

Introduction to EIGRP

Document ID: 13669

Contents

Introduction

Before You Begin

- Conventions

- Prerequisites

- Components Used

What is IGRP?

What is EIGRP?

How Does EIGRP Work?

EIGRP Concepts

- Neighbor Table

- Topology Table

- Feasible Successors

- Route States

- Packet Formats

- Route Tagging

Compatibility Mode

DUAL Example

Frequently Asked Questions

- Is configuring EIGRP as easy as configuring IGRP?

- Do I have debugging capabilities like IGRP?

- Are the same features available in IP-EIGRP as are available in IP-IGRP?

- How much bandwidth and processor resources does EIGRP use?

- Does IP-EIGRP support aggregation and variable length subnet masks?

- Does EIGRP support areas?

Related Information

Introduction

This paper is an introduction to the Interior Gateway Routing Protocol (IGRP) suite of routing protocols designed and developed by Cisco Systems. This paper should be used as an information-only document intended as a technology introduction and does not represent a protocol specification or product description.

Before You Begin

Conventions

Refer to Cisco Technical Tips Conventions for more information on document conventions.

Prerequisites

There are no specific prerequisites for this document.

Components Used

This document is not restricted to specific software and hardware versions.

What is IGRP?

IGRP is used in TCP/IP and Open System Interconnection (OSI) Internets. The original IP version was designed and deployed successfully in 1986. It is regarded as an IGP but has also been used extensively as an Exterior Gateway Protocol (EGP) for inter-domain routing. IGRP uses distance vector routing technology. The concept is that each router need not know all the router/link relationships for the entire network. Each router advertises destinations with a corresponding distance. Each router hearing the information adjusts the distance and propagates it to neighboring routers.

The distance information in IGRP is represented as a composite of available bandwidth, delay, load utilization, and link reliability. This allows fine tuning of link characteristics to achieve optimal paths.

What is EIGRP?

EIGRP is an enhanced version of IGRP. The same distance vector technology found in IGRP is also used in EIGRP, and the underlying distance information remains unchanged. The convergence properties and the operating efficiency of this protocol have improved significantly. This allows for an improved architecture while retaining existing investment in IGRP.

The convergence technology is based on research conducted at SRI International. The Diffusing Update Algorithm (DUAL) is the algorithm used to obtain loop-freedom at every instant throughout a route computation. This allows all routers involved in a topology change to synchronize at the same time. Routers that are not affected by topology changes are not involved in the recomputation. The convergence time with DUAL rivals that of any other existing routing protocol.

EIGRP has been extended to be network-layer-protocol independent, thereby allowing DUAL to support other protocol suites.

How Does EIGRP Work?

EIGRP has four basic components:

- Neighbor Discovery/Recovery
- Reliable Transport Protocol
- DUAL Finite State Machine
- Protocol Dependent Modules

Neighbor Discovery/Recovery is the process that routers use to dynamically learn of other routers on their directly attached networks. Routers must also discover when their neighbors become unreachable or inoperative. This process is achieved with low overhead by periodically sending small hello packets. As long as hello packets are received, a router can determine that a neighbor is alive and functioning. Once this is determined, the neighboring routers can exchange routing information.

The reliable transport is responsible for guaranteed, ordered delivery of EIGRP packets to all neighbors. It supports intermixed transmission of multicast or unicast packets. Some EIGRP packets must be transmitted reliably and others need not. For efficiency, reliability is provided only when necessary. For example, on a multi-access network that has multicast capabilities, such as Ethernet, it is not necessary to send hellos reliably to all neighbors individually. So EIGRP, sends a single multicast hello with an indication in the packet informing the receivers that the packet need not be acknowledged. Other types of packets, such as updates, require acknowledgment and this is indicated in the packet. The reliable transport has a provision to send multicast packets quickly when there are unacknowledged packets pending. This helps insure that convergence time remains low in the presence of varying speed links.

The DUAL finite state machine embodies the decision process for all route computations. It tracks all routes advertised by all neighbors. The distance information, known as a metric, is used by DUAL to select efficient loop free paths. DUAL selects routes to be inserted into a routing table based on feasible successors. A successor is a neighboring router used for packet forwarding that has a least cost path to a destination that is guaranteed not to be part of a routing loop. When there are no feasible successors but there are neighbors advertising the destination, a recomputation must occur. This is the process where a new successor is determined. The amount of time it takes to recompute the route affects the convergence time. Even though the recomputation is not processor-intensive, it is advantageous to avoid recomputation if it is not necessary. When a topology change occurs, DUAL will test for feasible successors. If there are feasible successors, it will use any it finds in order to avoid any unnecessary recomputation. Feasible successors are defined in more detail later in this document.

The protocol-dependent modules are responsible for network layer, protocol-specific requirements. For example, the IP-EIGRP module is responsible for sending and receiving EIGRP packets that are encapsulated in IP. IP-EIGRP is responsible for parsing EIGRP packets and informing DUAL of the new information received. IP-EIGRP asks DUAL to make routing decisions and the results of which are stored in the IP routing table. IP-EIGRP is responsible for redistributing routes learned by other IP routing protocols.

EIGRP Concepts

This section describes some details about Cisco's EIGRP implementation. Both data structures and the DUAL concepts are discussed.

Neighbor Table

Each router keeps state information about adjacent neighbors. When newly discovered neighbors are learned, the address and interface of the neighbor is recorded. This information is stored in the neighbor data structure. The neighbor table holds these entries. There is one neighbor table for each protocol dependent module. When a neighbor sends a hello, it advertises a HoldTime. The HoldTime is the amount of time a router treats a neighbor as reachable and operational. In other words, if a hello packet isn't heard within the HoldTime, then the HoldTime expires. When the HoldTime expires, DUAL is informed of the topology change.

The neighbor table entry also includes information required by the reliable transport mechanism. Sequence numbers are employed to match acknowledgments with data packets. The last sequence number received from the neighbor is recorded so out of order packets can be detected. A transmission list is used to queue packets for possible retransmission on a per neighbor basis. Round trip timers are kept in the neighbor data structure to estimate an optimal retransmission interval.

Topology Table

The Topology Table is populated by the protocol dependent modules and acted upon by the DUAL finite state machine. It contains all destinations advertised by neighboring routers. Associated with each entry is the destination address and a list of neighbors that have advertised the destination. For each neighbor, the advertised metric is recorded. This is the metric that the neighbor stores in its routing table. If the neighbor is advertising this destination, it must be using the route to forward packets. This is an important rule that distance vector protocols must follow.

Also associated with the destination is the metric that the router uses to reach the destination. This is the sum of the best advertised metric from all neighbors plus the link cost to the best neighbor. This is the metric that the router uses in the routing table and to advertise to other routers.

Feasible Successors

A destination entry is moved from the topology table to the routing table when there is a feasible successor. All minimum cost paths to the destination form a set. From this set, the neighbors that have an advertised metric less than the current routing table metric are considered feasible successors.

Feasible successors are viewed by a router as neighbors that are downstream with respect to the destination. These neighbors and the associated metrics are placed in the forwarding table.

When a neighbor changes the metric it has been advertising or a topology change occurs in the network, the set of feasible successors may have to be re-evaluated. However, this is not categorized as a route recomputation.

Route States

A topology table entry for a destination can have one of two states. A route is considered in the Passive state when a router is not performing a route recomputation. The route is in Active state when a router is undergoing a route recomputation. If there are always feasible successors, a route never has to go into Active state and avoids a route recomputation.

When there are no feasible successors, a route goes into Active state and a route recomputation occurs. A route recomputation commences with a router sending a query packet to all neighbors. Neighboring routers can either reply if they have feasible successors for the destination or optionally return a query indicating that they are performing a route recomputation. While in Active state, a router cannot change the next-hop neighbor it is using to forward packets. Once all replies are received for a given query, the destination can transition to Passive state and a new successor can be selected.

When a link to a neighbor that is the only feasible successor goes down, all routes through that neighbor commence a route recomputation and enter the Active state.

Packet Formats

EIGRP uses five packet types:

- Hello/Acks
- Updates
- Queries
- Replies
- Requests

As stated earlier, hellos are multicast for neighbor discovery/recovery. They do not require acknowledgment. A hello with no data is also used as an acknowledgment (ack). Acks are always sent using a unicast address and contain a non-zero acknowledgment number.

Updates are used to convey reachability of destinations. When a new neighbor is discovered, update packets are sent so the neighbor can build up its topology table. In this case, update packets are unicast. In other cases, such as a link cost change, updates are multicast. Updates are always transmitted reliably.

Queries and replies are sent when destinations go into Active state. Queries are always multicast unless they are sent in response to a received query. In this case, it is unicast back to the successor that originated the query. Replies are always sent in response to queries to indicate to the originator that it does not need to go into Active state because it has feasible successors. Replies are unicast to the originator of the query. Both queries and replies are transmitted reliably.

Request packets are used to get specific information from one or more neighbors. Request packets are used in route server applications. They can be multicast or unicast. Requests are transmitted unreliably.

Route Tagging

EIGRP has the notion of internal and external routes. Internal routes are ones that have been originated within an EIGRP autonomous system (AS). Therefore, a directly attached network that is configured to run EIGRP is considered an internal route and is propagated with this information throughout the EIGRP AS. External routes are ones that have been learned by another routing protocol or reside in the routing table as static routes. These routes are tagged individually with the identity of their origination.

External routes are tagged with the following information:

- The router ID of the EIGRP router that redistributed the route.
- The AS number where the destination resides.
- A configurable administrator tag.
- Protocol ID of the external protocol.
- The metric from the external protocol.
- Bit flags for default routing.

As an example, suppose there is an AS with three border routers. A border router is one that runs more than one routing protocol. The AS uses EIGRP as the routing protocol. Let's say two of the border routers, BR1 and BR2, use Open Shortest Path First (OSPF) and the other, BR3, uses Routing Information Protocol (RIP).

Routes learned by one of the OSPF border routers, BR1, can be conditionally redistributed into EIGRP. This means that EIGRP running in BR1 advertises the OSPF routes within its own AS. When it does so, it advertises the route and tags it as an OSPF learned route with a metric equal to the routing table metric of the OSPF route. The router-id is set to BR1. The EIGRP route propagates to the other border routers. Let's say that BR3, the RIP border router, also advertises the same destinations as BR1. Therefore BR3, redistributes the RIP routes into the EIGRP AS. BR2, then, has enough information to determine the AS entry point for the route, the original routing protocol used, and the metric. Further, the network administrator could assign tag values to specific destinations when redistributing the route. BR2 can use any of this information to use the route or re-advertise it back out into OSPF.

Using EIGRP route tagging can give a network administrator flexible policy controls and help customize routing. Route tagging is particularly useful in transit ASes where EIGRP would typically interact with an inter-domain routing protocol that implements more global policies. This combines for very scalable policy based routing.

Compatibility Mode


EIGRP provides compatibility and seamless inter-operation with IGRP routers. This is important so users can take advantage of the benefits of both protocols. The compatibility features do not require users to have a flag day to enable EIGRP. EIGRP can be enabled in strategic places carefully without disruption to IGRP performance.

There is an automatic redistribution mechanism used so IGRP routes are imported into EIGRP and vice versa. Since the metrics for both protocols are directly translatable, they are easily comparable as if they were routes that originated in their own AS. In addition, IGRP routes are treated as external routes in EIGRP so the tagging capabilities are available for custom tuning.

IGRP routes take precedence over EIGRP routes by default. This can be changed with a configuration command that does not require the routing processes to restart.

DUAL Example

The following network diagram illustrates how DUAL converges. The example focuses on destination N only. Each node shows its cost to N (in hops). The arrows show the node's successor. So, for example, C uses A to reach N and the cost is 2.



If the link between A and B fails, B sends a query informing its neighbors that it has lost its feasible successor. D receives the query and determines if it has any other feasible successors. If it does not, it has to start a route computation and enter the active state. However in this case, C is a feasible successor because its cost (2) is less than than D's current cost (3) to destination N. D can switch to C as its successor. Note A and C did not participate because they were unaffected by the change.

Now let's cause a route computation to occur. In this scenario, let's say the link between A and C fails. C determines that it has lost its successor and has no other feasible successors. D is not considered a feasible successor because its advertised metric (3) is greater than C's current cost (2) to reach destination N. C must perform a route computation for destination N. C sends a query to its only neighbor D. D replies because its successor has not changed. D does not need to perform a route computation. When C receives the reply it knows that all neighbors have processed the news about the failure to N. At this point, C can choose its new feasible successor D with a cost of (4) to reach destination N. Note that A and B were unaffected by the topology change and D needed to simply reply to C.

Frequently Asked Questions

Is configuring EIGRP as easy as configuring IGRP?

Yes, you configure EIGRP just like you configure IGRP. You configure a routing process and which networks the protocol should run over. Existing configuration files can be used.

Do I have debugging capabilities like IGRP?

Yes, there are both protocol independent and dependent **debug** commands that inform you what the protocol is doing. There is a suite of **show** commands that give you neighbor table status, topology table status, and EIGRP traffic statistics.

Are the same features available in IP–EIGRP as are available in IP–IGRP?

All of the features that you have used in IGRP are available in EIGRP. One feature to point out is multiple routing processes. You can use a single process that runs both IGRP and EIGRP. You can use multiple processes that run both. You can use one process that runs IGRP and another to run EIGRP. You can mix and match. This can help customize your routing to a particular protocol as your needs change.

How much bandwidth and processor resources does EIGRP use?

The bandwidth utilization issue has been addressed by implementing partial and incremental updates. Therefore, only when a topology change occurs does routing information get sent. Regarding processor utilization, the feasible successor technology greatly reduces the total processor utilization of an AS by requiring only the routers that were affected by a topology change to perform the route recomputation. Furthermore, the route recomputation only occurs for routes that were affected. Only those data structures are accessed and used. This greatly reduces search time in complex data structures.

Does IP–EIGRP support aggregation and variable length subnet masks?

Yes it does. IP–EIGRP performs route aggregation the same way IGRP does. That is, subnets of an IP network are not advertised over another IP network. The subnet routes are summarized into a single network number aggregate. In addition, IP–EIGRP will allow aggregation on any bit boundary in an IP address and can be configured at network interface granularity.

Does EIGRP support areas?

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. If one wants to bound the propagation of routing information, multiple routing processes can be configured to achieve a hierarchy. Since DUAL itself limits route propagation, multiple routing processes are typically used to define organizational boundaries.

Related Information

- [EIGRP Support Page](#)
- [Technical Support – Cisco Systems](#)

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2009 – 2010 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Aug 10, 2005

Document ID: 13669
