



TECHNICAL IMPLEMENTATION GUIDE

CISCO IOS EIGRP FOR MANAGED SERVICES BASIC OPERATIONS AND PROTOCOL CHARACTERISTICS

This document covers the following information about the Enhanced Interior Gateway Routing Protocol (EIGRP):

- Cisco IOS® Software-based technologies for managed services
- Cisco IOS EIGRP in managed services
- Four primary technologies employed by EIGRP
- Diffusing Update Algorithm (DUAL), and how it improves the operational efficiency of EIGRP
- How to use EIGRP to interconnect networks with different routing protocols as well as different routed protocols
- How to migrate gradually to EIGRP

The managed network services opportunity is projected to increase significantly in the coming years. A recent Cisco Systems® survey of 500 large corporations found substantial interest in managed IP services. At the top of the list were IP VPNs as a foundation for the integration of older networks and the addition of new services. Business customers—from the largest global corporations to midsize and smaller firms—are focusing on achieving cost efficiencies while adding new services. Recognizing the value of the network as a strategic tool, many companies are turning to service providers to manage their networks so they can focus more resources on their businesses. The IP VPN allows integration of older networks (such as ATM and Frame Relay) and provides a foundation for many new services. Examples include managed core services offerings and managed WAN and LAN services, which have been augmented by improved Web-based, user-friendly tools; service-level agreements (SLAs) and guarantees; and many newer IP-based applications, such as voice over IP (VoIP).

Cisco IOS Software technologies make possible a secure, highly available, cost-effective managed services environment within an IP VPN. Partnering with Cisco®, service providers can streamline their infrastructures to avoid unnecessary overhead, offer more services more efficiently, and position these service offerings successfully with the established Cisco global enterprise customer base. Features such as EIGRP make routing more efficient and turn IP/Multiprotocol Label Switching (MPLS) VPNs into a simpler, benefit-rich addition to customer networks.

Cisco IOS Software technologies for managed services environments serve as the foundation for high-speed routing and IP/MPLS, scalable IP VPNs, and robust network security, all integrated through a next-generation network management interface. These operate within several network topologies to fit the needs of different customers. Products in the Cisco IOS Software Family bring customizable networking solutions to headquarters, branch offices, and campuses, and extend full network capability to mobile workers, telecommuters, and remote data centers.

The EIGRP IP network routing solution is well-known to Cisco enterprise customers, who are taking advantage of its greater optimization of path routing, fast convergence, and lower CPU usage benefits. EIGRP allows service providers to more rapidly and cost-effectively deploy IP VPN services for their customers. In turn, enterprise customers can more quickly enjoy affordable, efficient, and secure managed VPNs.

Layer 3 IP VPNs built upon Cisco IP/MPLS technology are among the managed VPN services that many enterprise customers are considering. IP/MPLS VPNs deployed using site-to-site EIGRP—rather than external Border Gateway

Protocol (BGP) or static routes—eliminate the need for the enterprise network staff to learn new protocols. By integrating the capabilities of link-state and distance-vector protocols, EIGRP greatly increases operational efficiency in networks of all sizes. The protocol is highly scalable for large networks, giving bigger enterprises the confidence they need to implement managed services offered by service providers.

OVERVIEW OF EIGRP

The Enhanced Interior Gateway Routing Protocol (EIGRP) represents an evolution from its predecessor IGRP. This evolution resulted from changes in networking and the demands of diverse, large-scale internetworks. EIGRP integrates the capabilities of link-state protocols into distance-vector protocols. Additionally, EIGRP contains several important protocols that greatly increase its operational efficiency relative to other routing protocols. One of these protocols is the *Diffusing Update Algorithm* (DUAL) developed at SRI International by J.J. Garcia-Luna-Aceves. DUAL enables EIGRP routers to determine whether a path advertised by a neighbor is looped or loop-free, and allows a router running EIGRP to find alternate paths without waiting on updates from other routers.

EIGRP provides compatibility and transparent interoperation with IGRP routers. An automatic-redistribution mechanism allows IGRP routes to be imported into EIGRP, and conversely, so it is possible to add EIGRP gradually into an existing IGRP network. Because the metrics for both protocols are directly translatable, they are as easily comparable as if they were routes that originated in their own autonomous systems. In addition, EIGRP treats IGRP routes as external routes and provides a way for the network administrator to customize them.

This document provides an overview of the basic operations and protocol characteristics of EIGRP.

EIGRP CAPABILITIES AND ATTRIBUTES

Primary capabilities that distinguish EIGRP from other routing protocols include fast convergence, support for variable-length subnet mask, support for partial updates, and support for multiple network layer protocols.

A router running EIGRP stores all its neighbors' routing tables so that it can quickly adapt to alternate routes. If no appropriate route exists, EIGRP queries its neighbors to discover an alternate route. These queries propagate until an alternate route is found.

EIGRP support for variable-length subnet masks permits routes to be automatically summarized on a network number boundary. In addition, EIGRP can be configured to summarize on any bit boundary at any interface.

EIGRP does not make periodic updates. Instead, it sends partial updates only when the metric for a route changes. Propagation of partial updates is automatically bounded so that only those routers that need the information are updated. As a result of these two capabilities, EIGRP consumes significantly less bandwidth than IGRP.

EIGRP includes support for AppleTalk, IP, and Novell NetWare. The AppleTalk implementation redistributes routes learned from the Routing Table Maintenance Protocol (RTMP). The IP implementation redistributes routes learned from Open Shortest Path First (OSPF), Routing Information Protocol (RIP), Intermediate System-to-Intermediate System (IS-IS), Exterior Gateway Protocol (EGP), or Border Gateway Protocol (BGP). The Novell implementation redistributes routes learned from Novell RIP or Service Advertising Protocol (SAP).

UNDERLYING PROCESSES AND TECHNOLOGIES

To provide superior routing performance, EIGRP employs four important technologies that combine to differentiate it from other routing technologies: neighbor discovery and recovery, Reliable Transport Protocol (RTP), DUAL finite-state machine, and protocol-dependent modules.

The *neighbor discovery and recovery* mechanism helps enable routers to dynamically learn about other routers on their directly attached networks. Routers also must discover when their neighbors become unreachable or inoperative. This process is achieved with low overhead by periodically sending small hello packets. As long as a router receives hello packets from a neighboring router, it assumes that the neighbor is functioning, and the two can exchange routing information.

RTP is responsible for guaranteed, ordered delivery of EIGRP packets to all neighbors. It supports intermixed transmission of multicast or unicast packets. For efficiency, only certain EIGRP packets are transmitted reliably. On a multiaccess network that has multicast capabilities, such as Ethernet, it is not necessary to send hello packets reliably to all neighbors individually. For that reason, EIGRP sends a single multicast hello packet containing an indicator that informs the receivers that the packet does not need to be acknowledged. Other types of packets, such as updates, indicate in the packet that acknowledgment is required. RTP contains a provision for sending multicast packets quickly when unacknowledged packets are pending, helping ensure 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 by tracking all routes advertised by all neighbors. DUAL uses distance information to select efficient, loop-free paths and selects routes for insertion in a routing table based on feasible successors. A feasible successor is a neighboring router used for packet forwarding that is a least-cost path to a destination that is guaranteed not to be part of a routing loop. When a neighbor changes a metric, or when a topology change occurs, DUAL tests for feasible successors. If one is found, DUAL uses it to avoid recomputing the route unnecessarily. When no feasible successors exist but neighbors still advertise the destination, a recomputation (also known as a diffusing computation) must occur to determine a new successor. Although recomputation is not processor-intensive, it does affect convergence time, so it is advantageous to avoid unnecessary recomputations.

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

ROUTING CONCEPTS

EIGRP relies on four fundamental concepts: neighbor tables, topology tables, route states, and route tagging. Each of these is summarized in the discussions that follow.

Neighbor Tables

When a router discovers a new neighbor, it records the neighbor's address and interface as an entry in the *neighbor table*. One neighbor table exists for each protocol-dependent module. When a neighbor sends a hello packet, it advertises a hold time, which is the amount of time that a router treats a neighbor as reachable and operational. If a hello packet is not received within the hold time, the hold time expires and DUAL is informed of the topology change.

The neighbor-table entry also includes information required by RTP. Sequence numbers are employed to match acknowledgments with data packets, and the last sequence number received from the neighbor is recorded so that 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-table entry to estimate an optimal retransmission interval.

Topology Tables

The *topology table* contains all destinations advertised by neighboring routers. The protocol-dependent modules populate the table, and the table is acted on by the DUAL finite-state machine. Each entry in the topology table includes the destination address and a list of neighbors that have advertised the destination. For each neighbor, the entry records the advertised metric, which the neighbor stores in its routing table. An important rule that distance-vector protocols must follow is that if the neighbor advertises this destination, it must use the route to forward packets.

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

Route States

A topology-table entry for a destination can exist in one of two states: active or passive. A destination is in the *passive state* when the router is not performing a recomputation; it is in the *active state* when the router is performing a recomputation. If feasible successors are always available, a destination never has to go into the active state, thereby avoiding a recomputation.

A recomputation occurs when a destination has no feasible successors. The router initiates the recomputation by sending a query packet to each of its neighboring routers. The neighboring router can send a reply packet, indicating that it has a feasible successor for the destination, or it can send a query packet, indicating that it is participating in the recomputation. While a destination is in the active state, a router cannot change the destination routing-table information. After the router has received a reply from each neighboring router, the topology-table entry for the destination returns to the passive state, and the router can select a successor.

Route Tagging

EIGRP supports internal and external routes. Internal routes originate within an EIGRP autonomous system. 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 autonomous system. External routes are either learned by another routing protocol or they reside in the routing table as static routes. These routes are tagged individually with the identity of their origin.

External routes are tagged with the following information:

- Router ID of the EIGRP router that redistributed the route
- Autonomous-system number of the destination
- Configurable administrator tag
- ID of the external protocol
- Metric from the external protocol
- Bit flags for default routing

Route tagging allows the network administrator to customize routing and maintain flexible policy controls. Route tagging is particularly useful in transit autonomous systems, where EIGRP typically interacts with an interdomain routing protocol that implements more global policies, resulting in very scalable, policy-based routing.

EIGRP PACKET TYPES

EIGRP uses the following packet types: hello and acknowledgment, update, and query and reply.

Hello packets are multicast for neighbor discovery and recovery and do not require acknowledgment. An *acknowledgment* packet is a hello packet that has no data. Acknowledgment packets contain a nonzero acknowledgment number and always are sent by using a unicast address.

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

Query and reply packets are sent when a destination has no feasible successors. *Query* packets are always multicast. *Reply* packets are sent in response to query packets to instruct the originator not to recompute the route because feasible successors exist. Reply packets are unicast to the originator of the query. Both query and reply packets are transmitted reliably.

SUMMARY

Cisco's EIGRP is one of the most feature-rich and robust routing protocols ever developed. Its unique combination of features blends the best attributes of distance-vector protocols with the best attributes of link-state protocols. The result is a hybrid routing protocol that defies easy categorization with conventional protocols.

EIGRP also is remarkably easy to configure and use, as well as remarkably efficient and secure in operation. It can be used in conjunction with IPv4, AppleTalk, and Internetwork Packet Exchange (IPX). More importantly, its modular architecture readily enables Cisco to add support for other routed protocols that may be developed in the future.

REVIEW QUESTIONS

Q—Name the four primary technologies that are used by EIGRP.

A—EIGRP employs four primary technologies, including neighbor discover and recovery, Reliable Transport Protocol (RTP), Diffusing Update Algorithm (DUAL) finite-state machine, and a modular architecture that helps enable support for new protocols to be easily added to an existing network.

Q—Explain why EIGRP is more efficient in operation than IGRP.

A—Unlike most other distance-vector routing protocols, EIGRP does not mandate a periodic update of routing tables between neighboring routers. Instead, it employs a neighbor discovery and recovery mechanism to ensure that neighbors remain aware of each other's accessibility. As long as a router receives periodic hello packets from its neighbors, it can assume that those neighbors remain functional. More importantly, it can assume that all its routes that rely upon passage through those neighbors remain usable. Thus, EIGRP is much more efficient than conventional distance-vector routing protocols because it imposes much less overhead on routers and transmission facilities during normal operation.

Q—How does RTP enable improved convergence times?

A—RTP is responsible for providing guaranteed delivery of EIGRP packets between neighboring routers. However, not all EIGRP packets that neighbors exchange must be sent reliably. Some packets, such as hello packets, can be sent unreliably. More importantly, they can be multicast rather than having separate datagrams, with essentially the same payload being discretely addressed and sent to individual routers. This helps an EIGRP network converge quickly, even when its links are of varying speeds.

Q—Why does EIGRP tag certain routes?

A—EIGRP supports both internal and external routes. Routes that are internal to an autonomous system are completely contained within that autonomous system. External routes are those that are learned from neighbors that lie outside the autonomous system. External routes are tagged with information that identifies their origin. This enables a network administrator to develop customized interdomain routing policies.

REFERENCES

- . • Pepelnjak, Ivan, EIGRP Network Design Solutions, Indianapolis: Cisco Press, 2000.
- . • Sportack, Mark A., IP Routing Fundamentals, Indianapolis: Cisco Press, 1999.
- . • <http://www.cisco.com/univercd/cc/td/doc/cisintwk/idg4/nd2017.htm>



Corporate Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters

Cisco Systems International
BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters

Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the Cisco Web site at www.cisco.com/go/offices.

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus • Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2004 Cisco Systems, Inc. All rights reserved. CCSP, the Cisco Square Bridge logo, Cisco Unity, Follow Me Browsing, FormShare, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, Packet, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, Registrar, ScriptShare, SlideCast, SMARTnet, StrataView Plus, SwitchProbe, TeleRouter, The Fastest Way to Increase Your Internet Quotient, TransPath, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.