TECHNICAL IMPLEMENTATION GUIDE

CISCO IOS EIGRP IN MANAGED SERVICES SITE-OF-ORIGIN SUPPORT

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.

MPLS VPN PE-CE Scenario

This document describes the functions of EIGRP MPLS VPN provider edge-to-customer edge (PE-CE) site-of-origin (SoO) support. In order to support complex topologies in EIGRP PE-CE networks, EIGRP must be modified to support the SoO attribute as well as modifying the normal path-selection criteria used between BGP and EIGRP on the PE.

Table 1 gives definitions of terms used in this document.

Table 1. Definitions and Acronyms

Term or Acronym	Definition
Backdoor router	Routers that connect two or more sites that also connect to each other through MPLS/VPN EIGRP PE-CE links
Backdoor link	Link connecting two backdoor routers
CE	Customer edge router; the router belonging to the customer network that connects to the PE router to use the MPLS/VPN network
Cost community	BGP extended community that can be inserted anywhere into the best-path calculation
PE	Provider edge router; the entry point into the service provider network; it belongs to and is administered by the service provider, and is the redistribution point between EIGRP and BGP in the PE-CE network
Site	Collection of routers that have well-defined exit points to other "sites;" SoO is a special-purpose tag that identifies the site that injects a route into the network; typically used with MPLS/VPN PE-CE networks

Design Considerations

The goal of the EIGRP VPN Version 4 SoO design is to allow an EIGRP network to support complex topologies, such as MPLS/VPN links between sites with backdoor links, CEs dual-homed to different PEs, and PEs supporting CEs from different sites within the same VPN routing and forwarding (VRF). Path selection within the EIGRP network containing PE-CE links should be based on metric, allowing either the link to go through the VPN or the EIGRP backdoor link to act as the primary (best) link or as the backup (not best, but available to be used if the best fails) link.

In order to accomplish this goal, EIGRP must be 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 and EIGRP interaction takes place through the use of BGP cost communities.

Backward compatibility also must be maintained with existing implementations that are not capable of supporting the SoO attribute. Primarily three types of SoO or backdoor topologies must be supported (note: the PE must support the SoO feature in all the following cases):

- Customer's routers support SoO feature (SoO defined on backdoors)
- Customer's routers support SoO feature (without SoO defined on backdoors)
- Customer's routers do not support SoO feature

This document describes the operation of the SoO feature in each of these types of topologies.

All Routers Support SoO Feature with SoO Defined on Backdoor Links

If all the routers in the customer's sites between the PEs and the backdoor routers support the SoO feature and the SoO values are defined on both the PEs and the backdoor links, the PEs and the backdoor routers all play a role in supporting convergence across

the two (or more) sites. Routers that are not PEs or backdoor routers need to propagate the SoO value only on routes as they forward them to their neighbors, but they play no other role in convergence beyond the normal Diffusing Update Algorithm (DUAL) computations. The next two sections describe the operation of the PEs and backdoor routers in this environment.

PE Operation

To use the EIGRP SoO feature, the site where the route originates must be identified with the **ip vrf sitemap <route-map>** command applied to the interface of the PE connected to the CE.

When EIGRP on a PE receives routes from a CE on an interface with a SoO value defined, it checks each route received to determine whether there is a SoO value associated with the route that matches the interface SoO value. If the SoO values match, the route is filtered before it is placed in the topology table. This is done to stop routing loops if a route is originated in this site, sent across the VPN to another PE attached to this site or another site, and then relearned through a backdoor or dual-home link. This process is explained in more detail in example 1.

When EIGRP on the PE receives a route from a CE that does not contain a SoO value or contains a SoO value that does not match the interface SoO value, the route is accepted into the topology table so that it can be redistributed into BGP. When the PE redistributes an EIGRP route into BGP and the EIGRP route does not contain a SoO value, the SoO value defined on the interface used to reach the next hop (CE) is included in the extended communities attributes associated with the route. If the EIGRP topology-table entry already had a SoO value associated with the route, this SoO value is 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 also receives 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 also inserts a BGP cost community extended community attribute that is derived from the route type and metric. This cost community is used by BGP in the best-path calculation. This process is explained in more detail in the section “Changes to BGP and EIGRP Interaction.”

When EIGRP on the PE redistributes BGP VPN prefixes into the EIGRP topology table, it extracts the SoO value (if set) from the BGP extended community attributes and includes it with the prefix in the topology table. EIGRP tests the SoO value for each prefix when sending updates to the CE routers, filtering any prefixes that contain the same SoO value as defined on the interface connecting to the CE. 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 SoO value, the SoO value retrieved from the BGP table is included with the route as the PE sends it to the CE routers. Because in this scenario all routers in the customer sites support the SoO value, this SoO value stays 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 is not 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 also should match the SoO value used on the PEs for the same site (Figure 1).

When a backdoor router receives EIGRP updates (or replies) from a neighbor across the backdoor link (that is, from the backdoor router serving the other site), it checks each received route to verify that it does not contain a SoO value that matches the one defined on the interface. If it finds a route with a SoO value that matches, the route is rejected and is not put into the topology table.

Typically, the reason that a route would be received with a matching SoO value is if it is learned by the other site through the VPN connection and is being 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. The advantages of this configuration are clear when compared to the operation of the network when the SoO value is not defined on the backdoor links, a situation described in the section “All Routers Support SoO Feature Without SoO on Backdoor Links.” One disadvantage of defining SoO values on the backdoor is that if a site is split into partitioned sections, the path through the VPN, another site, and the backdoor link cannot be used as an alterative path to reach the other half of the local site. An example of this behavior is included in example 1.

Example 1

In Figure 1, routes originating in site 1 are tagged with the SoO value 1:1 when the PE 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 are filtered out when updates are sent from PE2 to CE2, however, stopping them from being relearned in site 1 through the VPN.

Routes originating in site 2 have a SoO value of 1:2 applied as BGP on PE2 redistributes them from EIGRP into the BGP table, and they retain this tag when PE3 redistributes them back into EIGRP. As these site 2 routes flow across site 3, they retain the SoO value of 1:2, so when they reach backdoor router 1 (BD1) they are 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 if the loss of a route occurs in site 2.

When a route is lost in site 2 (for example, an interface goes down), EIGRP starts 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 has no other EIGRP neighbors to query, and removes the route from its routing table and sends an infinity reply to CE3, which then replies to the neighbors that queried it. When the route is removed from the routing table, BGP on PE2 is notified and it starts the withdraw process to remove the route from PE1 and PE3. The BGP withdraw process is very slow, however, and may take several minutes to complete.

Figure 1. Backdoor Router

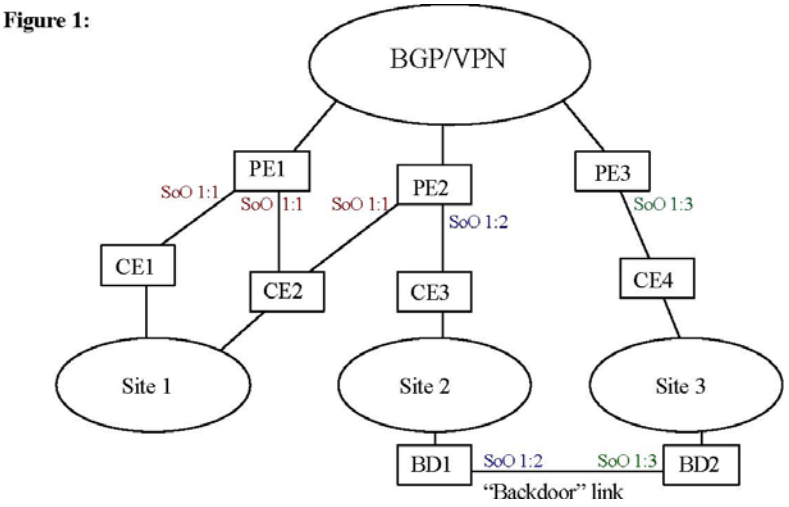

Queries also flow from the router that lost the interface across site 2 to BD1, which queries across the backdoor link to BD2. Because it takes much longer for PE2 to withdraw the route from PE3 with BGP than for the EIGRP query process to reach BD2,

BD2 still has 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 is 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 then completes and the route is removed from all site 2 routers.

When the BGP withdraw is finally received on PE3, the route is removed from the BGP table on PE3. This removes the route from the PE3 routing table, and EIGRP on PE3 goes active on the route and sends queries to CE4 and then throughout site 3 looking for an alternative path to the destination. Because there are no alternative paths for the prefix in site 3, the route is removed from all routers in site 3. At this point, the entire network is converged.

As mentioned in the section “Backdoor Router,” filtering based on SoO on the backdoor router does have a disadvantage. As shown in Figure 2, if one or more routers or links is down in site 2 such that site 2 is split into two partitioned sections, filtering on the backdoor link stops routers in site 2a from being able to reach destinations in site 2b. Because routes tagged with SoO value 1:2 are filtered on BD1, routers within site 2b cannot use this valid, alternative path to “heal” the broken site.

Figure 2. Backdoor Router & Link Failure

Figure 2

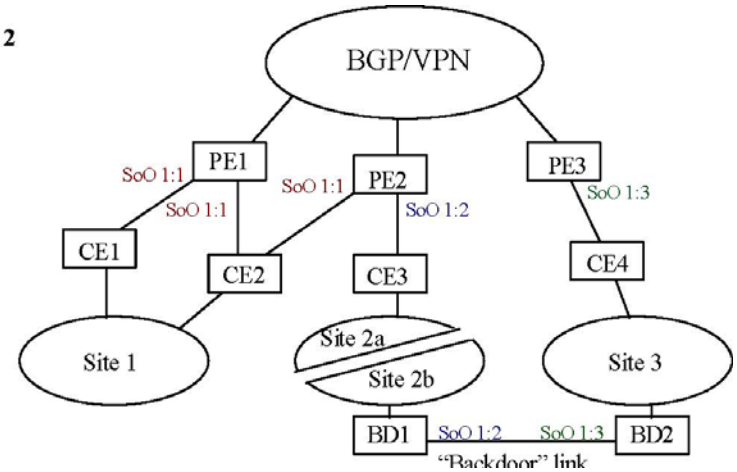

All Routers Support SoO Feature Without SoO on Backdoor Links

If the routers in the customer’s sites support the SoO feature but a SoO value is not defined on the backdoor links, only the PEs perform operations based on the SoO attributes. All other routers in the sites, including the backdoor routers, are responsible only for propagating the SoO values on the routes as they flow through the network, and performing their normal DUAL computations for route convergence.

PE Operation

PE operation in this scenario is virtually identical to that for the previous scenario. The only significant difference is that instead of the backdoor router filtering inbound EIGRP routes containing a SoO value matching the interface SoO value, this responsibility now lies solely with the PE.

When EIGRP on a PE receives routes from a CE on an interface with a SoO value defined, it checks each route received to determine whether there is a SoO value associated with the route that matches the interface SoO value. If the SoO values match, the route is filtered before it is placed in the topology table. This is done to stop routing loops if a route is originated in this site, sent across the VPN to another PE attached to this site or another site, and then relearned through a backdoor or dual-home link.

When EIGRP on the PE receives a route from a CE that does not contain a SoO value or contains a SoO value that does not match the interface SoO value, the route is accepted into the topology table so that it can be redistributed into BGP. When the PE redistributes an EIGRP route into BGP and the EIGRP route does not contain a SoO value, the SoO value defined on the interface used to reach the next hop (CE) is included in the extended communities attributes associated with the route. If the EIGRP topology-table entry already had a SoO value associated with the route, this SoO value is 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 also receives 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 also inserts a BGP cost community extended community attribute that is derived from the route type and metric. This cost community is used by BGP in the best-path calculation, which is described in more detail in the section “Changes to BGP and EIGRP Interaction.”

When EIGRP on the PE redistributes BGP and VPN prefixes into the EIGRP topology table, it extracts the SoO value (if set) from the BGP extended community attributes and includes it with the prefix in the topology table. EIGRP tests the SoO value for each prefix when sending updates to the CE routers, filtering any prefixes that contain the same SoO value as defined on the interface connecting to the CE. 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 SoO value, the SoO value retrieved from the BGP table is included with the route as the PE sends it to the CE routers. Because in this scenario all routers in the customer sites support the SoO value, this SoO value stays associated with the routes as they are sent across the site.

Example 2

In this example, site 2 prefixes have a SoO value of 1:2 associated with them on PE2 as it redistributes the EIGRP routes into the BGP table. When the site 2 routes are redistributed back into EIGRP on PE3, the SoO value of 1:2 is included with the routes. As the site 2 routes flow throughout the site 3 network, the SoO 1:2 value is retained with the routes. Eventually, the site 2 routes reach BD1 through the backdoor link. Because in this scenario a SoO value is not defined on the backdoor link, the updates containing the site prefixes learned in site 3 through the VPN link are advertised back into site 2 through the backdoor link between BD1 and BD2. Because a backdoor link normally has a large metric in order to make it a backup path, routers in site 2 typically select the paths within site 2 and do not use the path through the backdoor link to reach other site 2 prefixes through the VPN.

If a route loss occurs in site 2, the operation is slightly different than that in example 1. In this case (without the SoO on the backdoor link), the beginning of the process is the same. When the route is lost on a router in site 2, queries are sent looking for an alternative path. When the queries reach PE2, it replies with infinity and removes the route from the routing table. BGP is then notified of the route loss and starts the withdraw process to remove the prefix from PE1 and PE3 (Figure 3).

Figure 3. Backdoor Router & Healing Process

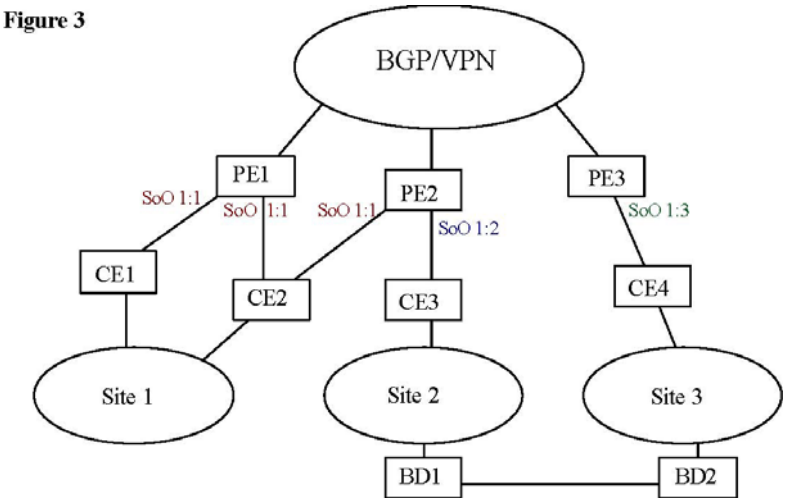

"Backdoor" link

When the query reaches BD2 across the backdoor link, however, the process changes from the description in example 1. BD2 replies with a valid metric based on its path to the prefix through site 3 and PE3. Because a SoO value is not defined on the interface of BD1 facing BD2, BD1 accepts the reply, installs a route to the lost prefix pointing at BD2, and sends a reply to the router that sent it the query. This reply process continues throughout the rest of site 2, until all the routers in site 2 (except PE2) have a route to the lost prefix pointing in the direction of BD1 and site 3.

PE2 does not accept this reply, however, because it filters any received updates (and replies) for routes containing the SoO value of 1:2, which this route learned from site 3 contains. This filtering is useful for significantly reducing the time required to correctly converge on the lost route. This difference is evident in the next section dealing with scenarios in which the customer routers do not support SoO.

Eventually, the BGP withdraw sent earlier from PE2 to PE3 is received and processed on PE3, causing PE3 to remove the prefix from the BGP table and routing table. This causes EIGRP to go active on the route, sending queries from PE3 to CE4 and the rest of site 3 and across the backdoor link into (and throughout) site 2, until all routers in both sites have been queried. When all the replies have been received, the lost route is removed from site 2 and site 3.

This temporary invalid routing was avoided in example 1 by defining the SoO value on the backdoor link. Without the SoO causing filtering of updates and replies on the backdoor link, however, the limitation stated in the first section (and demonstrated in Figure 2) is also removed.

If site 2 becomes partitioned, with half reachable by PE2 and the other half reachable by BD1, the path through the VPN and site 3 can be used to "heal" the partitioned site.

Customer Routers Do Not Support SoO Feature

The SoO feature also can be useful if the customer networks are not running a version of Cisco IOS Software that supports the SoO feature. The capability of the SoO feature is significantly decreased in this scenario, but it may meet the needs of many customers who cannot upgrade to a newer version of Cisco IOS Software with SoO support. In this case, the PEs are the only routers involved in the SoO operation, and the remainder of the network is responsible only for performing the DUAL computations for normal route convergence.

PE Operation

The PE operation in this scenario is identical to that in the other two, with the exception that it never receives routes from the CE with a SoO value associated with them, so it has nothing to filter. Routes still have the SoO value from the local interface applied to them before installation in the BGP table when redistributing from EIGRP into BGP, and the SoO value also is extracted from the BGP extended community attributes and added to the route in the EIGRP topology table when redistributing from BGP into EIGRP. EIGRP also filters routes redistributed from BGP if the SoO value on the routes matches the SoO defined on the local interface, as described in the earlier sections. The EIGRP and BGP interaction through the use of cost communities still allows EIGRP to support complex topologies, but with significant degradation in convergence when routes are lost.

Example 3

Because of the EIGRP imposition of the BGP cost community and modifications done to Administrator (Admin) Distance comparison routines in the SoO feature, complex topologies are supported even if the SoO feature is not supported in the customer sites. If routes from site 3 are propagated to site 2 through the VPN link as well as the backdoor link, the path with the best metric is chosen for traffic from site 2 to site 3. Because the backdoor link is normally set to a higher metric, paths through the VPN link are typically preferred over the backdoor link. When a route is lost, however, the link through the backdoor can be used as an alterative. Because of the lack of SoO support in the customer networks, convergence is more complicated than the other two scenarios.

If a route in site 2 is lost (interface down), the process starts just like in the previous examples. Queries are sent in all directions from the router that lost the interface looking for an alternative path to the destination. When the query reaches PE2, EIGRP has no other neighbors to query, so the route is deleted from the topology table and routing table and an infinity reply is sent back toward the site 2 router that started the query process. Meanwhile, BGP on PE2 is notified that the route was removed from the routing table and starts the withdraw process from PE1 and PE3. As stated earlier, this withdraw process can take minutes.

At the same time, queries flow though site 2 to BD1, and across the backdoor link to BD2. Because BD2 still has a path to the prefix through site 3 (because the BGP withdraw process has not yet completed), BD2 replies to BD1 with its metric through site 3. The reply with the path through site 3 propagates its way through site 2 until all routers in site 2 have installed a route to the lost prefix pointing in the direction of BD1 and site 3 (Figure 4).

Figure 4. Backdoor Router & Route Redistribution

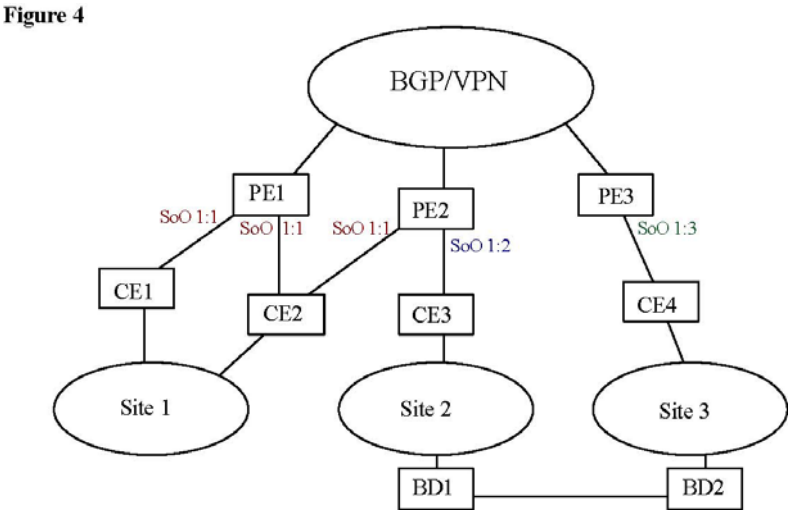

"Backdoor" Link

In contrast to example 2, PE2 receives the reply initially generated by BD2 and installs the route for the lost prefix into the routing table, pointing at CE3 as its next hop in the path through the backdoor link. The installation of this invalid route on PE2 also causes BGP to get notified of a "new" route, which BGP installs in its BGP table and eventually sends to PE3.

While this reply process in site 2 is happening, however, the previous BGP withdraw arrives and is processed by PE3, and the prefix is withdrawn from the PE3 BGP table and routing table. This route deletion causes EIGRP on PE3 to go active, sending queries into site 3 looking for an alternative path for this prefix. As this new query is making its way across site 3 and the backdoor link into site 2, the BGP update from PE2 to PE3, based on the installation of the invalid path that site 2 learned through the backdoor, is received and processed on PE3. This provides PE3 with another "new" path to the lost prefix through the VPN link, but with a worsened metric. This new route is redistributed into EIGRP and updates are sent through site 3, through the backdoor link into site 2. This causes all the routers in site 2 and site 3 to again install a route for the lost prefix pointing toward PE3.

That is, of course, until the BGP withdraw sent earlier (when the first query from PE3 through site 3 and site 2) is received and processed on PE3. PE3 then removes the route from BGP, removing it from EIGRP, which goes active, and this process continues until the route reaches infinity metric or 100 hop count.

This process demonstrates a typical redistribution loop generated when there is mutual redistribution between dissimilar routing protocols (particularly with distance-vector protocols.) The behavior is typically known as "counting to infinity." Because the SoO tag is not propagated through the networks (because the customer sites do not support the feature), the routing redistribution loop is not terminated until infinity is reached. The SoO tag used in the sections "All Routers Support SoO Feature with SoO on Backdoor Links" and "Customer Routers Do Not Support SoO Feature" allows EIGRP to recognize that a route is not new but is a route that originated in this site, and cannot be used as the source of a new update.

In this last case, where customer routers cannot support SoO, some other method must be used to minimize the impact of this redistribution behavior. One way to minimize the impact is to decrease the size of infinity. EIGRP defaults to a maximum hop count of 100, but this number can be changed using the following command on the PE:

router eigrp <as> address-family ipv4 vrf green

metric maximum-hops <#>

By decreasing the maximum hop count to a smaller number, “infinity” is reached much sooner and the redistribution loop is not sustained as long. The hop count should be set to a number large enough to account for the maximum path length though the site, including any extra hops that may be necessary in failure conditions. Note that using this technique does not eliminate the problem. Instead, it lessens the impact of the problem.

Changes to BGP and EIGRP Interaction

The BGP and EIGRP interaction on the PE must be modified in order to allow the proper path selection when comparing the native EIGRP route and the VPN-sourced route. Before the EIGRP SoO feature, 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 always wins the BGP best-path calculation and the BGP peer-derived route is not installed in the routing table, regardless of metric.

The EIGRP SoO support changes the BGP and EIGRP interaction through the use of the BGP cost community extended community attribute. BGP supports the capability of 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 the 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.

In order 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 is derived from the route type (internal and 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 is 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 continues and other BGP criteria are used to determine the best path.

The cost community attribute is supplied in the form Cost:POI:ID:value. The cost is the well-known cost community value (0x4301). The point of insertion (POI) is defined as the “absolute point of insertion” (128). The ID value is set to 128 for internal routes and 129 for external routes. Because BGP prefers the lower-cost community ID, if a prefix is learned both as an internal and external route, the internal route wins. And finally, the value portion of the cost community is set to the composite metric on the router where the redistribution into BGP is taking place.

Now, when BGP has a prefix in the BGP table that is locally sourced and it receives the same prefix from a BGP peer, BGP compares the cost community values of the two paths. Whichever path has the best (lowest) cost community value is selected as the best path.

End-User Interface

No command-line interface (CLI) changes are necessary to implement the EIGRP SoO feature. The existing **ip vrf sitemap SoO** interface command sets the SoO value associated with routes the PE receives through the interface connecting to the CE. This command also is used on the backdoor link to define the routes that belong to this site.

interface e0/0 ip vrf sitemap SoO

route-map SoO permit 10 set extcommunity SoO 1:1

As mentioned in example 3, it also is useful to decrease the maximum hop count used by EIGRP if the customer routers do not support the SoO feature. Refer to example 3 for suggestions about how to determine the correct hop count to use. The command to decrease the hop count follows:

router eigrp <as> address-family ipv4 vrf green

metric maximum-hops <#>

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