidth that can be used by EIGRP traffic using the command "(config-if)#ip bandwidth-percent eigrp".

- For a **multi-point subinterface**, the bandwidth value for each virtual circuit (VC) is the subinterface bandwidth / number of virtual circuits. If the VCs have different CIR values, the VCs with lower CIR may be saturated by EIGRP traffic. To solve this problem, the following approaches can be used:
  - Manually configure the bandwidth value of the subinterface as the **minimum CIR x number of virtual circuits** (this approach may cause inefficient usage of the bandwidth for VCs with higher CIR); or
  - Use **point-to-point configuration** instead of multi-point configuration, so that a different bandwidth value can be configured for each subinterface. This approach is the **recommended solution**.
- The bandwidth allocated to EIGRP traffic for a VC should be the same on both ends of the circuit.

3.7 EIGRP does not require any special configuration for it to work on different types of Layer 2 networks, other than the bandwidth usage control mentioned above. Other routing protocols, such as OSPF, require different configurations for different types of Layer 2 networks (e.g. Ethernet and Frame Relay).

## 4. Configuring EIGRP

4.1 The following commands are used for **basic EIGRP configuration**:

Command	Description
(config)#[no] router eigrp <AS number>	<ul style="list-style-type: none"> <li>• Turn on the EIGRP routing process and enter the router configuration mode "(config-router)#" (or turn off the EIGRP routing process).</li> <li>• All routers must use the same AS number in order to share their routing information. The AS number can be assigned arbitrarily within the range 1 - 65535.</li> <li>• No EIGRP routing process is defined by default.</li> </ul>
(config-router)#[no] network <ip address> [<mask>]	<ul style="list-style-type: none"> <li>• Add (or remove) the specified directly connected network to the EIGRP AS.</li> <li>• The router will then establish neighbor relationship with other EIGRP routers through the local interfaces that are configured with</li> </ul>

	<p>addresses under the specified network.</p> <ul style="list-style-type: none"> <li>• Optionally, a netmask or wildcard mask can be used for specifying the interfaces to be added to the routing process. If it is not specified, classful network address is used.</li> <li>• There is no limit to the number of "(config-router)#network" commands that can be issued on the router.</li> </ul>
<pre>(config-router)#[no]   passive-interface     &lt;int-type&gt; &lt;int-num&gt;  (config-router)#passive-interface   default</pre>	<ul style="list-style-type: none"> <li>• Disable (or re-enable) sending of hello messages out the specified interface. Therefore, neighbor relationship cannot be established, and no routing traffic will be sent or received.</li> <li>• An interface is non-passive by default.</li> <li>• The keyword "default" sets all interfaces as passive by default.</li> </ul>

#### 4.2 An example of **configuring EIGRP** is as follows:

<pre>Router(config)#<b>interface</b> ethernet 1 Router(config-if)#<b>ip address</b> 192.168.1.1   255.255.255.0 Router(config-if)#<b>no shutdown</b> Router(config-if)#<b>exit</b></pre>	<ul style="list-style-type: none"> <li>• <i>Define the IP address of the interface and the network directly connected.</i></li> </ul>
<pre>Router(config)#<b>interface</b> ethernet 2 Router(config-if)#<b>ip address</b> 192.168.2.1   255.255.255.0 Router(config-if)#<b>no shutdown</b> Router(config-if)#<b>exit</b></pre>	<ul style="list-style-type: none"> <li>• <i>Define the IP address of another interface and the network directly connected.</i></li> </ul>
<pre>Router(config)#<b>router</b> eigrp 1 Router(config-router)#<b>network</b> 192.168.0.0   0.0.255.255 Router(config-router)#<b>exit</b></pre>	<ul style="list-style-type: none"> <li>• <i>Enable EIGRP routing for the two interfaces.</i></li> </ul>

#### 4.3 The following commands are used for **configuring load-balancing**:

Command	Description
<pre>(config-router)#maximum-paths   &lt;number of paths&gt;</pre>	<ul style="list-style-type: none"> <li>• Set the maximum number (valid values are 1-6) of equal or unequal-cost routes for a destination network that will be kept in the routing table (or reset to the default value).</li> </ul>
<pre>(config-router)#no maximum-paths</pre>	<ul style="list-style-type: none"> <li>• The default value is 4.</li> </ul>

<pre>(config-router)#variance &lt;multiplier&gt;  (config-router)#no variance</pre>	<ul style="list-style-type: none"> <li>• Set the variance multiplier (or reset to the default value).</li> <li>• Any route with the metric lower than or equal to the lowest metric x variance will be added into the routing table.</li> <li>• If the "maximum-paths" has been reached, only the lowest metric routes are kept in the routing table.</li> <li>• The default variance multiplier is 1, i.e. only the routes with the path cost equals the lowest metric will be put in the routing table.</li> </ul>
<pre>(config-router)#[no] traffic-share min [across-interfaces]</pre>	<ul style="list-style-type: none"> <li>• Configure the router to use the route with the lowest metric when there are multiple routes to the same destination network in the routing table.</li> <li>• It means that the load-balancing function using unequal-cost routes is turned off. However, if the route with the lowest metric fails, the route with the next lowest metric can still be used immediately without waiting for convergence of the network.</li> <li>• The keyword "across-interfaces" allows you to configure multi-interface load splitting on different interfaces with equal cost paths.</li> <li>• The "no" form of the command disables this function.</li> </ul>
<pre>(config-router)#[no] traffic-share balanced</pre>	<ul style="list-style-type: none"> <li>• Configure the router to distribute traffic among all the routes for the same destination network in the routing table.</li> <li>• Routes with higher metrics are less-preferable and get less traffic.</li> <li>• The "no" form of the command disables this function.</li> <li>• This function is enabled by default. However, since the default variance multiplier is 1, the default behavior is load-balancing over equal-cost paths.</li> </ul>

4.4 An example of **configuring load-balancing** is as follows:

<pre>Router(config)#<b>router</b> eigrp 1 Router(config-router)#<b>network</b> 192.168.0.0 0.0.255.255 Router(config-router)#<b>maximum-paths</b> 6 Router(config-router)#<b>variance</b> 2</pre>	<ul style="list-style-type: none"> <li>• <i>Set the maximum number of routes per destination as 6.</i></li> <li>• <i>For each destination, install routes</i></li> </ul>
---	--

Router(config-router)# <b>traffic-share</b> <b>balanced</b> Router(config-router)# <b>exit</b>	<p>with metrics <math>\leq</math> the lowest metric <math>\times 2</math>.</p> <ul style="list-style-type: none"> <li>• Enable load-balancing over unequal-cost routes. This function is enabled by default. Therefore, this command is actually not necessary.</li> </ul>
--	--

4.5 The following commands are used for **configuring route summarization**:

Command	Description
(config-router)#[no] auto-summary	<ul style="list-style-type: none"> <li>• Enable (or disable) auto-summarization.</li> <li>• Prior to IOS version 12.2, auto-summarization is enabled by default for EIGRP. On newer versions IOS, it is disabled by default.</li> </ul>
(config-if)#[no] ip summary-address eigrp <AS number> <network-address> <netmask> [<admin-distance>]	<ul style="list-style-type: none"> <li>• Configure EIGRP to summarize the routes for the subnets covered by the specified network address and netmask, and advertise the summary route out the interface (or disable the configuration).</li> <li>• A route to the null interface (null0) for the summarized address will also be created.</li> <li>• Optionally, the administrative distance of the summary route can be specified.</li> </ul>
(config-if)#no ip unreachable	<ul style="list-style-type: none"> <li>• Disable the generation of ICMP unreachable messages.</li> <li>• By default, the router generates an ICMP unreachable message to the source of a packet dropped by the null interface. With this command, the packet will be dropped silently.</li> </ul>

4.6 An example of **configuring route summarization** is as follows:

<pre>Router(config)#<b>interface</b> ethernet 0 Router(config-if)#<b>ip address</b> 172.16.1.1 255.255.255.0 Router(config-if)#<b>exit</b> Router(config)#<b>interface</b> ethernet 1 Router(config-if)#<b>ip address</b> 172.16.2.1 255.255.255.0  Router(config)#<b>router</b> eigrp 1 Router(config-router)#<b>no auto-summary</b> Router(config-router)#<b>network</b> 172.16.1.0 Router(config-router)#<b>network</b> 172.16.2.0</pre>	<ul style="list-style-type: none"> <li>• <i>Configure interface Ethernet 0.</i></li> <li>• <i>Configure interface Ethernet 1.</i></li> <li>• <i>Disable auto-summarization.</i></li> <li>• <i>Enable EIGRP routing for the two interfaces.</i></li> </ul>
---	---

<pre>Router(config-router)#exit  Router(config)#interface serial 0 Router(config-if)#ip address 10.5.12.1 255.255.255.0 Router(config-if)#ip summary-address eigrp 1 172.16.0.0 255.255.0.0 Router(config-if)#exit</pre>	<ul style="list-style-type: none"> <li>• <i>Configure interface Serial 0.</i></li> <li>• <i>Advertise a summary route for 172.16.0.0/16, instead of the two more specific routes 172.16.1.0/24 and 172.16.2.0/24, out the interface Serial 0.</i></li> </ul>
--	--

4.7 The following commands are used for **tuning EIGRP parameters** and **configuring EIGRP stub routing**:

Command	Description
<pre>(config-if)#[no] ip hello-interval eigrp &lt;AS-number&gt; &lt;hello interval in seconds&gt;</pre>	<ul style="list-style-type: none"> <li>• Configure the hello interval for the specified EIGRP AS on the interface (or reset to the default value).</li> <li>• The smaller the hello interval (and the hold-time interval), the faster the network converges. However, more routing traffic will be generated.</li> <li>• The default hello interval is 60 seconds for NBMA networks with bandwidth &lt;= 1.5 Mbps, and 5 seconds for other networks.</li> </ul>
<pre>(config-if)#[no] ip hold-time eigrp &lt;AS number&gt; &lt;hold-time interval in seconds&gt;</pre>	<ul style="list-style-type: none"> <li>• Configure the hold-time interval for the specified EIGRP AS on the interface (or reset to the default value).</li> <li>• The hold-time interval should at least equal three times the hello interval.</li> <li>• The default hold-time interval is 180 seconds for NBMA networks with bandwidth &lt;= 1.5 Mbps, and 15 seconds for other networks.</li> </ul>
<pre>(config-if)#[no] ip bandwidth-percent eigrp &lt;AS number&gt; &lt;percent&gt;</pre>	<ul style="list-style-type: none"> <li>• Configure the maximum percentage of bandwidth that may be used by the specified EIGRP AS on the interface (or reset to the default value).</li> <li>• The default value is 50%.</li> </ul>
<pre>(config-router)#[no] eigrp stub [receive-only   connected   static   summary]</pre>	<ul style="list-style-type: none"> <li>• Configure the router as an EIGRP stub router (or disable the configuration).</li> <li>• The keyword "receive-only" restricts the router from advertising any of its routes. This keyword cannot be used together with any other keyword.</li> <li>• The keyword "connected" permits the router to advertise connected routes. It is enabled by default.</li> </ul>

	<ul style="list-style-type: none"> <li>• The keyword "static" permits the router to advertise redistributed static routes. It is disabled by default.</li> <li>• The keyword "summary" permits the router to advertise summary routes. It is enabled by default.</li> <li>• The keywords "connected", "static", and "summary" can be used in any combination. If the command is issued with any of these three keywords, the router will be permitted to advertise the corresponding types of routes, but restricted from advertising the types of routes not specified.</li> <li>• No special configuration is required on the distribution router.</li> </ul>
--	---

4.8 The following commands are used for **debugging EIGRP information**:

Command	Description
#[no] debug eigrp packet	<ul style="list-style-type: none"> <li>• Enable (or disable) debugging of transmission and receipt of EIGRP packets.</li> </ul>
#[no] debug eigrp neighbor	<ul style="list-style-type: none"> <li>• Enable (or disable) debugging of EIGRP neighbor information.</li> </ul>
#[no] debug ip eigrp	<ul style="list-style-type: none"> <li>• Enable (or disable) debugging of EIGRP routing updates (including routes and metrics information).</li> </ul>
#[no] debug ip eigrp notifications	<ul style="list-style-type: none"> <li>• Enable (or disable) debugging of EIGRP events and notifications.</li> </ul>

4.9 The following commands are used for **showing EIGRP information**:

Command	Description
#show ip route eigrp	<ul style="list-style-type: none"> <li>• Show routing table entries learned through EIGRP.</li> </ul>
#show ip route <destination network address> [<subnet mask>]	<ul style="list-style-type: none"> <li>• Show detail information of the specified route including: <ul style="list-style-type: none"> <li>• Destination network address and netmask.</li> <li>• Routing protocol (i.e. EIGRP) and AS number from which this route is learned.</li> <li>• Administrative distance.</li> <li>• Route metric.</li> <li>• Redistribution protocol.</li> <li>• The last time the route was updated (in hh:mm:ss), the source of the update, and the</li> </ul> </li> </ul>

	<p>interface received the update.</p> <ul style="list-style-type: none"> <li>• Next hop IP address.</li> <li>• Delay (in microseconds), bandwidth (in kbps), reliability (1-255), load (1-255), MTU (in bytes) of the route.</li> <li>• etc.</li> </ul>
#show ip protocols	<ul style="list-style-type: none"> <li>• Show the parameters and current state of each active routing protocol process.</li> <li>• For EIGRP, the following information are shown: <ul style="list-style-type: none"> <li>• AS number.</li> <li>• Whether the incoming and outgoing filtering list have been set.</li> <li>• The protocol(s) that is being redistributed.</li> <li>• Whether automatic route summarization has been enabled.</li> <li>• The networks for which the routing process is currently injecting routes.</li> </ul> </li> <li>• Routing Information Sources. For each source, the following information are shown: <ul style="list-style-type: none"> <li>• IP address of the source/neighbor.</li> <li>• Administrative distance.</li> <li>• Time of the last update received from the source.</li> </ul> </li> <li>• etc.</li> </ul>
#show ip eigrp interfaces [<int-type> <int-num>] [<AS number>]	<ul style="list-style-type: none"> <li>• Show EIGRP-related interface information for all EIGRP interfaces, or the specified interface, or the interfaces configured with the specified AS.</li> <li>• The following information are shown for each interface: <ul style="list-style-type: none"> <li>• Interface name.</li> <li>• Number of directly connected neighbors (peers), etc.</li> </ul> </li> </ul>
#show ip eigrp neighbors [<AS number>] [detail]	<ul style="list-style-type: none"> <li>• Show EIGRP neighbor information.</li> <li>• The following information are shown for each neighbor: <ul style="list-style-type: none"> <li>• IP address of the neighbor.</li> <li>• Interface that is connected to the neighbor.</li> <li>• AS number.</li> <li>• Hold-time, i.e. length of time that the router will wait to hear from the neighbor before declaring it is down.</li> <li>• Uptime, i.e. elapsed time (in hours:minutes:seconds) since the router first heard from this neighbor.</li> </ul> </li> </ul>

	<ul style="list-style-type: none"> <li>• Number of EIGRP packets (update, query, and reply) that the router is waiting to send.</li> <li>• Sequence number of the last update, query, or reply packet that was received from this neighbor.</li> <li>• Smooth round-trip time (SRTT) (in ms) required for an EIGRP packet to be sent to the neighbor and an acknowledgment of that packet to be received.</li> <li>• Retransmission time-out (RTO) (in ms).</li> <li>• If the keyword "detail" is specified, the following additional information will also be shown for each neighbor: <ul style="list-style-type: none"> <li>• The IOS version of the neighbor.</li> <li>• The number of times that a packet has been retransmitted.</li> <li>• The number of times an attempt was made to retransmit a packet.</li> <li>• Elapsed time (in hours:minutes:seconds) since the neighbor has restarted.</li> </ul> </li> </ul>
<pre>#show ip eigrp topology [&lt;AS number&gt;   {&lt;ip-address&gt; [&lt;netmask&gt;]}] [summary]</pre>	<ul style="list-style-type: none"> <li>• Show all entries in the EIGRP topology database, or only the entries of the specified AS or destination IP address.</li> <li>• The following information are shown for each entry (destination network): <ul style="list-style-type: none"> <li>• Entry state (passive, active, update, query, or reply). In a stable network, all entries should be in the passive state. Active state indicates that the router has lost its path to the destination and is searching for a replacement. Update/query/reply state indicates that an update/query/reply packet was sent.</li> <li>• Destination network address and netmask.</li> <li>• Feasible distance, i.e. the best metric to reach the destination.</li> <li>• Number of successors, i.e. the number of next hops that are selected into the routing table.</li> <li>• For each known next hop to the destination, the following information are shown: <ul style="list-style-type: none"> <li>• IP address of the peer which advertises the path (i.e. the next hop router).</li> <li>• The interface from which the path is learned (i.e. the outgoing interface).</li> <li>• The EIGRP metric that represents the cost from the router to the destination.</li> <li>• The EIGRP metric that the peer advertises (i.e. the cost from the peer to</li> </ul> </li> </ul> </li> </ul>

	the destination). It is known as the advertised distance.
<code>#show ip eigrp traffic [&lt;AS number&gt;]</code>	<ul style="list-style-type: none"><li>• Show the number of EIGRP packets (hello, update, query, reply, or ack) sent and received for each AS, or the specified AS.</li></ul>

## Table of Contents of CCNP BSCI 642-801 Exam Notes

<b>Chapter 1.</b>	<b>Cisco Certifications</b>	<b>1</b>
1.	CISCO CERTIFICATION PATHS	1
2.	CISCO CERTIFICATION NETWORK ASSOCIATE (CCNA)	2
3.	CISCO CERTIFICATION NETWORK PROFESSIONAL (CCNP)	3
4.	CISCO CERTIFICATION INTERNETWORK EXPERT (CCIE) IN ROUTING AND SWITCHING	3
<b>Chapter 2.</b>	<b>Advanced IP Addressing</b>	<b>5</b>
1.	IP ADDRESSING FUNDAMENTALS	5
2.	SUBNETTING AND VARIABLE LENGTH SUBNET MASK (VLSM)	9
3.	CLASSLESS INTER-DOMAIN ROUTING (CIDR) AND ROUTE SUMMARIZATION	12
4.	NETWORK ADDRESS TRANSLATION (NAT)	15
5.	IP VERSION 6	25
6.	IP NETWORK DESIGN	32
<b>Chapter 3.</b>	<b>Routing Principles</b>	<b>35</b>
1.	ROUTING FUNDAMENTALS	35
2.	ROUTING PROTOCOLS	38
3.	CONFIGURING STATIC ROUTES AND ODR	47
4.	CONFIGURING RIP AND IGRP	51
5.	CONFIGURING ROUTE SUMMARIZATION	58
6.	TROUBLESHOOTING ROUTES	59
<b>Chapter 4.</b>	<b>Enhanced IGRP (EIGRP)</b>	<b>61</b>
1.	EIGRP BASICS	61
2.	EIGRP OPERATION	64
3.	SCALABILITY & PERFORMANCE	67
4.	CONFIGURING EIGRP	69
<b>Chapter 5.</b>	<b>Open Shortest Path First (OSPF)</b>	<b>78</b>
1.	OSPF BASICS	78
2.	OSPF OPERATION	83
3.	CONFIGURING OSPF - BASICS	88
4.	CONFIGURING OSPF OVER NBMA NETWORKS	94
5.	CONFIGURING OSPF ACROSS MULTIPLE AREAS	103
<b>Chapter 6.</b>	<b>Integrated IS-IS</b>	<b>115</b>
1.	IS-IS BASICS	115

2.	IS-IS OPERATION	122
3.	CONFIGURING IS-IS	131
<b>Chapter 7.</b>	<b>Manipulating Routing Updates</b>	<b>140</b>
1.	ROUTE REDISTRIBUTION	140
2.	ROUTE FILTERING WITH DISTRIBUTE LIST	145
3.	ROUTE-MAP	147
<b>Chapter 8.</b>	<b>Border Gateway Protocol (BGP)</b>	<b>155</b>
1.	BGP BASICS	155
2.	BGP OPERATION	161
3.	SCALABILITY & PERFORMANCE	168
4.	CONFIGURING BGP	171
<b>Appendix</b>		<b>183</b>
1.	ACCESS-LIST CONFIGURATION	183
2.	USEFUL WEBSITES	189
3.	COMMONLY USED WELL-KNOWN TCP AND UDP PORTS	191

## **CCNP BSCI 642-801 Exam Notes – All you need to pass the exam**

Copyright©2005 by the KP Lab Limited. All rights reserved. No part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of the publisher.

**ISBN** 988-97323-3-5  
**Publisher** KP Lab Limited  
**Author** K. Wan  
**Web Site** [www.kp-lab.com](http://www.kp-lab.com)  
**e-mail** [kplab@pacific.net.hk](mailto:kplab@pacific.net.hk)

### **About the Author**

K. Wan, MSc., CISSP, CCNP, CCSE, MCSE, MCDBA, SCSA, SCNA, SCJP, has more than ten years' experience in system and security administration on various computing platforms. He is currently an IT infrastructure and security manager working in Hong Kong.

**IT Certification Examination Study Guides published by KP Lab:**

1. CISSP Exam Notes  
ISBN: 988-97323-1-9  
Free Chapter:  
<http://www.kp-lab.com/download.htm>  
Complete Book (\$16.80):  
<http://www.kp-lab.com/cissp.htm>
  
2. CCNA 640-801 Exam Notes  
ISBN: 988-97323-2-7  
Free Chapter:  
<http://www.kp-lab.com/download.htm>  
Complete Book (\$10.39):  
<http://www.kp-lab.com/ccna.htm>
  
3. CCNP BSCI 642-801 Exam Notes  
ISBN: 988-97323-3-5  
Free Chapter:  
<http://www.kp-lab.com/download.htm>  
Complete Book (\$15.95):  
[http://www.kp-lab.com/ccnp\\_bsci.htm](http://www.kp-lab.com/ccnp_bsci.htm)