WHITE PAPER

CISCO IOS EIGRP FOR MANAGED SERVICES IN MPLS VPN PE-CE OPERATION

Introduction

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 Enhanced Interior Gateway Routing Protocol (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.

Many enterprise customers are considering managed VPN services such as Layer 3 IP VPNs built on Cisco® IP/Multiprotocol Label Switching (IP/MPLS), available in Cisco IOS® Software. IP/MPLS VPNs deployed using site-to-site Enhanced Interior Gateway Routing Protocol (EIGRP)—rather than external Border Gateway Protocol (BGP) or static routes—eliminate the need for 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. It is highly scalable for large networks, giving bigger enterprises the confidence they need to implement managed services offered by service providers.

This white paper provides an overview of how EIGRP operates when enterprise customers buy Layer 3 MPLS VPN services from MPLS VPN service providers that support EIGRP as a provider edge-customer edge (PE-CE) protocol. This paper assumes that the reader is familiar with the basic workings of MPLS VPNs, VPNv4 support with Multiprotocol BGP (MP-BGP) extensions, and EIGRP.

Summary

An enterprise with many sites can build a private WAN by deploying routers at each site and then interconnecting those routers with a private backbone. If the enterprise actually owns all of the transmission media and switches (which constitute the backbone), it has a truly private network. More commonly, the transmission media and at least some of the backbone switches are owned by a service provider, and this network infrastructure is shared among multiple enterprise networks. In this case, each enterprise network is not a private network, but a VPN.

Layer 3 VPNs using BGP have been the most widely deployed MPLS technology. They use virtual routing instances to create a separate routing table for each subscriber, and use BGP to establish peering relations and to signal the VPN-associated labels with each of the corresponding provider edge routers. This results in a highly scalable implementation—core provider routers are not required to maintain information about the VPNs. BGP VPNs are useful when subscribers want Layer 3 connectivity and wish to offload their routing overhead to a service provider.

Cisco IOS Software supports various PE-CE protocols. These protocols include EIGRP, Open Shortest Path First (OSPF), Routing Information Protocol Version 2 (RIPv2), External BGP (eBGP), and static routes. For IP network routing, EIGRP is well-known to Cisco enterprise customers, who benefit from its greater optimization of path routing, fast convergence, and lower CPU utilization. Deploying 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 securely managed VPNs.

Challenge

Why enterprise customers are looking at MPLS VPNs:

- **Cost**—Depending on the service providers, VPNs can be less costly than buying leased lines or Frame Relay circuits.
- **Less impact to networks**—Customers would like to migrate to VPNs for cost reasons while maintaining their existing routing designs, and have the flexibility to have WAN backbones that consist of both VPNs and leased circuits.

Why service providers support EIGRP as PE-CE:

- **Already in enterprises**—There is a large installed base of enterprise customers with EIGRP.

- **Customer preference**—Enterprise customers who are moving toward buying VPN services from an MPLS provider would prefer to run EIGRP all the way to the service provider's boundary rather than deploying eBGP. In this scenario, customers do not have to learn any other protocol and can preserve the routing environment they are used to.

Solution

Cisco IOS Software makes a secure, highly available, cost-effective managed services environment possible for IP VPNs. 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 Cisco's established global enterprise customer base. Features like EIGRP make routing more efficient and turn IP/MPLS VPNs into a simpler and benefit-rich addition to customer networks.

More than 60 percent of Cisco enterprise customers use EIGRP in their networks. Reasons include:

- Ease of use—Turn it on and EIGRP works as advertised. This is the primary reason that customers use EIGRP.
- Easy to understand/low complexity—The protocol is by far the least complicated to learn and deploy, as compared to any other major interior gateway protocols (IGPs).
- Capability—EIGRP is suitable for deployment in all scenarios (including hub and spoke, broadcast, and nonbroadcast multiaccess [NBMA]). In addition, the protocol is continuously enhanced.
- Scalability—EIGRP scales to the current and future needs of enterprise customers.
- Subsecond convergence—A unique aspect of EIGRP is "feasible successors." Since the backup routes are precomputed, it takes less than a second to converge around such failures.
- Reliability—Due to its reduced complexity, given a good network design, EIGRP does not require a high amount of attention from end users.
- Capital investment—Investment in Cisco products is preserved because EIGRP is widely available across multiple Cisco platforms that are suitable for both enterprises and service providers.

How EIGRP MPLS VPN PE-CE Works

EIGRP MPLS VPN support allows native EIGRP to run on PE-CE links, requiring no upgrade of existing enterprise customer equipment or configurations. All necessary equipment or configuration changes are consolidated to the provider edge routers (Figure 1). BGP redistributes routes into EIGRP using route type and metric information extracted from BGP extended community information.

Figure 1. EIGRP Routes Advertised to BGP Backbone

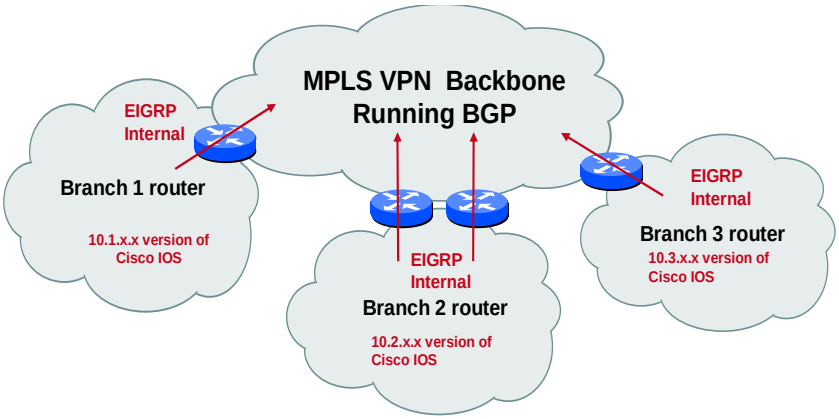

Without EIGRP PE-CE support, normal redistribution of EIGRP into BGP (and vice versa at the provider edge) would result in intersite EIGRP routes appearing as external routes in the target customer edge cloud. The loss of the original route attributes would result in all routes traversing the MPLS VPN backbone becoming less preferable than the routes that do not traverse the MPLS VPN backbone.

To solve this problem, redistribution of EIGRP metrics are preserved across the MPLS VPN backbone though use of MP-BGP extended community attributes. The EIGRP route type and vector metric information is encoded in a series of well-known attributes. These attributes are transported across the MPLS VPN backbone and used to recreate the EIGRP route when received by the target provider edge router.

General Operation

The VPN Backbone as a Transport

The MPLS VPN backbone is treated as another transport to pass EIGRP route information from one customer site to its peering site. EIGRP routes are redistributed into BGP with extended community information appended to the BGP route. BGP then carries this route over the MPLS VPN backbone, with the EIGRP-specific information encoded in the BGP extended community attributes. The EIGRP route information appears as any other MPLS label-encapsulated data within the VPN backbone. Routing protocols within the MPLS VPN backbone have nothing to do with the customer routes.

Once the peering customer site receives the route, BGP redistributes the route into EIGRP. EIGRP then extracts the BGP extended community information and reconstructs the route as it appeared in the original customer site.

Provider Edge Router: Non-EIGRP-Originated Routes

On the provider edge router, if a route is received via BGP and the route has no extended community information for EIGRP, the route will be advertised to the customer edge router as an external EIGRP route using the default metric. If no default metric is configured, the route will not be advertised to the customer edge router.

Provider Edge Router: EIGRP-Originated Internal Routes

If a route is received via BGP and the route has extended community information for EIGRP, the route type is set to “internal” if the source’s autonomous system matches. If the source autonomous system fails to match the configured autonomous system for the given VPN routing and forwarding (VRF) Cisco IOS route table instance, then the rules for non-EIGRP-originated routes will hold.

The internal route is then reconstructed and advertised to the customer edge router as an internal EIGRP route using the extended community information

Provider Edge Router: EIGRP-Originated External Routes

If a route is received via BGP and the route has extended community information for EIGRP, the route type is set to “external” if the source autonomous system matches. If the source autonomous system fails to match the configured autonomous system for the given VRF, the rules for non-EIGRP originated routes will hold.

The external route is then reconstructed and advertised to the customer edge router as an external EIGRP route using the extended community information.

Multiple VRF Support

On a provider edge router, one instance of EIGRP can support multiple EIGRP MPLS VPN VRFs. Support for each VRF translates into its own EIGRP process. The number of EIGRP processes is dependent on the available system resources and the number of supported VRFs on a given platform. There is always an EIGRP process created for the default routing table.

Extended Communities Defined for EIGRP VPNv4

There are no EIGRP adjacencies, EIGRP updates, or EIGRP queries sent across the MPLS VPN backbone. Only EIGRP metric information is carried across the MPLS VPN backbone via the MP-BGP extended communities.

The provider edge router is part of the EIGRP network; therefore, all EIGRP protocol-specific behavior with MPLS VPNs is no different than with any other regular EIGRP network. As mentioned before, the MPLS VPN backbone is treated as a transport.

In order for EIGRP to recreate metrics derived from the originating customer site, the original metric is encoded in MP-BGP extended communities by the provider edge router. These extended communities may then be transported across the MPLS VPN backbone by BGP.

Metric Propagation

Routes are recreated by the provider edge router and sent to the customer edge router as an EIGRP route (Figure 2). The same route type and cost bases as the original route are used to recreate the EIGRP route. The metric of the recreated route is increased by the link cost of the interface. The MPLS VPN backbone or the backdoor link can be made the preferred path by adjusting the metrics.

Figure 2. EIGRP VPN Metric Propagation

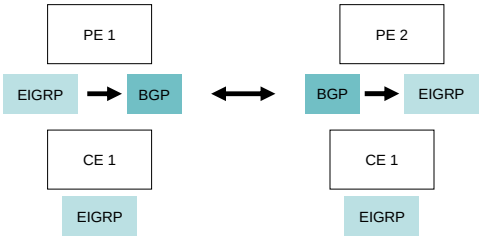

For a given Network X, in the CE1 router, the following happens:

CE1 advertises Network X to PE1 via EIGRP

PE1 EIGRP (VRF-Blue) redistributes Network X to BGP (VRF-Blue)

PE1 BGP (VRF-Blue) requests external attributes from EIGRP (VRF-Blue)

PE1 BGP sends Network X to PE2 with attached extended community attributes

PE2 BGP (VRF-Blue) redistributes Network X to EIGRP (VRF-Blue)

PE2 EIGRP (VRF-Blue) requests external attributes from BGP (VRF-Blue)

PE2 EIGRP (VRF-Blue) rebuilds the route as a native EIGRP route

PE2 advertises Network X to CE2 via EIGRP

CE2 receives Network X from PE2 as a native EIGRP route

Note: At this point, CE2 has the same information as CE1 for Network X

Configuration Highlights

EIGRP configurations on the provider edge router use the VPNv4 address family submode of router mode, similar in lines with BGP VPNv4 address family support.

Address Family for VRF Configuration

EIGRP VPNv4 specific router-mode commands are available in the EIGRP VPNv4 address family submode. This is entered via the following router-mode command:

[no] address-family ipv4 vrf <name>

Note that address family has a number of possible submodes, although only IPv4 is currently supported.

Single Instance EIGRP with Single VPNv4 Example

router EIGRP 1

address-family ipv4 vrf vrf-red

autonomous-system 69

"router sub-mode commands"

exit-address-family

In this example, there are two EIGRP processes created—AS-1 for the default routing table and AS-69 for the “red” routing table.

Example: Single Instance EIGRP with Multiple VPNv4

router EIGRP 1

address-family ipv4 vrf red

autonomous-system 57

"router sub-mode commands"

exit-address-family

address-family ipv4 vrf blue

autonomous -system 46

"router sub-mode commands"

exit-address-family

In this example, there are three EIGRP processes created—AS-1 for the default routing table, AS-57 for the “red” routing table, and AS-46 for the “blue” routing table.

Therefore, EIGRP-VPNv4 is bounded to the scope of the VRF it is configured for.

router EIGRP 42

address-family ipv4 vrf vrf-red

autonomous-system 42

exit-address-family

address-family ipv4 vrf vrf-green

autonomous-system 42

exit-address-family

In the above example, all three EIGRP-VPNv4 processes are unique and will not share neighbors, routing information, or topology information.

Route Redistribution

EIGRP VPNv4 allows route redistribution using the following address family mode command:

[no] redistribute <*protocol*> [metric *[metric-value*]]

EIGRP MPLS VPNs allow for flexibility in the way native and non-native EIGRP routes are handled. This control is accomplished either via the *redistribute* command or the *default-metric* command. The following function is supported:

redistribute BGP <as>

Only BGP routes with BGP extended community information will be distributed into EIGRP. EIGRP looks for BGP extended community information and, if found, uses this information to recreate the original EIGRP route. If the extended community information is missing, EIGRP will not learn the route from BGP.

redistribute BGP <as> metric B D R L M

The configured metric values are used only for BGP routes redistributed into EIGRP. EIGRP looks for BGP extended community information and, if found, uses this information to recreate the original EIGRP route. If the extended community information is missing, EIGRP uses the metric values that have been configured.

default-metric B D R L M [Bandwidth/Delay/Reliability/Load/MTU]

The configured metric values are used for any non-EIGRP routes being redistributed into EIGRP. If a route is a BGP route, EIGRP will look for BGP extended community information and, if found, will use this information to recreate the original EIGRP route. If the extended community information is missing, EIGRP uses the metric values configured.

EIGRP MPLS VPN PE-CE Site of Origin and BGP Cost Community Support for EIGRP MPLS VPN PE-CE with Backdoor Links

The goal of EIGRP Site of Origin (SoO) is to allow an EIGRP network to support complex topologies, such as sites that are connected to each other via backdoor links and MPLS/VPN links, customer edge routers that are dual-homed to different provider edge routers, and provider edge routers supporting customer edge routers from different sites within the same VRF. Path selection within the EIGRP network containing PE-CE links should be based on metrics, allowing either the link through the VPN or the EIGRP backdoor link to act as primary (best) link or act as backup (available if the best should fail) link.

To accomplish this goal, EIGRP is capable of retrieving the SoO attribute on routes redistributed from BGP, filtering routes with certain SoO values, and propagating the SoO values as routes are sent throughout the EIGRP network. Additionally, EIGRP and BGP must be capable of interacting in a way that avoids the normal path selection behavior of BGP, which prefers locally sourced routes over BGP-derived routes. This BGP/EIGRP interaction occurs through the use of BGP cost communities, which is explained in the sections below.

Routers Support EIGRP SoO Feature with SoO Defined on Backdoor Links

If all of the routers in the customer's sites between the provider edge routers and the backdoor routers support the SoO feature, and the SoO values are defined on both the provider edge routers and the backdoor links, the provider edge routers and the backdoor routers will all play a role in supporting convergence across the two (or more) sites. Routers that are not provider edge routers or backdoor routers must only propagate the SoO value on routes as they forward them to their neighbors, but they play no

other role in convergence beyond the normal dual-attachment stations. The next two sections describe the operation of the provider edge routers and backdoor routers in this environment.

Provider Edge Operation

To use the EIGRP SoO feature, the site where the route originates must be identified via the “ip vrf sitemap <route-map>” command applied to the interface of the provider edge router connected to the customer edge router. The syntax of this command and the associated route-map command is included in a later section.

When EIGRP on a provider edge router receives routes from a customer edge router on an interface with an SoO value that has been defined, it checks each route received to determine whether there is an SoO value associated with the route that matches the interface SoO value. If the SoO values match, the route will be filtered before it is placed in the topology table. This is done to stop routing loops (if a route originated in this site is sent across the VPN to another provider edge router attached to this site or another site, and is then relearned through a backdoor or dual-home link). This process will be explained in more detail in the section below.

When EIGRP on the provider edge router receives a route from a customer edge router that doesn’t contain an SoO, value or contains an SoO value that doesn’t match the interface SoO value, the route will be accepted into the topology table so that it can be redistributed into BGP. When the provider edge router redistributes an EIGRP route into BGP and the EIGRP route doesn’t contain an SoO value, the SoO value defined on the interface used to reach the next hop (to the customer edge router) will be included in the extended communities attributes associated with the route. If the EIGRP topology table entry already had an SoO value associated with the route, this value will be included with the route when it is redistributed into the BGP table instead of the interface SoO value. Any BGP peer receiving these prefixes in an update from this BGP speaker will also receive the SoO value associated with each prefix, identifying the site where each prefix originated.

When the EIGRP route is redistributed into the BGP table, EIGRP will also insert a BGP cost community extended community attribute that is derived from the route type and metric. This cost community will be used by BGP in the best path calculation, which is described in more detail below.

When EIGRP on the provider edge router redistributes BGP/VPN prefixes into the EIGRP topology table, it will extract the SoO value (if this has been set) from the BGP extended community attributes and include it with the prefix in the topology table. EIGRP will test the SoO value for each prefix when sending updates to the customer edge routers, filtering any prefixes that contain the same SoO value as defined on the interface connecting to the customer edge router. This filtering is performed to stop transient routing loops due to relearning through the VPN routes that originated in this site.

For routes with SoO values that differ from the local interface’s SoO value, the SoO value retrieved from the BGP table will be included with the route as the provider edge router sends it to the customer edge routers. In this scenario, all routers in the customer sites support the SoO value, so this value will stay associated with the routes as they are sent across the site.

Backdoor Router

A backdoor router is an EIGRP router that connects one site to another but not through the MPLS VPN network. Typically, a backdoor link is used as a backup path between EIGRP sites if the VPN link is down or not available. The metric on the backdoor link would normally be set high enough so that the path through the backdoor will not be selected unless there is a VPN link failure. In this scenario, the backdoor router has the SoO value defined on the backdoor link connecting it to the other site, identifying its local site-id, which should also match the SoO value used on the provider edge routers for the same site (Figure 3).

When a backdoor router receives EIGRP updates (or replies) from a neighbor across the backdoor link (from the backdoor router serving the other site), it checks each received route to verify that it does not contain an SoO value that matches the one defined on the interface. If it finds a route with an SoO value that matches, the route is rejected and not put into the topology table. Typically, the reason that a route would be received with a matching SoO value would be that it was learned by the other site via the VPN connection and advertised back to the original site over the backdoor link. By filtering these routes based on the SoO value at the backdoor link, short term invalid routing is avoided.

In Figure 3, routes originating in Site 1 are tagged with the SoO value 1:1 when the provider edge router redistributes them into BGP. When the routes are redistributed from BGP into EIGRP on PE2 and PE3, the SoO value is pulled out of the BGP table and retained on the routes as they are sent to Site 2 and Site 3. The routes with the SoO value 1:1 will be filtered out when updates are sent from PE2 to CE2, however, stopping them from being relearned in Site 1 via the VPN.

Routes originating in Site 2 will have an SoO value of 1:2 applied as BGP on PE2 redistributes them from EIGRP into the BGP table. They will retain this tag when PE3 redistributes them back into EIGRP. As these Site 2 routes flow across Site 3, they will retain the SoO value of 1:2, so when they reach BD1 they will be recognized as routes originated in Site 2 and filtered out. The reason that this filtering may be desirable is best understood by following the operation in the event of a loss of a route in Site 2. In this case (if an interface goes down), EIGRP will start its normal active process, sending queries to routers in every direction looking for an alterative path to this destination. When the query reaches PE2, PE2 will have no other EIGRP neighbors to query and will remove the route from its routing table and send an infinity reply to CE3, which will then reply to the neighbors that queried it. When the route is removed from the routing table, BGP on PE2 will be notified and will start the withdraw process to remove the route from PE1 and PE3. The BGP withdraw process is slow, however, and may take several minutes to complete.

Figure 3. Backdoor Routing

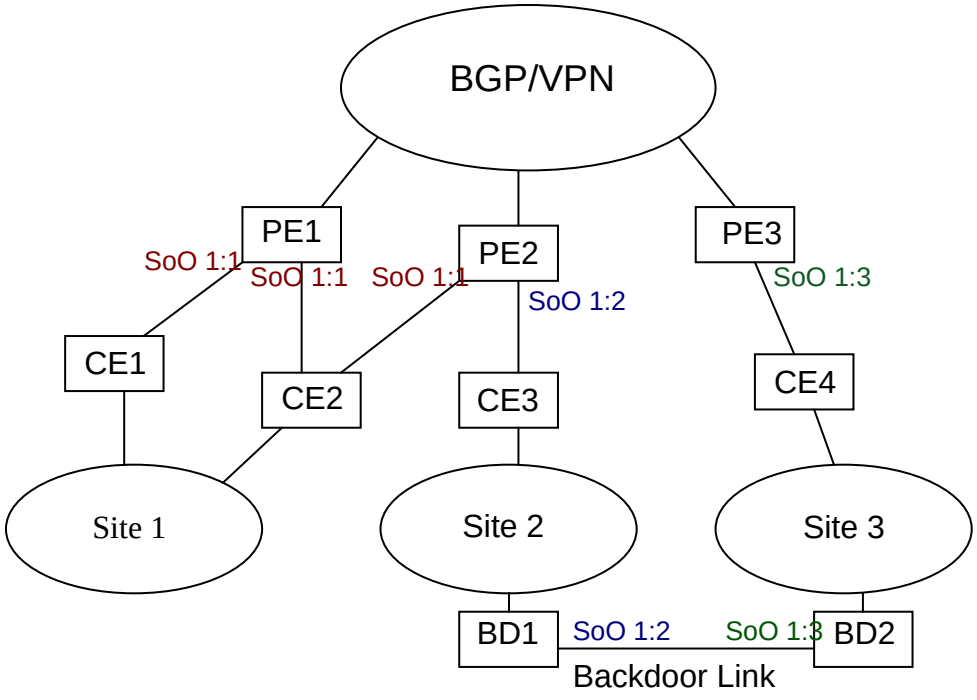

Queries will also flow from the router that lost the interface across Site 2 to the backdoor router BD1, which queries across the backdoor link to BD2. Since it takes much longer for PE2 to withdraw the route from PE3 via BGP than for the EIGRP query

process to reach BD2, BD2 will still have a path to the lost route pointing through Site 3 toward PE3. Because BD1 is filtering routes received on the backdoor link with the SoO value 1:2, the reply from BD2 to BD1 will be converted to infinity, thus not allowing the path to the lost route through Site 3 to be installed in Site 2. The query process in Site 2 will then be completed and the route will be removed from all Site 2 routers.

When the BGP withdraw is finally received on PE3, the route will be removed from the BGP table on PE3. This will remove the route from PE3's routing table and EIGRP on PE3 will go active on the route and send queries to CE4, and then throughout Site 3 looking for an alternative path to the destination. Since there are no alternative paths for the prefix in Site 3, the route will be removed from all routers in Site 3. At this point, the entire network is converged.

Changes to BGP/EIGRP Interaction

The BGP/EIGRP interaction on the provider edge router must be modified to allow the proper path selection when comparing the native EIGRP route and the VPN-sourced route. Before the EIGRP SoO feature was available, BGP always selected a locally sourced route (such as the native EIGRP route redistributed into BGP) over any route learned from a BGP peer. Therefore, if the EIGRP route exists in the BGP table before the same prefix is learned from a BGP peer, the locally sourced route will always win the BGP best-path calculation and the BGP peer-derived route will not be installed in the routing table, regardless of metrics.

The EIGRP SoO support changes the BGP/EIGRP interaction through the use of the BGP cost community extended community attribute. BGP supports putting a special extended community attribute (cost community) at various points in the BGP best-path calculation so that it can be used as a determining factor in the path selection process. The cost community evaluation can be inserted at any point in the best-path calculation, including at the very beginning. By putting the cost community check at the beginning of the best-path calculation, EIGRP can influence path selection before the locally sourced test by populating the cost community attribute as routes are injected into BGP.

To use the cost community check to solve the backdoor problem, EIGRP populates each route with the cost community as the route is injected into the BGP table. The cost community value will be derived from the route type (internal or external) and composite metric. The information is encoded so that internal routes are preferred over external routes and, if the route types are the same, the best composite metric will be determined to be the best path. If the cost community values are the same (same route type and metric), then the BGP best-path calculation will continue and other BGP criteria will be used to determine the best path.

When BGP has a prefix in the BGP table that is locally sourced and it receives the same prefix from a BGP peer, BGP will compare the cost community values of the two paths. The path that has the best (or lowest) cost community value will be selected as the best path.

EIGRP Max-Prefix Limit—Protecting VPNs on a Provider Edge Router

EIGRP's max-prefix limits the number of prefixes that are received from a neighbor or peer and/or the number of prefixes that are received as an aggregate from all the neighbors for a given EIGRP process. The scenarios include prefix limit during normal redistribution and prefix limit in case of a MPLS VPN PE-CE. The goal is to prevent resource starvation, affecting the CPU and memory, on the router because of improper redistribution (accidentally redistributing the entire BGP table into EIGRP, for example) and an excessive number of prefixes learned from a customer edge router or the particular VPN in case of an MPLS VPN PE-CE scenario.

The redistribution limit applies to all IP redistributed prefixes. The user can configure the maximum number of prefixes that can be redistributed into EIGRP and the threshold value (with a default of 75 percent). Syslog messages are generated when the threshold value is reached and when the maximum number of prefixes are reached. The limit is enforced when the configured prefix limit is reached.

Users can choose to run the max-prefix limit in the warning mode only. In this case, only the syslog messages are generated and the number of prefixes is not limited when the max-prefix limit is reached.

Max-Prefix Limit for EIGRP MPLS VPN PE-CE

Currently, the service provider can use “maximum route limit” on a per-VRF basis to restrict the number of routes that are installed in the routing table. However, this does not stop the processing of excessive prefixes from the customer edge router devices. The processing of excessive prefixes could be due to the wrong redistribution on the customer edge side or more than normal growth in a customer’s network. The storage of excessive prefixes in the topology database could lead to CPU and memory starvation, even though one can limit the number of routes in the VRF.

The EIGRP max-prefix limit provides a protection mechanism against resource starvation in case of excess prefixes being received from customer edge routers that are within the same EIGRP process/VPN. As a result, resource starvation due to excessive prefixes will only affect the customer edge routers or VPN/VRF/EIGRP process in question. Customers on other VPNs (and the provider edge router in general) will not be affected because of malfunctioning customer edge routers.

There are two scenarios on the provider edge router:

- Limiting the number of prefixes on a per-EIGRP process/VPN—In this case, the prefix limit applies to the entire EIGRP process/VPN in question and adjacencies are brought down for all the neighbors that belong to that VPN on a given provider edge router. At the same time, when the prefix limit is reached, no further routes are accepted via MP-BGP (from the MPLS VPN backbone).
- Limiting the number of prefixes from a neighbor/customer edge router for a given VPN/EIGRP process—In this case, the adjacency is brought down only for a specific customer edge router that has exceeded the prefix limit. The rest of the customer edge routers that belong to the same VPN on the provider edge router work normally. Here, the prefix limit is configured on a per-neighbor basis.

The options and the network behavior include:

Prefix Limit Configured at the EIGRP Process/VPN Level

- Applies to prefixes that are learned from the directly attached customer edge routers (for a given VPN) and prefixes that are learned via the VPN backbone (MP-BGP or self-generated prefixes on the provider edge router).
- Adjacency brought down with all the customer edge routers in that given VPN on that provider edge router.
- Provider edge router does not learn routes from connected the customer edge routers or via the VPN backbone.

Prefix Limit Configured on a Per-Customer-Edge/Neighbor Basis

- Only applies to prefixes that are learned from the customer edge router in question.
- Adjacency brought down only with that customer edge router.

EIGRP SNMP Support

EIGRP MIBs provide the following five object groups or tables on a per-autonomous system and per-VPN basis:

- EIGRP VPN table
- EIGRP traffic statistics
- EIGRP topology data
- EIGRP neighbor data
- EIGRP interface data

The VRF table ID is incorporated into the indexing of each MIB object within each of these groups. This structure allows for identifying how many VPNs are configured on a router by walking the MIB tree and requesting only the table ID object. The association of a VRF table ID with its configured textual VPN name is made within the VPN table object.

Features

Most enterprise customers who have deployed EIGRP and are moving toward buying the VPN services from an MPLS provider prefer to run EIGRP all the way to other sites rather than eBGP or static routes. They are accustomed to EIGRP and do not have to learn other protocols. Cisco has the complete EIGRP MPLS VPN PE-CE solution, which includes:

- MPLS VPN support for EIGRP between the provider edge and customer edge routers
- EIGRP limits the number of redistributed routes/max-prefix limit support
- EIGRP MPLS VPN PE-CE SoO
- EIGRP SNMP MIB
- BGP cost community support for EIGRP MPLS VPN PE-CE with backdoor links

For platform support and releases, please search the Feature Navigator at:

http://www.cisco.com/go/fn

Conclusion

Cisco has a complete solution for EIGRP MPLS VPN PE-CE. The solution includes support for complex topology designs that include backdoor links between customer sites. Cisco’s EIGRP MPLS VPN PE-CE solution provides an easy migration for enterprise customers moving from private backbones and leased lines to MPLS VPN links and to support networks that are a combination of these links. The solution is robust enough for service providers to protect their customers from a malfunctioning router or VPN. Given the ease of use and reduced complexity of EIGRP, it is easy for both enterprises and service providers alike to deploy this robust, flexible, and manageable solution.

the Cisco Website at www.cisco.com/go/offices.

Printed in the USA