EIGRP Frequently Asked Questions

Document ID: 13681

Questions

Introduction

This document contains frequently asked questions (FAQs) about IP Enhanced Interior Gateway Routing Protocol (EIGRP).

Q. Does EIGRP require a default network to propagate a default route?

A. Although EIGRP can propagate a default route using the default network method, it is not required. EIGRP redistributes default routes directly.

Q. Should I always use the eigrp log−neighbor−changes command when I configure EIGRP?

A. Yes, this command makes it easy to determine why an EIGRP neighbor was reset. This reduces troubleshooting time.

Q. Does EIGRP support secondary addresses?

A. EIGRP does support secondary addresses. Since EIGRP always sources data packets from the primary address, Cisco recommends that you configure all routers on a particular subnet with primary addresses that belong to the same subnet. Routers do not form EIGRP neighbors over secondary networks. Therefore, if all of the primary IP addresses of routers do not agree, problems can arise with neighbor adjacencies.

Q. What debugging capabilities does EIGRP have?

A. There are protocol−independent and −dependent **debug** commands. There is also a suite of **show** commands that display neighbor table status, topology table status, and EIGRP traffic statistics. Some of these commands are:

- **show ip eigrp neighbors**
- **show ip eigrp interfaces**
- **show ip eigrp topology**
- **show ip eigrp traffic**

Q. What does the word serno mean on the end of an EIGRP topology entry when you issue the show ip eigrp topology command?

A. For example:

```
show ip eigrp topology
P 172.22.71.208/29, 2 successors, FD is 46163456
via 172.30.1.42 (46163456/45651456), Serial0.2, serno 7539273
via 172.30.2.49 (46163456/45651456), Serial2.6, serno 7539266
```

Serno stands for serial number. When DRDBs are threaded to be sent, they are assigned a serial number from a circular number space shared by all interfaces. If you display the topology table at the time an entry is threaded, it shows you the serial number associated with the DRDB.

Threading is the technique used inside the router to queue items up for transmission to neighbors. The updates are not created until it is time for them to go out the interface. Prior to that, a linked list of pointers to items to send is created (for example, the thread).

These sernos are local to the router and are not passed with the routing update.

Q. What percent of bandwidth and processor resources does EIGRP use?

A. EIGRP version 1 introduced a feature that prevents any single EIGRP process from using more than fifty percent of the configured bandwidth on any link during periods of network convergence. Each AS or protocol (for instance, IP, IPX, or Appletalk) serviced by EIGRP is a separate process. You can use the **ip bandwidth−percent eigrp interface** configuration command in order to properly configure the bandwidth statement on each WAN interface. Refer to the EIGRP White Paper for more details on how this feature works.

In addition, the implementation of partial and incremental updates means that EIGRP sends routing information only when a topology change occurs. This feature significantly reduces bandwidth use.

The feasible successor feature of EIGRP reduces the amount of processor resources used by an autonomous system (AS). It requires only the routers affected by a topology change to perform route re−computation. The route re−computation only occurs for routes that were affected, which reduces search time in complex data structures.

Q. Does EIGRP support aggregation and variable length subnet masks?

A. Yes, EIGRP supports aggregation and variable length subnet masks (VLSM). Unlike Open Shortest Path First (OSPF), EIGRP allows summarization and aggregation at any point in the network. EIGRP supports aggregation to any bit. This allows properly designed EIGRP networks to scale exceptionally well without the use of areas. EIGRP also supports automatic summarization of network addresses at major network borders.

Q. Does EIGRP support areas?

A. No, a single EIGRP process is analogous to an area of a link−state protocol. However, within the process, information can be filtered and aggregated at any interface boundary. In order to bound the propagation of routing information, you can use summarization to create a hierarchy.

Q. Can I configure more than one EIGRP autonomous system on the same router?

A. You can configure more than one EIGRP autonomous system on the same router, but Cisco does not recommend it. Multiple EIGRP autonomous systems on the same router that use mutual redistribution can cause discrepancies in the EIGRP topology table. Cisco recommends you configure only one EIGRP autonomous system in any single autonomous system. Also use another protocol, such as BGP, in order to connect the two EIGRP autonomous systems.

Q. If there are two EIGRP processes that run and two equal paths are learned, one by each EIGRP process, do both routes get installed?

A. No, only one route is installed. The router installs the route that was learned through the EIGRP process with the lower Autonomous System (AS) number. In Cisco IOS releases earlier than 12.2(7)T, the router installed the path with the latest timestamp received from either of the EIGRP processes. The change in behavior is tracked by Cisco bug ID CSCdm47037.

Q. What does the EIGRP stuck in active message mean?

A. When EIGRP returns a stuck in active (SIA) message, it means that it has not received a reply to a query. EIGRP sends a query when a route is lost and another feasible route does not exist in the topology table. The SIA is caused by two sequential events:

- The route reported by the SIA has gone away.
- An EIGRP neighbor (or neighbors) have not replied to the query for that route.

When the SIA occurs, the router clears the neighbor that did not reply to the query. When this happens, determine which neighbor has been cleared. Keep in mind that this router can be many hops away. Refer to What Does the EIGRP DUAL−3−SIA Error Message Mean? for more information.

Q. What does the neighbor statement in the EIGRP configuration section do?

A. The **neighbor** command is used in EIGRP in order to define a neighboring router with which to exchange routing information. Due to the current behavior of this command, EIGRP exchanges routing information with the neighbors in the form of unicast packets whenever the **neighbor** command is configured for an interface. EIGRP stops processing all multicast packets that come inbound on that interface. Also, EIGRP stops sending multicast packets on that interface.

The ideal behavior of this command is for EIGRP to start sending EIGRP packets as unicast packets to the specified neighbor, but not stop sending and receiving multicast packets on that interface. Since the command does not behave as intended, the **neighbor** command should be used carefully, understanding the impact of the command on the network.

Q. Why does the EIGRP passive−interface command remove all neighbors for an interface?

A. The **passive−interface** command disables the transmission and receipt of EIGRP hello packets on an interface. Unlike IGRP or RIP, EIGRP sends hello packets in order to form and sustain neighbor adjacencies. Without a neighbor adjacency, EIGRP cannot exchange routes with a neighbor. Therefore, the **passive−interface** command prevents the exchange of routes on the interface. Although EIGRP does not send or receive routing updates on an interface configured with the **passive−interface** command, it still includes the address of the interface in routing updates sent out of other non−passive interfaces. Refer to How Does the Passive Interface Feature Work in EIGRP? for more information.

Q. Why are routes received from one neighbor on a point−to−multipoint interface that runs EIGRP not propagated to another neighbor on the same point−to−multipoint interface?

A. The split horizon rule prohibits a router from advertising a route through an interface that the router itself uses to reach the destination. In order to disable the split horizon behavior, use the **no ip split−horizon eigrp *as−number*** interface command. Some important points to remember about EIGRP split horizon are:

- Split horizon behavior is turned on by default.
- When you change the EIGRP split horizon setting on an interface, it resets all adjacencies with EIGRP neighbors reachable over that interface.

- Split horizon should only be disabled on a hub site in a hub−and−spoke network.
- Disabling split horizon on the spokes radically increases EIGRP memory consumption on the hub router, as well as the amount of traffic generated on the spoke routers.
- The EIGRP split horizon behavior is not controlled or influenced by the **ip split−horizon** command.

For more details on split horizon and poison reverse, refer to Split Horizon and Poison Reverse. For more information on commands, refer to EIGRP Commands.

Q. When I configure EIGRP, how can I configure a network statement with a mask?

A. The optional network−mask argument was first added to the network statement in Cisco IOS Software Release 12.0(4)T. The mask argument can be configured in any format (such as in a network mask or in wild card bits). For example, you can use **network 10.10.10.0 255.255.255.252** or **network 10.10.10.0 0.0.0.3**.

Q. I have two routes: 172.16.1.0/24 and 172.16.1.0/28. How can I deny 172.16.1.0/28 while I allow 172.16.1.0/24 in EIGRP?

A. In order to do this you need to use a prefix−list as shown here:

```
router eigrp 100
       network 172.16.0.0
       distribute-list prefix test in
       auto-summary
       no eigrp log-neighbor-changes
       !
       ip prefix-list test seq 5 permit 172.16.1.0/24
```

This allows only the 172.16.1.0/24 prefix and therefore denies 172.16.1.0/28.

Note: The use of ACL and distribute−list under EIGRP does not work in this case. This is because ACLs do not check the mask, they just check the network portion. Since the network portion is the same, when you allow 172.16.1.0/24, you also allow 172.16.1.0/28.

Q. I have a router that runs Cisco Express Forwarding (CEF) and EIGRP. Who does load−balancing when there are multiple links to a destination?

A. The way in which CEF works is that CEF does the switching of the packet based on the routing table which is populated by the routing protocols such as EIGRP. In short, CEF does the load−balancing once the routing protocol table is calculated. Refer to How does load−balancing work? for more details on load balancing.

Q. How can I use only one path when a router has two equal cost paths?

A. Configure the bandwidth value on the interfaces to default, and increase the delay on the backup interface so that the router does not see two equal cost paths.

Q. What is the difference in metric calculation between EIGRP and IGRP?

A. The EIGRP metric is obtained when you multiply the IGRP metric by 256. The IGRP uses only 24 bits in its update packet for the metric field, but EIGRP uses 32 bits in its update

packet for the metric field. For example, the IGRP metric to a destination network is 8586, but the EIGRP metric is 8586 x 256 = 2,198,016. Integer division is used when you divide 10^7 by minimum BW, so the calculation involves integer division, which leads to a variation from manual calculation.

Q. What is the EIGRP Stub Routing feature?

A. The Stub routing feature is used to conserve bandwidth by summarizing and filtering routes. Only specified routes are propagated from the remote (Stub) router to the distribution router because of the Stub routing feature. For more information about the stub routing feature, refer to EIGRP Stub Routing. The EIGRP stub feature can be configured on the switch with the **eigrp stub** command, and it can be removed with the **no eigrp stub**. When you remove the **eigrp stub** command from the switch, the switch that runs the IP Base image throws the error:

```
EIGRP is restricted to stub configurations only
```

This issue can be resolved if you upgrade to Advanced Enterprise Images. This error is documented in CSCeh58135.

Q. How can I send a default route to the Stub router from the hub?

A. Do this under the outbound interface on the hub router with the command **ip summary−address eigrp X 0.0.0.0 0.0.0.0**. This command suppresses all the more specific routes and only sends the summary route. In the case of the 0.0.0.0 0.0.0.0, it means it suppresses everything, and the only route that is in the outbound update is 0.0.0.0/0. One drawback to this method is that EIGRP installs a 0.0.0.0/0 route to Null0 is the local routing table with an admin distance of 5.

Q. How EIGRP behaves over a GRE tunnel compared to a directly connected network?

A. EIGRP will use the same administrative distance and metric calculation for the GRE tunnel. The cost calculation is based on bandwidth and delay. The bandwidth and delay of the GRE tunnel will be taken from the tunnel interface configured on the router. The tunnel will also be treated like a directly connected network. If there are two paths to reach a network either through a VLAN interface or tunnel interface, EIGRP prefers the Virtual−Access Interface (VAI) VLAN interface because the VLAN interface has greater bandwidth than the tunnel interface. In order to influence the routing through the tunnel interface, increase the bandwidth parameter of the tunnel interface, or increase the delay parameter of the VLAN interface.

Q. What is an offset−list, and how is it useful?

A. The offset−list is an feature used to modify the composite metrics in EIGRP, for example, bandwidth and delay. Since the bandwidth taken for the calculation is local to the router, the value applied in the offset−list is not applicable to bandwidth, so the offset−list affects only the delay value. The value configured in the offset−list command is added to the delay value calculated by the router for the route matched by an access−list. An offset−list is the preferred method to influence a particular path that is advertised and/or chosen.

Q. How can I tag external routes in EIGRP?

A. You can tag only routes that EIGRP has learned from another routing protocol.

Q. What are the primary functions of the PDM?

A. EIGRP supports 3 protocol suites: IP, IPv6, and AppleTalk. Each of them has its own PDM. These are the primary functions of PDM:

- Maintaining the neighbor and yopology tables of EIGRP routers that belong to that protocol suite
- Building and translating protocol specific packets for DUAL
- Interfacing DUAL to the protocol specific routing table
- Computing the metric and passing this information to DUAL; DUAL handles only the picking of the feasible successors (FSs)
- Implement filtering and access−lLists.
- Perform redistribution functions to/from other routing protocols.

Q. What are the various load−balancing options available in EIGRP?

A. The offset list can be used to modify the metrics of routes that EIGRP learns through a particular interface, or PBR can be used.

Related Information

- **EIGRP Support Page**
- **Technical Support & Documentation − Cisco Systems**

Updated: Jul 21, 2008 Document ID: 13681