



MPLS VPN--BGP Local Convergence

Last Updated: December 15, 2011

This document provides information about reducing the downtime of a provider edge (PE) to customer edge (CE) link failure. It describes how to reroute PE-egress traffic onto a backup path to the CE before BGP has reconverged. The MPLS VPN--BGP Local Convergence feature is also referred to as “local protection.”

This document explains how to use PE-CE local convergence. For information on using BGP PIC Edge for BGP local convergence support, see BGP PIC Edge for IP and MPLS-VPN.



Note

The MPLS VPN--BGP Local Convergence feature affects only traffic exiting the Virtual Private Network (VPN). Therefore, it cannot fully protect traffic end-to-end by itself.

- [Finding Feature Information, page 1](#)
- [Prerequisites for MPLS VPN--BGP Local Convergence, page 2](#)
- [Restrictions for MPLS VPN--BGP Local Convergence, page 2](#)
- [Information About MPLS VPN--BGP Local Convergence, page 2](#)
- [How to Enable MPLS VPN--BGP Local Convergence, page 4](#)
- [Configuration Examples for MPLS VPN--BGP Local Convergence, page 8](#)
- [Additional References, page 14](#)
- [Feature Information for MPLS VPN--BGP Local Convergence, page 15](#)

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest feature information and caveats, see the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the Feature Information Table at the end of this document.

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.



Americas Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

Prerequisites for MPLS VPN--BGP Local Convergence

- Before MPLS VPN --BGP Local Convergence link protection can be enabled, the customer site must be connected to the provider site by more than one path.
- Both the main forwarding path and the redundant backup path must have been installed within Border Gateway Protocol (BGP), and BGP must support lossless switchover between operational paths.
- Any of the supported routing protocols can be used between the PE and CE as long as the path is redistributed into BGP. The supported protocols for IPv4 are External BGP (eBGP), Routing Information Protocol (RIP), Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), and static routing. The supported protocols for IPv6 are External BGP (eBGP) and static routing.
- All PE routers that are serving as backup to the link must have assigned a unique Route Distinguisher to each VRF table involved with the link to ensure that the route reflectors advertise all available paths.
- Although not required, it is recommended that the backup PE (shown as “PE2” in the figure below) runs the same Cisco IOS release that is running on the PE (“PE1”) whose link with the CE will be protected; that is, Cisco IOS Release 12.2(33) SRC, 12.2(33)SB, Cisco IOS 15.0(1)M, Cisco IOS 15.0(1)S, or a more recent version of those products.

Restrictions for MPLS VPN--BGP Local Convergence

- The MPLS VPN--BGP Local Convergence feature affects only traffic exiting the VPN. Therefore, it cannot fully protect traffic end-to-end by itself.
- This link protection cannot be initiated *during* a high availability (HA) stateful switchover (SSO). But links already configured with this protection *before* the switchover begins will remain protected after the switchover.
- If you perform an in-service software downgrade from an image that does include this link protection to an image that does not support this feature, active protection will be halted when BGP routes are refreshed.
- Any next-hop core tunneling technology that is supported by BGP is also supported for protection, including Multiprotocol Label Switching (MPLS), IP/Layer 2 Tunneling Protocol version 3 (L2TPv3), and IP/generic routing encapsulation (GRE). Enabling a Carrier Supporting Carrier (CsC) protocol between the PE and CE is also supported. Interautonomous system option A (back-to-back virtual routing and forwarding (VRF)) is supported because it is essentially the same as performing the PE-CE link protection in both autonomous systems. However, interautonomous system options B and C protection are not supported.
- The MPLS VPN--BGP Local Convergence feature for IPv4 supports the eBGP, RIP, EIGRP, OSPF, and static routing protocols only.
- The MPLS VPN--BGP Local Convergence feature for IPv6 supports the eBGP and static routing protocols only.

Information About MPLS VPN--BGP Local Convergence

- [How Link Failures Are Handled with BGP](#), page 3
- [How Links Are Handled with the MPLS VPN--BGP Local Convergence Feature](#), page 3

- [How Link Failures Are Detected, page 4](#)

How Link Failures Are Handled with BGP

Within a Layer 3 VPN network, the failure of a PE-CE link can cause a loss of connectivity (LoC) to a customer site, which is detrimental to time-sensitive applications. Several factors contribute to the duration of such an outage:


- The time to detect the failure
- The programming of the forwarding
- The convergence of BGP (in large networks, the restored traffic arrival time at its destination varies according to the prefix)

When BGP detects a PE-CE link failure, it removes all of the BGP paths through the failing link. BGP runs the best-path algorithm on the affected prefixes and selects alternate paths for each prefix. These new paths (which typically include a remote PE) are installed into forwarding. The local labels are removed and BGP withdrawals are sent to all BGP neighbors. As each BGP neighbor receives the withdrawal messages (typically indirectly using routereflectors), the best-path algorithm is called and the prefixes are switched to an alternate path. Only then is connectivity restored.

How Links Are Handled with the MPLS VPN--BGP Local Convergence Feature

The MPLS VPN--BGP Local Convergence feature requires that the prefixes to be protected on a PE-CE link have at least one backup path that does not include that link. (See the figure below.) The customer site must have backup paths to the provider site.

Figure 1 Network Configured with Primary and Backup Paths




The MPLS VPN--BGP Local Convergence feature reduces LoC time by sending the broken link's traffic over a backup path (as shown in the figure below) instead of waiting for total network convergence. The local label is maintained for 5 minutes while prefixes switch from the failing local path to the backup path. Because the label is not freed as had been the usual practice, forwarding continues to take place.

The best-path algorithm selects the backup path. Thus, the local label has been applied in place of the failed BGP best-path label (which is sometimes called "label swapping"). Traffic is restored locally while the

network propagation of the BGP withdrawal messages takes place. Eventually, the egress PE router converges and bypasses the local repair.

Figure 2 Network Using the Backup Path After a PE-CE Link Failure on the Primary Path



Note

After the 5-minute label preservation, the local labels are freed. Any BGP prefix that is remote and is not part of a CsC network does not have a local label and is removed. The delay in local label deletion does not modify normal BGP addition and deletion of BGP paths. Rather, BGP reprograms the new backup bestpath into forwarding as usual.

How Link Failures Are Detected

Local protection relies on BGP being notified of the interface failure. Detection can occur using either the interface drivers or the routing tables. If an interface or route goes down, the corresponding path in the routing table is removed and BGP will be notified using the routing application programming interfaces (APIs).

However, when the routing table cannot detect the failure (as when a Layer 2 switch goes down), BGP determines that a neighbor is down through use of its hold-down timer. However, that determination can be extremely slow because of the 3-minute default for BGP session timeout.

You can reduce the detection delay by either reducing the BGP session timeout interval (as described in the *Configuring Internal BGP Features* document) or by enabling the Bidirectional Forwarding Detection (BFD) protocol within eBGP between the PE and CE. For complete instructions to enable BFD, see the *Bidirectional Forwarding Detection* document.

How to Enable MPLS VPN--BGP Local Convergence

**Note**

To configure a VPN routing and forwarding (VRF) instance for IPv4 and IPv6 VPNs or to upgrade your existing single-protocol IPv4-only VRF to a multiprotocol VRF configuration, see MPLS VPN--VRF CLI for IPv4 and IPv6 VPNs.

- [Configuring MPLS VPN--BGP Local Convergence with IPv4](#), page 5
- [Configuring MPLS VPN--BGP Local Convergence with IPv6](#), page 6
- [Troubleshooting Tips](#), page 8

Configuring MPLS VPN--BGP Local Convergence with IPv4

Ensure that the CE is already connected to the PE by a minimum of two paths.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **ip vrf *vrf-name***
4. **rd *route-distinguisher***
5. **protection local-prefixes**
6. **do show ip vrf detail**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Router> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Router# configure terminal	Enters global configuration mode.
Step 3	ip vrf <i>vrf-name</i> Example: Router(config)# ip vrf vpn1	Enters VRF configuration mode. <ul style="list-style-type: none"> • If no VRF routing table and Cisco Express Forwarding table had been previously created for this named VRF, then this command also creates them, giving both tables the specified value for the <i>vrf-name</i> argument (in this example, the name is vpn1).

Command or Action	Purpose
<p>Step 4 <code>rd route-distinguisher</code></p> <p>Example:</p> <pre>Router(config-vrf)# rd 100:3</pre>	<p>(Optional) Establishes the route distinguisher for the named VRF.</p> <ul style="list-style-type: none"> If no route distinguisher had been previously established for the named VRF, then you must enter this command. The route distinguisher value can be either an: <ul style="list-style-type: none"> Autonomous system number followed by a colon and an arbitrary number (for example, 100:3) <p>or</p> <ul style="list-style-type: none"> IP address followed by a colon and an arbitrary number (for example, 192.168.122.15:1)
<p>Step 5 <code>protection local-prefixes</code></p> <p>Example:</p> <pre>Router(config-vrf)# protection local-prefixes</pre>	<p>Allows a preconfigured backup path to carry traffic if the PE-CE link breaks by preserving the local prefixes while BGP reconverges.</p>
<p>Step 6 <code>do show ip vrf detail</code></p> <p>Example:</p> <pre>Router(config-vrf)# do show ip vrf detail</pre>	<p>(Optional) Verifies that the MPLS VPN--BGP Local Convergence feature has been configured.</p>

Configuring MPLS VPN--BGP Local Convergence with IPv6

Ensure that the CE is already connected to the PE by a minimum of two paths.

SUMMARY STEPS

1. `enable`
2. `configure terminal`
3. `vrf definition vrf-name`
4. `rd route-distinguisher`
5. `address-family [ipv4 | ipv6]`
6. `protection local-prefixes`
7. `do show ip vrf detail`

DETAILED STEPS

Command or Action	Purpose
<p>Step 1 enable</p> <p>Example:</p> <pre>Router> enable</pre>	<p>Enables privileged EXEC mode.</p> <ul style="list-style-type: none"> • Enter your password if prompted.
<p>Step 2 configure terminal</p> <p>Example:</p> <pre>Router# configure terminal</pre>	<p>Enters global configuration mode.</p>
<p>Step 3 vrf definition <i>vrf-name</i></p> <p>Example:</p> <pre>Router(config)# vrf definition vrf2</pre>	<p>Enters VRF configuration mode.</p> <ul style="list-style-type: none"> • If no VRF routing table and Cisco Express Forwarding table had been previously created for this named VRF, then this command also creates them, giving both tables the specified value for the <i>vrf-name</i> argument (in this example, the name is vrf2).
<p>Step 4 rd <i>route-distinguisher</i></p> <p>Example:</p> <pre>Router(config-vrf)# rd 100:3</pre>	<p>(Optional) Establishes the route distinguisher for the named VRF.</p> <ul style="list-style-type: none"> • If no route distinguisher had been previously established for the named VRF, then you must enter this command. • The route distinguisher value can be either an: <ul style="list-style-type: none"> ◦ Autonomous system number followed by a colon and an arbitrary number (for example, 100:3) <p>or</p> <ul style="list-style-type: none"> ◦ IP address followed by a colon and an arbitrary number (for example, 192.168.122.15:1)
<p>Step 5 address-family [ipv4 ipv6]</p> <p>Example:</p> <pre>Router(config-vrf)# address-family ipv6</pre>	<p>Enters VRF address family configuration mode and specifies the IPv4 or IPv6 protocol.</p>
<p>Step 6 protection local-prefixes</p> <p>Example:</p> <pre>Router(config-vrf-af)# protection local-prefixes</pre>	<p>Allows a preconfigured backup path to carry traffic if the PE-CE link breaks by preserving the local prefixes while BGP reconverges.</p>

Command or Action	Purpose
Step 7 do show ip vrf detail Example: Router(config-vrf-af)# do show ip vrf detail	(Optional) Verifies that the MPLS VPN to BGP Local Convergence feature has been configured.

- [Examples, page 8](#)

Examples

To verify that local link protection has been enabled, enter the **show ip vrf detail** command. If the protection is enabled, the status message “Local prefix protection enabled” will be shown in the display:

```
Router# show ip vrf detail
VRF vpn1 (VRF Id = 1); default RD 100:1; default VPNID <not set>
Interfaces:
  AT1/0/1.1
VRF Table ID = 1
  Export VPN route-target communities
    RT:100:1
  Import VPN route-target communities
    RT:100:1          RT:100:2
No import route-map
No export route-map
VRF label distribution protocol: not configured
VRF label allocation mode: per-prefix
  Local prefix protection enabled
```

Troubleshooting Tips

- Ensure that a minimum of two paths are present for the protected prefix in BGP in steady state condition on the PE. The path using the protected PE should be the BGP best-path before failover occurs. To display the configuration, enter the **show ip bgp vpnv4 vrf vpn ip-prefix** command.
- Ensure that local protection has been enabled in the protected PE by entering the **show ip vrf detail** command, as shown in the [Examples, page 8](#).
- When route reflectors exist in the topology, ensure that each VRF has a unique route distinguisher.

Configuration Examples for MPLS VPN--BGP Local Convergence

- [Example MPLS VPN--BGP Local Convergence, page 9](#)
- [Example MPLS VPN--BGP Local Convergence for 6VPE 6PE, page 11](#)

Example MPLS VPN--BGP Local Convergence

The following examples show how MPLS VPN--BGP local convergence can prevent traffic loss after a link failure. You can display a detailed view of local link protection before, during, and after BGP convergence by using the **show bgp vpnv4** and **show mpls forwarding-table vrf** commands as shown in the following three-stage example.



Note

The **show bgp vpnv4 unicast** command is equivalent to the **show ip bgp vpnv4** command.

Example 1: Before the Link Failure

Both a primary path and a backup path have been configured:

```
Router# show bgp vpnv4 unicast all 172.16.0.1
BGP routing table entry for 100:1:172.16.0.1/32, version 2
Paths: (2 available, best #2, table v1)
Flag: 0x820
  Advertised to update-groups:
    1
  100, imported path from 100:2:172.16.0.1/32
    172.16.0.6 (metric 21) from 172.16.0.7 (172.16.0.7)
      Origin incomplete, metric 0, localpref 100, valid, internal
      Extended Community: RT:100:0
      Originator: 172.16.0.6, Cluster list: 172.16.0.7
      mpls labels in/out 16/17
  100
    172.16.1.1 from 172.16.1.1 (172.16.0.1)
      Origin incomplete, metric 0, localpref 100, valid, external, best
      Extended Community: RT:100:0
      mpls labels in/out 16/nolabel
BGP routing table entry for 100:2:172.16.0.1/32, version 9
Paths: (1 available, best #1, no table)
Flag: 0x820
  Not advertised to any peer
  100
    172.16.0.6 (metric 21) from 172.16.0.7 (172.16.0.7)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:100:0
      Originator: 172.16.0.6, Cluster list: 172.16.0.7
      mpls labels in/out nolabel/17
```

Label information for both paths can be displayed:

```
Router# show bgp vpnv4 unicast all labels
Network          Next Hop      In label/Out label
Route Distinguisher: 100:1 (v1)
  172.16.0.1/32   172.16.0.6    16/17
                   172.16.1.1    16/nolabel
  172.16.0.5/32   172.16.0.4    nolabel/23
  172.16.0.22/32  0.0.0.0       17/nolabel(v1)
  172.16.0.44/32  172.16.0.4    nolabel/24
  172.16.0.66/32  172.16.0.6    nolabel/21
  172.16.1.0/24   172.16.1.1    18/nolabel
                   0.0.0.0       18/nolabel(v1)
  172.16.5.0/24   172.16.0.4    nolabel/25
  172.16.8.0/24   172.16.0.6    19/23
                   172.16.1.1    19/nolabel
Route Distinguisher: 100:2
  172.16.0.1/32   172.16.0.6    nolabel/17
  172.16.0.66/32  172.16.0.6    nolabel/21
  172.16.8.0/24   172.16.0.6    nolabel/23
```

The PE1 (see the first figure above) forwarding table contains BGP best-path information:

```
Router# show mpls forwarding-table vrf v1 172.16.0.1 detail
Local      Outgoing Prefix          Bytes Label  Outgoing  Next Hop
Label      Label    or Tunnel Id   Switched     interface
16         No Label 172.16.0.1/32[V] 570         Et0/0     172.16.1.1
          MAC/Encaps=14/14, MRU=1504, Label Stack{
          AABCC000B00AABCC000C000800
          VPN route: v1
          No output feature configured
```

Example 2: After the Link Failure and Before BGP Convergence

After the link failure on only one path, the backup path remains available (see the second figure above):

```
Router# show bgp vpnv4 unicast all 172.16.0.1

BGP routing table entry for 100:1:172.16.0.1/32, version 19
Paths: (1 available, best #1, table v1)
  Not advertised to any peer
  100, imported path from 100:2:172.16.0.1/32
    172.16.0.6 (metric 21) from 172.16.0.7 (172.16.0.7)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:100:0
      Originator: 172.16.0.6, Cluster list: 172.16.0.7
      mpls labels in/out 16/17
BGP routing table entry for 100:2:172.16.0.1/32, version 9
Paths: (1 available, best #1, no table)
  Not advertised to any peer
  100
    172.16.0.6 (metric 21) from 172.16.0.7 (172.16.0.7)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:100:0
      Originator: 172.16.0.6, Cluster list: 172.16.0.7
      mpls labels in/out nolabel/17
```

The label information for the backup path label can be displayed:

```
Router# show bgp vpnv4 unicast all labels

Network      Next Hop      In label/Out label
Route Distinguisher: 100:1 (v1)
 172.16.0.1/32 172.16.0.6    16/17
 172.16.0.5/32 172.16.0.4    nolabel/23
 172.16.0.22/32 0.0.0.0       17/nolabel(v1)
 172.16.0.44/32 172.16.0.4    nolabel/24
 172.16.0.66/32 172.16.0.6    nolabel/21
 172.16.1.0/24 172.16.0.6    nolabel/22
 172.16.5.0/24 172.16.0.4    nolabel/25
 172.16.8.0/24 172.16.0.6    19/23
Route Distinguisher: 100:2
 172.16.0.1/32 172.16.0.6    nolabel/17
 172.16.0.66/32 172.16.0.6    nolabel/21
 172.16.1.0/24 172.16.0.6    nolabel/22
 172.16.8.0/24 172.16.0.6    nolabel/23
```

The PE 1 (see the first figure above) forwarding table contains new label and next-hop information to direct traffic onto the backup path:

```
Router# show mpls forwarding-table vrf v1 172.16.0.1 detail

Local      Outgoing Prefix          Bytes Label  Outgoing  Next Hop
Label      Label    or Tunnel Id   Switched     interface
16         17       172.16.0.1/32[V] 0            Et1/0     172.16.3.2
          MAC/Encaps=14/22, MRU=1496, Label Stack{21 17}
          AABCC000D00AABCC000C018847 0001500000011000
          VPN route: v1
          No output feature configured
```

Example 3: After Local Label Expiration and BGP Reconvergence

Because the local label preservation window has expired, the replacement local label is now gone from the PE 1 forwarding table information:

```
Router# show mpls forwarding-table vrf v1 172.16.0.1 detail
Local      Outgoing  Prefix          Bytes Label  Outgoing  Next Hop
Label      Label     or Tunnel Id   Switched     interface
None       17        172.16.0.1/32[V] 0             Et1/0      172.16.3.2
          MAC/Encaps=14/22, MRU=1496, Label Stack{21 17}
          AABCC000D00AABCC000C018847 0001500000011000
          VPN route: v1
          No output feature configured
```

The new BGP information reverts to the configuration shown in the first figure above:


```
Router# show bgp vpnv4 unicast all 172.16.0.1
BGP routing table entry for 100:1:172.16.0.1/32, version 23
Paths: (1 available, best #1, table v1)
  Not advertised to any peer
  100, imported path from 100:2:172.16.0.1/32
    172.16.0.6 (metric 21) from 172.16.0.7 (172.16.0.7)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:100:0
      Originator: 172.16.0.6, Cluster list: 172.16.0.7
      mpls labels in/out nolabel/17
BGP routing table entry for 100:2:172.16.0.1/32, version 9
Paths: (1 available, best #1, no table)
  Not advertised to any peer
  100
    172.16.0.6 (metric 21) from 172.16.0.7 (172.16.0.7)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:100:0
      Originator: 172.16.0.6, Cluster list: 172.16.0.7
      mpls labels in/out nolabel/17
Router# show bgp vpnv4 unicast all labels

Network          Next Hop          In label/Out label
Route Distinguisher: 100:1 (v1)
 172.16.0.1/32    172.16.0.6        nolabel/17
 172.16.0.5/32    172.16.0.4        nolabel/23
 172.16.0.22/32   0.0.0.0           17/nolabel(v1)
 172.16.0.44/32   172.16.0.4        nolabel/24
 172.16.0.66/32   172.16.0.6        nolabel/21
 172.16.1.0/24    172.16.0.6        nolabel/22
 172.16.5.0/24    172.16.0.4        nolabel/25
 172.16.8.0/24    172.16.0.6        nolabel/23
Route Distinguisher: 100:2
 172.16.0.1/32    172.16.0.6        nolabel/17
 172.16.0.66/32   172.16.0.6        nolabel/21
 172.16.1.0/24    172.16.0.6        nolabel/22
 172.16.8.0/24    172.16.0.6        nolabel/23
```

Example MPLS VPN--BGP Local Convergence for 6VPE 6PE

You can display a detailed view of local link protection before, during, and after BGP local convergence for Cisco IOS VPN IPv6 provider edge routers (6VPE) and Cisco IOS IPv6 provider edge routers (6PE) over MPLS by using the **show bgp vpnv6** and **show mpls forwarding-table vrf** commands as shown in the following three-stage example.

The figure below shows an MPLS VPN with BGP local convergence configured. The PE to CE routing protocol is eBGP, and the PE to route reflector (RR) sessions are BGP VPNv6. The protected prefix is the CE 1 loopback (2001:0DB8::/128). The primary path is from PE 1 to CE 1. The secondary path is from PE 1, through P and PE3, to CE 1.



Example 1: Before the Link Failure

Both a primary path and a backup path have been configured for the prefix 2001:0DB8::/128. The inlabel/outlabel settings for the two paths are 28/28 and 28/nolabel.

```
Router# show bgp vpnv6 unicast all 2001:0DB8::/128
BGP routing table entry for [1:1]2001:0DB8::/128, version 5
Paths: (2 available, best #2, table v1)
  Advertised to update-groups:
    2
  100, imported path from [2:2]2001:0DB8::/128
    ::FFFF:10.6.6.6 (metric 21) from 10.7.7.7 (10.7.7.7)
    Origin incomplete, metric 0, localpref 100, valid, internal
    Extended Community: RT:1:1
    Originator: 10.6.6.6, Cluster list: 10.7.7.7
    mpls labels in/out 28/28
  100
  2001:0DB8:0:ABCD::1 (FE80::A8BB:CCFF:FE00:B00) from 2001:0DB8:0:ABCD::1 (10.1.1.1)
    Origin incomplete, metric 0, localpref 100, valid, external, best

    Extended Community: RT:1:1
    mpls labels in/out 28/nolabel
BGP routing table entry for [2:2]2001:0DB8::/128, version 11
Paths: (1 available, best #1, no table)
  Not advertised to any peer
  100
  ::FFFF:10.6.6.6 (metric 21) from 10.7.7.7 (10.7.7.7)
  Origin incomplete, metric 0, localpref 100, valid, internal, best
  Extended Community: RT:1:1
  Originator: 10.6.6.6, Cluster list: 10.7.7.7
  mpls labels in/out nolabel/28
```

The PE 1 forwarding table contains new label and next-hop information to direct traffic onto the backup path:

```
Router#
show mpls forwarding-table vrf v1 2001:0DB8::/128 detail
Local      Outgoing Prefix          Bytes Label  Outgoing Next Hop
Label      Label    or Tunnel Id   Switched    interface
28         No Label 2001:0DB8::/128[V] 804         Et0/0
FE80::A8BB:CCFF:FE00:B00
  MAC/Encaps=14/14, MRU=1504, Label Stack{}
  AABBC000B00AABCC000C0086DD
  VPN route: v1
  No output feature configured
```

Example 2: After the Link Failure

After the link failure, the backup path is still available, the original path is removed from BGP, and the backup path is activated:

```
Router# show mpls forwarding-table vrf v1 2001:0DB8::/128 detail
Local      Outgoing Prefix          Bytes Label  Outgoing Next Hop
Label      Label    or Tunnel Id   Switched    interface
28         28       2001:0DB8::/128[V] 0           Et1/0      10.3.0.2
  MAC/Encaps=14/22, MRU=1496, Label Stack{23 28}
  AABBC000D00AABCC000C018847 000170000001C000
  VPN route: v1
  No output feature configured
```

After a configured length of time, the local label expires. The output from the **show mpls forwarding-table** command also verifies that the local label has expired:

```
Router# show mpls forwarding-table vrf v1 2001:0DB8::/128 detail
Local      Outgoing Prefix          Bytes Label  Outgoing Next Hop
Label      Label    or Tunnel Id   Switched    interface
None       28       2001:0DB8::/128[V] 0           Et1/0      10.3.0.2
  MAC/Encaps=14/22, MRU=1496, Label Stack{23 28}
```

```

AABBCC000D00AABBCC000C018847 000170000001C000
VPN route: v1
No output feature configured

```

Example 3: After the Link Is Restored

When the link is restored the original path is added to BGP and the traffic switches back to this path:

```

Router# show bgp vpnv6 unicast all 2001:0DB8::/128
BGP routing table entry for [1:1]2001:0DB8::/128, version 28
Paths: (2 available, best #1, table v1)
  Advertised to update-groups:
    2
  100
    2001:0DB8:0:ABCD::1 (FE80::A8BB:CCFF:FE00:B00) from 2001:0DB8:0:ABCD::1 (10.1.1.1)
      Origin incomplete, metric 0, localpref 100, valid, external, best
      Extended Community: RT:1:1
      mpls labels in/out 16/nolabel
  100, imported path from [2:2]2001:0DB8::/128
    ::FFFF:10.6.6.6 (metric 21) from 10.7.7.7 (10.7.7.7)
      Origin incomplete, metric 0, localpref 100, valid, internal
      Extended Community: RT:1:1
      Originator: 10.6.6.6, Cluster list: 10.7.7.7
      mpls labels in/out 16/28
BGP routing table entry for [2:2]2001:0DB8::/128, version 11
Paths: (1 available, best #1, no table)
  Not advertised to any peer
  100
    ::FFFF:10.6.6.6 (metric 21) from 10.7.7.7 (10.7.7.7)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:1:1
      Originator: 10.6.6.6, Cluster list: 10.7.7.7
      mpls labels in/out nolabel/28
Router# show mpls for vrf v1 2001:0DB8::/128 detail
Local      Outgoing Prefix          Bytes Label  Outgoing  Next Hop
Label      Label    or Tunnel Id    Switched     interface
16         No Label 2001:0DB8::/128[V] 0          Et0/0      FE80::A8BB:CCFF:FE00:B00
          MAC/Encaps=14/14, MRU=1504, Label Stack{}
          AABBCC000B00AABBCC000C0086DD
          VPN route: v1
          No output feature configured

```

Additional References

Related Documents

Related Topic	Document Title
Cisco IOS commands	Cisco IOS Master Commands List, All Releases
BGP configuration	Configuring a Basic BGP Network
Protocol for quickly detecting failed forwarding paths	Bidirectional Forwarding Detection
Configuration of BGP PIC Edge for IP and MPLS-VPN	BGP PIC Edge for IP and MPLS-VPN
Configuration of internal BGP features	Configuring Internal BGP Features
Configuration of VRF under the specific cases of IPv4 and IPv6 situations	MPLS VPN--VRF CLI for IPv4 and IPv6 VPNs

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	--

MIBs

MIB	MIBs Link
No new or modified MIBs are supported by this feature, and support for existing MIBs has not been modified by this feature.	To locate and download MIBs for selected platforms, Cisco software releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

RFCs

RFC	Title
RFC 2547	BGP/MPLS VPNs

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 MPLS VPN--BGP Local Convergence

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 MPLS VPN--BGP Local Convergence**

Feature Name	Releases	Feature Information
MPLS VPN--BGP Local Convergence	12.2(33)SRC 12.2(33)SB 15.0(1)M	<p>This feature reduces the downtime of a PE-CE link failure by rerouting PE-egress traffic onto a backup path to the CE before BGP has reconverged.</p> <p>In 12.2(33)SRC, this feature was introduced on the Cisco 7200 and the Cisco 7600.</p> <p>In 12.2(33)SB, this feature became available on the Cisco 7300 series and the Cisco 10000 series routers.</p> <p>This feature was integrated into Cisco IOS Release 15.0(1)M.</p> <p>The following command was introduced: protection local-prefixes.</p>
MPLS VPN--BGP Local Convergence for 6VPE/6PE	15.0(1)S	<p>This feature implements MPLS VPN--BGP local convergence for Cisco IOS 6VPE and Cisco IOS 6PE over MPLS.</p> <p>The following command was modified: protection local-prefixes.</p>

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

© 2011 Cisco Systems, Inc. All rights reserved.