



EIGRP Support for 6PE/6VPE

The EIGRP Support for 6PE/6VPE feature enables native IPv6 Enhanced Interior Gateway Routing Protocol (EIGRP) routes to preserve their original characteristics (metric and other attributes like type, delay, bandwidth, and maximum transmission unit [MTU]) while being redistributed from one IPv6 EIGRP site to another over a service-provider VPN cloud or an IPv6 provider edge (6PE) Multiprotocol Label Switching-VPN (MPLS-VPN) network. The Border Gateway Protocol (BGP) is used as the external routing protocol to transfer IPv6 EIGRP routes across the VPN cloud or the 6PE MPLS-VPN network. This module explains the EIGRP 6PE/6VPE feature.

- [Information About EIGRP Support for 6PE/6VPE, page 1](#)
- [Additional References for EIGRP Support for 6PE/6VPE, page 5](#)
- [Feature Information for EIGRP Support for 6PE/6VPE, page 5](#)


Information About EIGRP Support for 6PE/6VPE

EIGRP Support for 6PE/6VPE Overview

The EIGRP Support for 6PE/6VPE feature enables native IPv6 Enhanced Interior Gateway Routing Protocol (EIGRP) routes to preserve their original characteristics (metric and other attributes like type, delay, bandwidth, and the maximum transmission unit [MTU]) while being redistributed from one IPv6 EIGRP site to another over a service-provider VPN cloud or an IPv6 provider edge (6PE) Multiprotocol Label Switching-VPN (MPLS-VPN) network. The EIGRP 6PE/6VPE implementation allows native IPv6 EIGRP routes to run on provider-edge (PE) and customer-edge (CE) devices while preserving their original characteristics. MPLS

supports only IPv4. Therefore, the Border Gateway Protocol (BGP) is used to redistribute IPv6 routes over a service-provider VPN cloud or a 6PE MPLS-VPN network.

Figure 1: EIGRP 6PE/6VPE Functionality



The figure above shows two EIGRP sites connected through a service-provider VPN cloud. The two sites are also connected through a backdoor link, which is a backdoor device that is not part of the VPN network. Typically, a backdoor link is a high-cost, low-speed satellite link that is used as a backup path between the EIGRP sites in case the primary VPN is down or unavailable. The CE devices in the EIGRP sites are connected to the service-provider PE devices. The ideal path for both the EIGRP sites to connect with each other is through the VPN cloud. Before the EIGRP Support for 6PE/6VPE feature, a route redistributed from the CE in EIGRP site 1, through the PEs in the VPN cloud, into EIGRP site 2 would appear as an external route in the target CE cloud. This behavior is due to the route losing its original route attributes when it traverses through an external cloud. External routes have higher administrative distances than internal routes. Therefore, the EIGRP route that now appears as an external route in the target CE cloud has a higher administrative distance than the backdoor link, which is an internal route. The external EIGRP route becomes the less preferred route, resulting in the backdoor link becoming the primary route between the EIGRP sites, which is an undesirable behavior.

The EIGRP Support for 6PE/6VPE feature enables EIGRP route metrics to be preserved across a service-provider VPN network or a 6PE MPLS-VPN backbone through the use of BGP Extended Communities. When an EIGRP route with preserved metrics is received by the target PE device, the route is advertised to the receiving CE device as an internal route, along with the original metric and type. Because an internal route

has a lower administrative distance than an external route, the internal route becomes the preferred route of communication between the two EIGRP sites.

BGP Extended Communities

For the Enhanced Interior Gateway Routing Protocol (EIGRP) to recreate route metrics derived from the originating customer site, the original metrics are encoded into Border Gateway Protocol (BGP) Extended Communities by the provider-edge (PE) device that receives the routes from the transmitting customer-edge (CE) device. These extended communities are then transported across the Multiprotocol Label Switching-VPN (MPLS-VPN) backbone by BGP from one customer site to the other (peering customer site). After the peering customer site receives the routes, BGP redistributes the routes into EIGRP. EIGRP, then, extracts the BGP Extended Community information and reconstructs the routes as they appeared in the original customer site.

The following rules govern BGP Extended Communities:

Non-EIGRP-Originated Routes: If a non-EIGRP-originated route is received through BGP and the route has no extended community information for EIGRP, BGP advertises the route to the receiving CE as an external EIGRP route by using the route's default metric. If no default metric is configured, BGP does not advertise the route to the CE.

EIGRP-Originated Internal Routes: If an EIGRP-originated internal route is received through BGP and the route has extended community information for EIGRP, the PE sets the route type to "internal" if the source autonomous system number matches the autonomous system number configured for this VPN routing and forwarding (VRF) instance. BGP, then, reconstructs and advertises the route to the receiving CE as an internal EIGRP route by using the extended community information. If there is no autonomous system match, these routes are treated as non-EIGRP-originated routes.

EIGRP-Originated External Routes: If an EIGRP-originated external route is received through BGP and the route has extended community information for EIGRP, the PE sets the route type to "external" if the source autonomous system number matches the autonomous system number configured for this VRF instance. BGP, then, reconstructs and advertises this external route to the receiving CE as an external EIGRP route by using the extended community information. If there is no autonomous system match, these routes are treated as non-EIGRP-originated routes.

Preserving Route Metrics

The EIGRP 6PE/6VPE feature manages native and non-native Enhanced Interior Gateway Routing Protocol (EIGRP) routes by using the **redistribute** and the **default metric** commands, respectively. By using the **redistribute bgp as-number** command, you can ensure that only Border Gateway Protocol (BGP) routes with BGP Extended Community information are distributed into EIGRP. EIGRP uses this information to recreate the original EIGRP route. If the BGP Extended Community information is missing and the default metric is not specified, EIGRP will not learn the route from BGP.

By using the **redistribute bgp as-number metric-type type-value** command, you can ensure that the metric values configured using this command are used only for BGP routes redistributed into EIGRP. EIGRP looks for BGP Extended Community information, and if this information is found, EIGRP uses this information to recreate the original EIGRP route. If the Extended Community information is missing, EIGRP uses the metric values configured using this command to determine whether the route is the preferred route.

By using the **default-metric bandwidth delay reliability loading mtu** command, you can ensure that the metric values configured using this command are used for any non-EIGRP routes being redistributed into EIGRP. If the received route is a BGP route, EIGRP looks for BGP Extended Community information, and if this

information is found, EIGRP uses this information to recreate the original EIGRP route. If the extended community information is missing, EIGRP uses the metric values configured to determine whether the route is the preferred route.

EIGRP 6PE/6VPE SoO

The EIGRP 6PE/6VPE Site of Origin (SoO) functionality allows an Enhanced Interior Gateway Routing Protocol (EIGRP) network to support complex topologies, such as Multiprotocol Label Switching-VPN (MPLS-VPN) links between sites with backdoor links, customer-edge (CE) devices that are dual-homed to different provider-edge (PE) devices, and PEs supporting CEs from different sites within the same VPN routing and forwarding (VRF) instance. Path selection within the EIGRP network containing PE-CE links is based on route metrics that allow either the link through the VPN or the EIGRP backdoor to act as the primary (best) link or the backup link, if the primary link fails. EIGRP accomplishes this path selection by retrieving the Site of Origin (SoO) attribute from routes redistributed from the Border Gateway Protocol (BGP) network. This BGP/EIGRP interaction takes place through the use of the BGP Cost Community Extended Community attribute.

When routes are redistributed into EIGRP from a BGP network, BGP Cost Community Extended Community attributes are added to the routes. These attributes include the SoO attribute. The SoO attribute is used to identify the site of origin of a route and prevent advertisement of the route back to the source site. To enable the EIGRP SoO functionality, you must configure the **ip vrf sitemap** command on the PE interface that is connected to the CE device. This command enables SoO filtering on the interface. When EIGRP on the PE device receives CE routes on the interface that has a SoO value defined, EIGRP checks each route 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. This filtering is done to stop routing loops.

When EIGRP on the PE receives a route that does not contain an SoO value or contains an SoO value that does not match the interface SoO value, the route will be accepted into the topology table so that it can be redistributed into BGP. When the PE redistributes an EIGRP route that does not contain an SoO value into BGP, the SoO value that is defined on the interface used to reach the next hop (CE) is included in the Extended Communities attribute associated with the route. If the EIGRP topology table entry already has an SoO value associated with the route, this SoO value, instead of the interface SoO value, will be included with the route when it is redistributed into the BGP table. Any BGP peer that receives these prefixes will also receive the SoO value associated with each prefix, identifying the site, where each prefix originated.

The EIGRP SoO functionality ensures that BGP does not follow its normal path-selection behavior, where locally derived routes (such as native EIGRP routes redistributed into BGP) are preferred over BGP-derived routes.

For more information on the Site of Origin functionality, see the “EIGRP MPLS VPN PE-CE Site of Origin” chapter in the *IP Routing: EIGRP Configuration Guide*.

Backdoor Devices

Backdoor devices are EIGRP devices that connect one EIGRP site to another, but not through the Multiprotocol Label Switching-VPN (MPLS-VPN) network. Typically, a backdoor link is used as a backup path between peering EIGRP sites if the MPLS-VPN link is down or unavailable. The metric on the backdoor link is set high enough so that the path through the backdoor will not be selected unless there is a VPN link failure. You can define Site of Origin (SoO) values on the backdoor device on interfaces connecting the device to the peering sites, thus identifying the local-site identity of the link.

When a backdoor device receives EIGRP updates or replies from a neighbor, the device checks each received route to verify that the route does not contain an SoO value that matches the ones defined on its interfaces. If the device finds a route with a SoO value that matches the value defined on any of its interfaces, the route is rejected and not included in the topology table. Typically, the reason that a route is received with a matching SoO value is that the route is learned by the other peering site through the MPLS-VPN connection and is being advertised back to the original site over the backdoor link. By filtering such routes based on the SoO value defined on the backdoor link, you can avoid short-term, invalid routing.

Additional References for EIGRP Support for 6PE/6VPE

Related Documents

Related Topic	Document Title
Cisco IOS commands	Cisco IOS Master Command List, All Releases
EIGRP commands	Cisco IOS IP Routing: EIGRP Command Reference
EIGRP FAQs	EIGRP Frequently Asked Questions
EIGRP technology white papers	Enhanced Interior Gateway Routing Protocol

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

Feature Information for EIGRP Support for 6PE/6VPE

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 1: Feature Information for EIGRP Support for 6PE/6VPE

Feature Name	Releases	Feature Information
EIGRP Support for 6PE/6VPE	15.3(2)S	<p>The EIGRP Support for 6PE/6VPE feature enables native IPv6 Enhanced Interior Routing Protocol (EIGRP) routes to preserve their original characteristics while being redistributed from one IPv6 EIGRP site to another over a service-provider VPN cloud or an IPv6 Provider Edge (6PE) Multiprotocol Label Switching-VPN (MPLS-VPN) network.</p> <p>No commands were introduced or modified by this feature.</p>