



BGP PIC Edge for IP and MPLS-VPN

First Published: November 20, 2009

Last Updated: March 31, 2011

The BGP PIC Edge for IP and MPLS-VPN feature improves BGP convergence after a network failure. This convergence is applicable to both core and edge failures and can be used in both IP and MPLS networks. The BGP PIC Edge for IP and MPLS-VPN feature creates and stores a backup/alternate path in the routing information base (RIB), forwarding information base (FIB), and Cisco Express Forwarding so that when a failure is detected, the backup/alternate path can immediately take over, thus enabling fast failover.



Note

In this document, the BGP PIC Edge for IP and MPLS-VPN feature is called BGP PIC.

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 for BGP PIC”](#) section on page 19.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to <http://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

Contents

- [Prerequisites for BGP PIC, page 2](#)
- [Restrictions for BGP PIC, page 2](#)
- [Information About BGP PIC, page 3](#)
- [How to Configure BGP PIC, page 12](#)
- [Configuration Examples for BGP PIC, page 14](#)
- [Additional References, page 17](#)
- [Feature Information for BGP PIC, page 19](#)

Prerequisites for BGP PIC

- Ensure that the Border Gateway Protocol (BGP) and the IP or Multiprotocol Label Switching (MPLS) network is up and running with the customer site connected to the provider site by more than one path (multihomed).
- Ensure that the backup/alternate path has a unique next hop that is not the same as the next hop of the best path.
- Enable the Bidirectional Forwarding Detection (BFD) protocol to quickly detect link failures of directly connected neighbors.

Restrictions for BGP PIC

- With BGP Multipath, the BGP Prefix-Independent Convergence (PIC) feature is already supported.
- In MPLS VPNs, the BGP PIC feature is not supported with MPLS VPN Inter-Autonomous Systems Option B.
- The BGP PIC feature supports prefixes only for IPv4, IPv6, VPNv4, and VPNv6 address families.
- The BGP PIC feature cannot be configured with Multicast or L2VPN Virtual Routing and Forwarding (VRF) address families.
- If the route reflector is only in the control plane, then you do not need BGP PIC, because BGP PIC addresses data plane convergence.
- When two PE routers become each other's backup/alternate path to a CE router, traffic might loop if the CE router fails. Neither router will reach the CE router, and traffic will continue to be forwarded between the PE routers until the time-to-live (TTL) timer expires.
- The BGP PIC feature does not support Nonstop Forwarding with Stateful Switchover (NSF/SSO). However, ISSU is supported if both Route Processors have the BGP PIC feature configured.
- The BGP PIC feature solves the traffic forwarding only for a single network failure at both the edge and the core.
- The BGP PIC feature does not work with the BGP Best External feature. If you try to configure the BGP PIC feature after configuring the BGP Best External feature, you receive an error.

Information About BGP PIC

- [Benefits of the BGP PIC Edge for IP and MPLS-VPN Feature, page 3](#)
- [How BGP Converges Under Normal Circumstances, page 3](#)
- [How BGP PIC Improves Convergence, page 4](#)
- [How a Failure Is Detected, page 6](#)
- [How BGP PIC Achieves Subsecond Convergence, page 6](#)
- [How BGP PIC Improves Upon the Functionality of MPLS VPN—BGP Local Convergence, page 6](#)
- [Configuration Modes for Enabling BGP PIC, page 6](#)
- [BGP PIC Scenarios, page 7](#)

Benefits of the BGP PIC Edge for IP and MPLS-VPN Feature

- An additional path for failover allows faster restoration of connectivity if a primary path is invalid or withdrawn.
- Reduction of traffic loss.
- Constant convergence time so that the switching time is the same for all prefixes.

How BGP Converges Under Normal Circumstances

Under normal circumstances, BGP can take several seconds to a few minutes to converge after a network change. At a high level, BGP goes through the following process:

1. BGP learns of failures through either Interior Gateway Protocol (IGP) or BFD events or interface events.
2. BGP withdraws the routes from the routing information base (RIB), and the RIB withdraws the routes from the forwarding information base (FIB) and distributed FIB (dFIB). This process clears the data path for the affected prefixes.
3. BGP sends withdraw messages to its neighbors.
4. BGP calculates the next best path to the affected prefixes.
5. BGP inserts the next best path for affected prefixes into the RIB, and the RIB installs them in the FIB and dFIB.

This process takes a few seconds or a few minutes to complete, depending on the latency of the network, the convergence time across the network, and the local load on the devices. The data plane converges only after the control plane converges.

How BGP PIC Improves Convergence

The BGP PIC functionality is achieved by an additional functionality in the BGP, RIB, Cisco Express Forwarding, and MPLS.

- **BGP Functionality**

BGP PIC affects prefixes under IPv4 and VPNv4 address families. For those prefixes, BGP calculates an additional second best path, along with the primary best path. (The second best path is called the backup/alternate path.) BGP installs the best and backup/alternate paths for the affected prefixes into the BGP RIB. The backup/alternate path provides a fast reroute mechanism to counter a singular network failure. BGP also includes the alternate/backup path in its application programming interface (API) to the IP RIB.

- **RIB Functionality**

For BGP PIC, RIB installs an alternate path per route if one is available. With the BGP PIC functionality, if the RIB selects a BGP route containing a backup/alternate path, it installs the backup/alternate path with the best path. The RIB also includes the alternate path in its API with the FIB.

- **Cisco Express Forwarding Functionality**

With BGP PIC, Cisco Express Forwarding stores an alternate path per prefix. When the primary path goes down, Cisco Express Forwarding searches for the backup/alternate path in a prefix independent manner. Cisco Express Forwarding also listens to BFD events to rapidly detect local failures.

- **MPLS Functionality**

MPLS Forwarding is similar to Cisco Express Forwarding, in that it stores alternate paths and switches to an the alternate path if the primary path goes down.

When the BGP PIC feature is enabled, BGP calculates a backup/alternate path per prefix and installs it into BGP RIB, IP RIB, and FIB. This improves convergence after a network failure. There are two types of network failures that the BGP PIC feature detects:

- **Core node/link failure (internal Border Gateway Protocol [iBGP] node failure):** If a PE node/link fails, then the failure is detected through IGP convergence. IGP conveys the failure through the RIB to the FIB.
- **Local link/immediate neighbor node failure (external Border Gateway Protocol [eBGP] node/link failure):** To detect a local link failure or eBGP single-hop peer node failure in less than a second, you must enable BFD. Cisco Express Forwarding looks for BFD events to detect a failure of an eBGP single-hop peer.

Convergence in the Data Plane

Upon detection of a failure, Cisco Express Forwarding detects the alternate next hop for all prefixes affected by the failure. The data plane convergence is achieved in subseconds depending on whether the BGP PIC implementation exists in the software or hardware.

Convergence in the Control Plane

Upon detection of failure, BGP learns about the failure through IGP convergence or BFD events and sends withdraw messages for the prefixes, recalculating the best and backup/alternate paths, and advertising the next best path across the network.

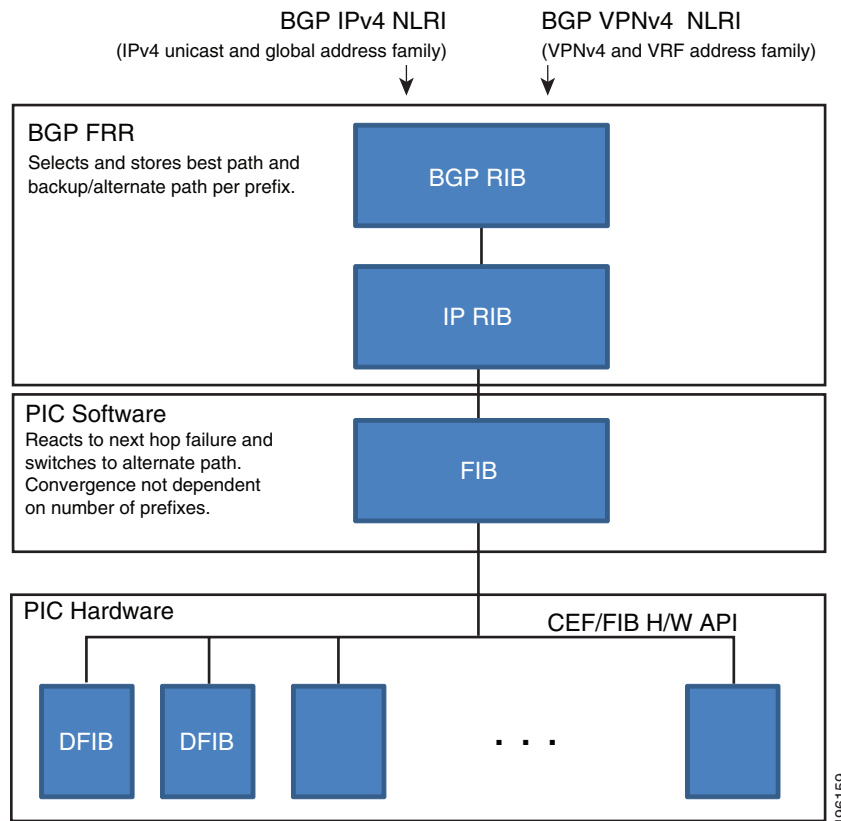
BGP Fast Reroute's Role in the BGP PIC Feature

BGP Fast Reroute (FRR) provides a best path and a backup/alternate path in BGP, RIB, and Cisco Express Forwarding. BGP FRR provides a very fast reroute mechanism into the RIB and Cisco Express Forwarding on the backup BGP next hop to reach a destination when the current best path is not available.

BGP FRR precomputes a second best path in BGP and gives it to the RIB and Cisco Express Forwarding as a backup/alternate path, and Cisco Express Forwarding programs it into line cards.

Therefore, BGP FRR sets up the best path and backup/alternate path. The BGP PIC feature provides the ability for Cisco Express Forwarding to quickly switch the traffic to the other egress ports if the current next hop or the link to this next hop goes down. This is illustrated in [Figure 1](#).

Figure 1 BGP PIC Edge and BGP FRR



How a Failure Is Detected

A failure in the iBGP (remote) peer is detected by IGP; it may take a few seconds to detect the failure. Convergence can occur in subseconds or seconds, depending on whether PIC is enabled on the line cards.

If the failure is with directly connected neighbors (eBGP), and if you use BFD to detect when a neighbor has gone down, the detection happens within a subsecond and the convergence can occur in subseconds or seconds, depending on whether PIC is enabled on the line cards.

How BGP PIC Achieves Subsecond Convergence

The BGP PIC feature works at the Cisco Express Forwarding level, and Cisco Express Forwarding can be processed in both hardware line cards and in the software.

- For platforms that support Cisco Express Forwarding processing in the line cards, the BGP PIC feature can converge in subseconds. The Cisco 7600 router and Cisco 10000 router supports Cisco Express Forwarding processing in the line cards and in the software, and thus can attain subsecond convergence.
- For platforms that do not use Cisco Express Forwarding in hardware line cards, Cisco Express Forwarding is achieved in the software. The BGP PIC feature will work with the Cisco Express Forwarding through the software and achieve convergence within seconds. The Cisco 7200 router supports Cisco Express Forwarding in the software and thus can achieve convergence in seconds rather than milliseconds.

How BGP PIC Improves Upon the Functionality of MPLS VPN—BGP Local Convergence

The BGP PIC feature is an enhancement to the [MPLS VPN—BGP Local Convergence](#) feature, which provides a failover mechanism that recalculates the best path and installs the new path in forwarding after a link failure. The feature maintains the local label for 5 minutes to ensure that the traffic uses the backup/alternate path, thus minimizing traffic loss.

The BGP PIC feature improves the LoC time to under a second by calculating a backup/alternate path in advance. When a link failure occurs, the traffic is sent to the backup/alternate path.

When you configure the BGP PIC feature, it will override the functionality of the [MPLS VPN—BGP Local Convergence](#) feature. You do not have to remove the **protection local-prefixes** command from the configuration.

Configuration Modes for Enabling BGP PIC

Because many service provider networks contain many VRFs, the BGP PIC feature allows you to configure the BGP PIC feature for all VRFs at once.

- VPNv4 address family configuration mode protects all the VRFs.
- VRF-IPv4 address family configuration mode protects only IPv4 VRFs.
- Router configuration mode protects prefixes in the global routing table.

BGP PIC Scenarios

The following scenarios explain how you can configure the BGP PIC functionality to achieve fast convergence:

- [IP PE-CE Link and Node Protection on the CE Side \(Dual PEs\), page 7](#)
- [IP PE-CE Link and Node Protection on the CE Side \(Dual CEs and Dual PE Primary and Backup Nodes\), page 8](#)
- [IP MPLS PE-CE Link Protection for the Primary or Backup/Alternate Path, page 9](#)
- [IP MPLS PE-CE Node Protection for Primary or Backup/Alternate Path, page 10](#)

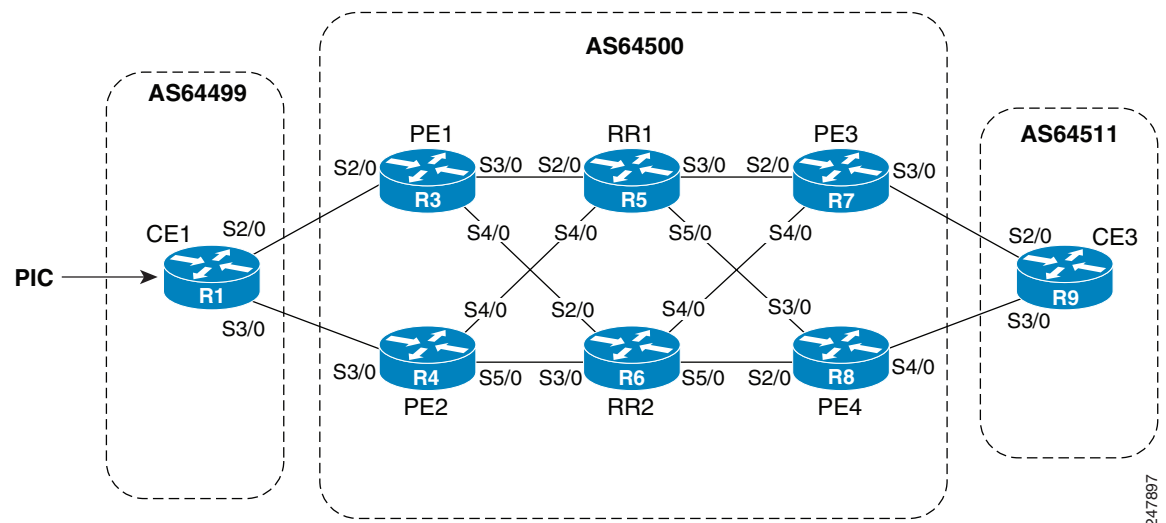
IP PE-CE Link and Node Protection on the CE Side (Dual PEs)

Figure 2 shows a network that uses the BGP PIC feature. The network includes the following components:

- eBGP sessions exist between the PE and CE routers.
- Traffic from CE1 uses PE1 to reach network 192.168.9.0/24 through router CE3.
- CE1 has two paths:
 - PE1 as the primary path.
 - PE2 as the backup/alternate path.

CE1 is configured with the BGP PIC feature. BGP computes PE1 as the best path and PE2 as the backup/alternate path and installs both routes into the RIB and Cisco Express Forwarding plane. When the CE1-PE1 link goes down, Cisco Express Forwarding detects the link failure and points the forwarding object to the backup/alternate path. Traffic is quickly rerouted due to local fast convergence in Cisco Express Forwarding.

Figure 2 Using BGP PIC to Protect the PE-CE Link



IP PE–CE Link and Node Protection on the CE Side (Dual CEs and Dual PE Primary and Backup Nodes)

Figure 3 shows a network that uses the BGP PIC feature on CE1. The network includes the following components:

- eBGP sessions exist between the PE and CE routers.
- Traffic from CE1 uses PE1 to reach network 192.168.9.0/24 through router CE3.
- CE1 has two paths:
 - PE1 as the primary path.
 - PE2 as the backup/alternate path.
- An iBGP session exists between the CE1 and CE2 routers.

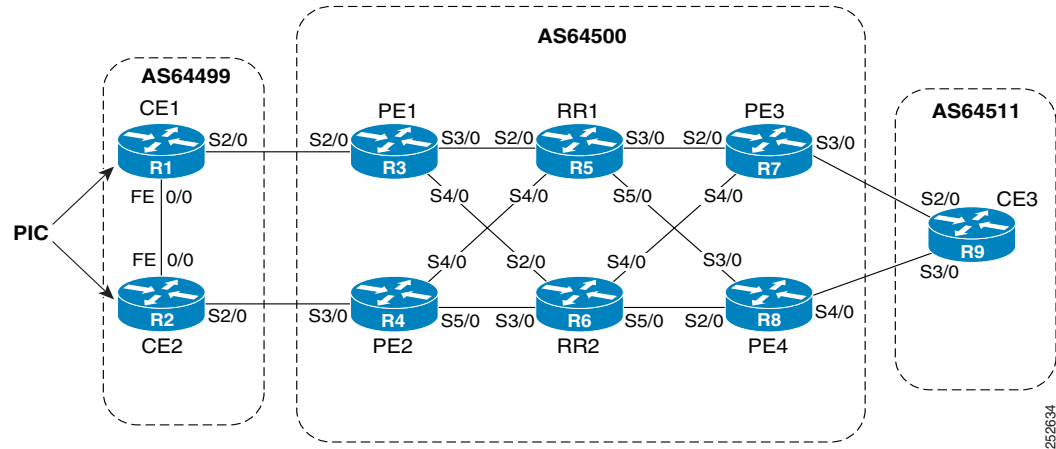
In this example, CE1 and CE2 are configured with the BGP PIC feature. BGP computes PE1 as the best path and PE2 as the backup/alternate path and installs both the routes into the RIB and Cisco Express Forwarding plane.

There should not be any policies set on CE1 and CE2 for the eBGP peers PE1 and PE2. Both CE routers must point to the eBGP route as next hop. On CE1, the next hop to reach CE3 is through PE1, so PE1 is the best path to reach CE3. On CE2, the best path to reach CE3 is PE2. CE2 advertises itself as the next hop to CE1, and CE1 does the same to CE2. As a result, CE1 has two paths for the specific prefix and it usually selects the directly connected eBGP path over the iBGP path according to the best path selection rules. Similarly, CE2 has two paths—an eBGP path through PE2 and an iBGP path through CE1–PE1.

When the CE1–PE1 link goes down, Cisco Express Forwarding detects the link failure and points the forwarding object to the backup/alternate node CE2. Traffic is quickly rerouted due to local fast convergence in Cisco Express Forwarding.

If the CE1–PE1 link or PE1 goes down and BGP PIC is enabled on CE1, BGP recomputes the best path, removing the next hop PE1 from RIB and reinstalling CE2 as the next hop into the RIB and Cisco Express Forwarding. CE1 automatically gets a backup/alternate repair path into Cisco Express Forwarding and the traffic loss during forwarding is now in subseconds, thereby achieving fast convergence.

Figure 3 Using BGP PIC in a Dual CE, Dual PE Network



IP MPLS PE–CE Link Protection for the Primary or Backup/Alternate Path

Figure 3 shows a network that uses the BGP PIC feature on CE1 and CE2. The network includes the following components:

- eBGP sessions exist between the PE and CE routers.
- The PE routers are VPNv4 iBGP peers with reflect routers in the MPLS network.
- Traffic from CE1 uses PE1 to reach the network 192.168.9.0/24 through router CE3.
- CE3 is dual-homed with PE3 and PE4.
- PE1 has two paths to reach CE3 from the reflect routers:
 - PE3 is the primary path with the next hop as a PE3 address.
 - PE4 is the backup/alternate path with the next hop as a PE4 address.

In this example, all the PE routers can be configured with the BGP PIC feature under IPv4 or VPNv4 address families.

For BGP PIC to work in BGP for PE–CE link protection, set the policies on PE3 and PE4 for prefixes received from CE3 so that one of the PE routers acts as the primary and the other as the backup/alternate. Usually, this is done using local preference and giving better local preference to PE3. In the MPLS cloud, traffic internally flows through PE3 to reach CE3. Thus, PE1 has PE3 as the best path and PE4 as the second path.

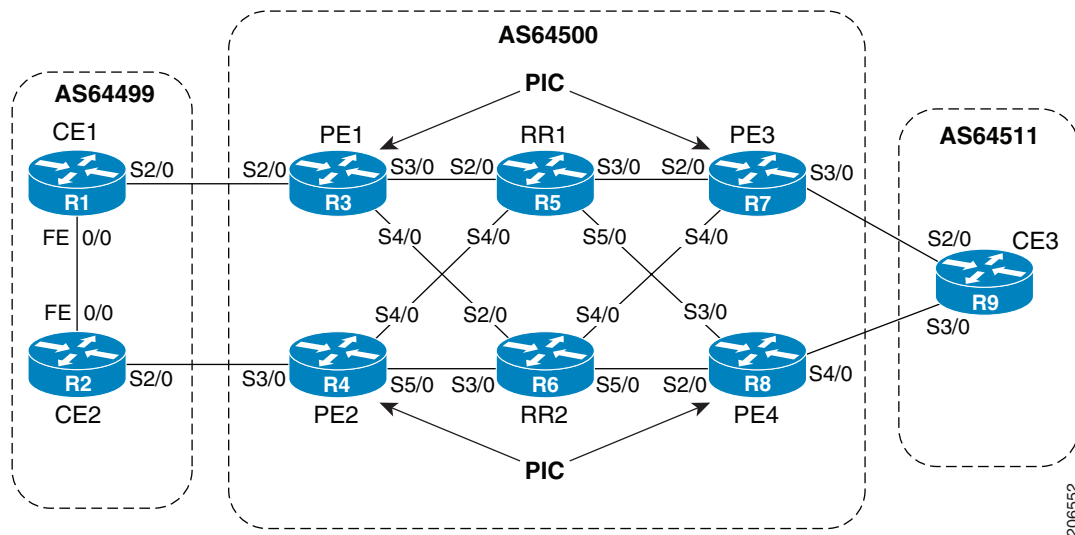
When the PE3–CE3 link goes down, Cisco Express Forwarding detects the link failure, and PE3 recomputes the best path, selects PE4 as the best path, and sends a withdraw message for the PE3 prefix to the reflect routers. Some of the traffic goes through PE3–PE4 until BGP installs PE4 as the best path route into the RIB and Cisco Express Forwarding. PE1 receives the withdraw, recomputes the best path, selects PE4 as the best path, and installs the routes into the RIB and Cisco Express Forwarding plane.

Thus, with BGP PIC enabled on PE3 and PE4, Cisco Express Forwarding detects the link failure and does in-place modification of the forwarding object to the backup/alternate node PE4 that already exists in Cisco Express Forwarding. PE4 knows that the backup/alternate path is locally generated and routes the traffic to the egress port connected to CE3. This way, traffic loss is minimized and fast convergence is achieved.

IP MPLS PE–CE Node Protection for Primary or Backup/Alternate Path

Figure 4 shows a network that uses the BGP PIC feature on all the PE routers in an MPLS network.

Figure 4 Enabling BGP PIC on All PEs Routers in the MPLS Network



The network includes the following components:

- eBGP sessions exist between the PE and CE routers.
- The PE routers are VPNv4 iBGP peers with reflect routers in the MPLS network.
- Traffic from CE1 uses PE1 to reach the network 192.168.9.0/24 through router CE3.
- CE3 is dual-homed with PE3 and PE4.
- PE1 has two paths to reach CE3 from the reflect routers:
 - PE3 is the primary path with the next hop as a PE3 address.
 - PE4 is the backup/alternate path with the next hop as a PE4 address.

In this example, all the PE routers are configured with the BGP PIC feature under IPv4 and VPNv4 address families.

For BGP PIC to work in BGP for the PE–CE node protection, set the policies on PE3 and PE4 for the prefixes received from CE3 such that one of the PE routers acts as primary and the other as backup/alternate. Usually, this is done using local preference and giving better local preference to PE3. In the MPLS cloud, traffic internally flows through PE3 to reach CE3. So, PE1 has PE3 as the best path and PE4 as the second path.

When PE3 goes down, PE1 knows about the removal of the host prefix by IGP in subseconds, recomputes the best path, selects PE4 as the best path, and installs the routes into the RIB and Cisco Express Forwarding plane. Normal BGP convergence will happen while BGP PIC is redirecting the traffic through PE4, and packets are not lost.

Thus, with BGP PIC enabled on PE3, Cisco Express Forwarding detects the node failure on PE3 and points the forwarding object to the backup/alternate node PE4. PE4 knows that the backup/alternate path is locally generated and routes the traffic to the egress port using the backup/alternate path. This way, traffic loss is minimized.

No Local Policies Set on the PE Routers

PE1 and PE2 point to the eBGP CE paths as the next hop with no local policy. Each of the PE routers receives the other's path, and BGP calculates the backup/alternate path and installs it into Cisco Express Forwarding, along with its own eBGP path towards CE as the best path. The limitation of the MPLS PE-CE link and node protection solutions is that you cannot change BGP policies. They should work without the need for a best-external path.

Local Policies Set on the PE Routers

Whenever there is a local policy on the PE routers to select one of the PE routers as the primary path to reach the egress CE, the **bgp advertise-best-external** command is needed on the backup/alternate node PE3 to propagate the external CE routes with a backup/alternate label into the route reflectors and the far-end PE routers.

Cisco Express Forwarding Recursion

Recursion is the ability to find the next longest matching path when the primary path goes down.

When the BGP PIC feature is not installed, and if the next hop to a prefix fails, Cisco Express Forwarding finds the next path to reach the prefix by recursing through the FIB to find the next longest matching path to the prefix. This is useful if the next hop is multiple hops away and there is more than one way of reaching the next hop.

However, with the BGP PIC feature, you may want to disable Cisco Express Forwarding recursion for the following reasons:

- Recursion slows down convergence when Cisco Express Forwarding searches all the FIB entries.
- BGP PIC Edge already precomputes an alternate path, thus eliminating the need for Cisco Express Forwarding recursion.

When the BGP PIC functionality is enabled, Cisco Express Forwarding recursion is disabled by default for two conditions:

- For next hops learned with a /32 network mask (host routes)
- For next hops that are directly connected

For all other cases, Cisco Express Forwarding recursion is enabled.

As part of the BGP PIC functionality, you can issue the **bgp recursion host** command to disable or enable Cisco Express Forwarding recursion for BGP host routes.

To disable or enable Cisco Express Forwarding recursion for BGP directly connected next hops, you can issue the **disable-connected-check** command.

How to Configure BGP PIC

- [Configuring BGP PIC, page 12](#)

Configuring BGP PIC

Because many service provider networks contain many VRFs, the BGP PIC feature allows you to configure the BGP PIC feature for all VRFs at once.

- VPNv4 address family configuration mode protects all the VRFs.
- VRF-IPv4 address family configuration mode protects only IPv4 VRFs.
- Router configuration mode protects prefixes in the global routing table.

For a full configuration example that includes configuring multiprotocol VRFs and shows output to verify that the feature is enabled, see the [“Example: Configuring BGP PIC” section on page 14](#).

Prerequisites

- If you are implementing the BGP PIC feature in an MPLS VPN, ensure that the network is working properly before configuring the BGP PIC feature. See [Configuring MPLS Layer 3 VPNs](#) for more information.
- If you are implementing the BGP PIC feature in an MPLS VPN, configure multiprotocol VRFs, which allow you to share route-target policies (import and export) between IPv4 and IPv6 or to configure separate route-target policies for IPv4 and IPv6 VPNs. For information about configuring multiprotocol VRFs, see [MPLS VPN—VRF CLI for IPv4 and IPv6 VPNs](#).
- Ensure that the CE router is connected to the network by at least two paths.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **router bgp** *autonomous-system-number*
4. **address-family ipv4** [**unicast** | **vrf** *vrf-name*]
or
address-family vpnv4 [**unicast**]
5. **bgp additional-paths install**
6. **neighbor** *ip-address* **remote-as** *autonomous-system-number*
7. **neighbor** *ip-address* **activate**
8. **bgp recursion host**
9. **neighbor ip-address fall-over** [**bfd** | **route-map** *map-name*]
10. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	<p>enable</p> <p>Example: Router> enable</p>	<p>Enables privileged EXEC mode.</p> <ul style="list-style-type: none"> Enter your password if prompted.
Step 2	<p>configure terminal</p> <p>Example: Router# configure terminal</p>	<p>Enters global configuration mode.</p>
Step 3	<p>router bgp <i>autonomous-system-number</i></p> <p>Example: Router(config)# router bgp 40000</p>	<p>Enters router configuration mode for the specified routing process.</p>
Step 4	<p>address-family ipv4 [unicast vrf <i>vrf-name</i>]</p> <p>or</p> <p>address-family vpnv4 [unicast]</p> <p>Example: Router(config-router)# address-family ipv4 unicast</p> <p>or</p> <p>Example: Router(config-router)# address-family vpnv4</p>	<p>Specifies the IPv4 or VPNv4 address family and enters address family configuration mode.</p> <ul style="list-style-type: none"> The unicast keyword specifies the IPv4 or VPNv4 unicast address family. The vrf keyword and <i>vrf-name</i> argument specify the name of the virtual routing and forwarding (VRF) instance to associate with subsequent IPv4 address family configuration mode commands.
Step 5	<p>bgp additional-paths install</p> <p>Example: Router(config-router-af)# bgp additional-paths install</p>	<p>Calculates a backup/alternate path and installs it into the RIB and Cisco Express Forwarding.</p>
Step 6	<p>neighbor <i>ip-address</i> remote-as <i>autonomous-system-number</i></p> <p>Example: Router(config-router-af)# neighbor 192.168.1.1 remote-as 45000</p>	<p>Adds the IP address of the neighbor in the specified autonomous system to the IPv4 multiprotocol BGP neighbor table of the local router.</p> <ul style="list-style-type: none"> By default, neighbors that are defined using the neighbor remote-as command in router configuration mode exchange only IPv4 unicast address prefixes. To exchange other address prefix types, neighbors must also be activated using the neighbor activate command in address family configuration mode for the other prefix types.
Step 7	<p>neighbor <i>ip-address</i> activate</p> <p>Example: Router(config-router-af)# neighbor 192.168.1.1 activate</p>	<p>Enables the neighbor to exchange prefixes for the IPv4 unicast address family with the local router.</p>

	Command or Action	Purpose
Step 8	bgp recursion host Example: Router(config-router-af)# bgp recursion host	(Optional) Enables the recursive-via-host flag for IPv4, VPNv4, and VRF address families. <ul style="list-style-type: none"> When the BGP PIC feature is enabled, Cisco Express Forwarding recursion is disabled. Under most circumstances, you do not want to enable recursion when BGP PIC is enabled.
Step 9	neighbor ip-address fall-over [bfd route-map map-name] Example: Router(config-router-af)# neighbor 192.168.1.1 fall-over bfd	Enables BFD protocol support to detect when a neighbor has gone away, which can occur within a subsecond.
Step 10	end Example: Router(config-router-af)# end	Exits address family configuration mode and returns to privileged EXEC mode.

Configuration Examples for BGP PIC

- [Example: Configuring BGP PIC, page 14](#)
- [Example: Displaying Backup/Alternate Paths for BGP PIC, page 15](#)

Example: Configuring BGP PIC

The following example shows how to configure the BGP PIC feature in VPNv4 address family configuration mode, which enables the feature on all VRFs. In the following example, there are two VRFs defined: blue and green. All the VRFs, including those in VRFs blue and green, are protected by backup/alternate paths.

```
vrf definition test1
 rd 400:1
  route-target export 100:1
  route-target export 200:1
  route-target export 300:1
  route-target export 400:1
  route-target import 100:1
  route-target import 200:1
  route-target import 300:1
  route-target import 400:1
  address-family ipv4
  exit-address-family
exit
!
interface Ethernet1/0
 vrf forwarding test1
 ip address 10.0.0.1 255.0.0.0
 exit
router bgp 3
 no synchronization
 bgp log-neighbor-changes
 redistribute static
```

```

redistribute connected
neighbor 10.6.6.6 remote-as 3
neighbor 10.6.6.6 update-source Loopback0
neighbor 10.7.7.7 remote-as 3
neighbor 10.7.7.7 update-source Loopback0
no auto-summary
!
address-family vpnv4
  bgp additional-paths install
  neighbor 10.6.6.6 activate
  neighbor 10.6.6.6 send-community both
  neighbor 10.7.7.7 activate
  neighbor 10.7.7.7 send-community both
exit-address-family
!
address-family ipv4 vrf blue
  import path selection all
  import path limit 10
  no synchronization
  neighbor 10.11.11.11 remote-as 1
  neighbor 10.11.11.11 activate
exit-address-family
!
address-family ipv4 vrf green
  import path selection all
  import path limit 10
  no synchronization
  neighbor 10.13.13.13 remote-as 1
  neighbor 10.13.13.13 activate
exit-address-family

```

The following **show vrf detail** command output shows that the BGP PIC feature is enabled:

```
Router# show vrf detail
```

```

VRF test1 (VRF Id = 1); default RD 400:1; default VPNID <not set>
  Interfaces:
    Se4/0
  Address family ipv4 (Table ID = 1 (0x1)):
    Export VPN route-target communities
      RT:100:1          RT:200:1          RT:300:1
      RT:400:1
    Import VPN route-target communities
      RT:100:1          RT:200:1          RT:300:1
      RT:400:1
    No import route-map
    No export route-map
    VRF label distribution protocol: not configured
    VRF label allocation mode: per-prefix
    Prefix protection with additional path enabled
  Address family ipv6 not active.

```

Example: Displaying Backup/Alternate Paths for BGP PIC

The command output in the following example shows that the VRFs in VRF blue have backup/alternate paths:

```
Router# show ip bgp vpnv4 vrf blue 10.0.0.0
```

```

BGP routing table entry for 10:12:12.0.0.0/24, version 88
Paths: (4 available, best #1, table blue)
  Additional-path

```

```

Advertised to update-groups:
  6
1, imported path from 12:23:12.0.0.0/24
  10.3.3.3 (metric 21) from 10.6.6.6 (10.6.6.6)
    Origin incomplete, metric 0, localpref 200, valid, internal, best
    Extended Community: RT:12:23
    Originator: 10.3.3.3, Cluster list: 10.0.0.1 , recursive-via-host
    mpls labels in/out nolabel/37
1, imported path from 12:23:12.0.0.0/24
  10.13.13.13 (via green) from 10.13.13.13 (10.0.0.2)
    Origin incomplete, metric 0, localpref 100, valid, external
    Extended Community: RT:12:23 , recursive-via-connected
1, imported path from 12:23:12.0.0.0/24
  10.3.3.3 (metric 21) from 10.7.7.7 (10.7.7.7)
    Origin incomplete, metric 0, localpref 200, valid, internal
    Extended Community: RT:12:23
    Originator: 10.3.3.3, Cluster list: 10.0.0.1 , recursive-via-host
    mpls labels in/out nolabel/37
1
  10.11.11.11 from 10.11.11.11 (1.0.0.1)
    Origin incomplete, metric 0, localpref 100, valid, external, backup/repair
    Extended Community: RT:11:12 , recursive-via-connected

```

The command output in the following example shows that the VRFs in VRF green have backup/alternate paths:

```

Router# show ip bgp vpnv4 vrf green 12.0.0.0

BGP routing table entry for 12:23:12.0.0.0/24, version 87
Paths: (4 available, best #4, table green)
Additional-path
Advertised to update-groups:
  5
1, imported path from 11:12:12.0.0.0/24
  10.11.11.11 (via blue) from 10.11.11.11 (1.0.0.1)
    Origin incomplete, metric 0, localpref 100, valid, external
    Extended Community: RT:11:12 , recursive-via-connected
1
  10.3.3.3 (metric 21) from 10.7.7.7 (10.7.7.7)
    Origin incomplete, metric 0, localpref 200, valid, internal
    Extended Community: RT:12:23
    Originator: 10.3.3.3, Cluster list: 10.0.0.1 , recursive-via-host
    mpls labels in/out nolabel/37
1
  10.13.13.13 from 10.13.13.13 (10.0.0.2)
    Origin incomplete, metric 0, localpref 100, valid, external, backup/repair
    Extended Community: RT:12:23 , recursive-via-connected
1
  10.3.3.3 (metric 21) from 10.6.6.6 (10.6.6.6)
    Origin incomplete, metric 0, localpref 200, valid, internal, best
    Extended Community: RT:12:23
    Originator: 10.3.3.3, Cluster list: 10.0.0.1 , recursive-via-host
    mpls labels in/out nolabel/37

```

The command output in the following example shows the BGP routing table entries for the backup and alternate paths:

```

Router# show ip bgp 10.0.0.0 255.255.0.0
BGP routing table entry for 10.0.0.0/16, version 123
Paths: (4 available, best #3, table default)
Additional-path
Advertised to update-groups:
  2          3
Local

```



```

10.0.101.4 from 10.0.101.4 (10.3.3.3)
  Origin IGP, localpref 100, weight 500, valid, internal
Local
10.0.101.3 from 10.0.101.3 (10.4.4.4)
  Origin IGP, localpref 100, weight 200, valid, internal
Local
10.0.101.2 from 10.0.101.2 (10.1.1.1)
  Origin IGP, localpref 100, weight 900, valid, internal, best
Local
10.0.101.1 from 10.0.101.1 (10.5.5.5)
  Origin IGP, localpref 100, weight 700, valid, internal, backup/repair

```

The command output in the following example shows the routing information base entries for the backup and alternate paths:

```
Router# show ip route repair-paths 10.0.0.0 255.255.0.0
```

```

Routing entry for 10.0.0.0/16
  Known via "bgp 10", distance 200, metric 0, type internal
  Last update from 10.0.101.2 00:00:56 ago
  Routing Descriptor Blocks:
  * 10.0.101.2, from 10.0.101.2, 00:00:56 ago
    Route metric is 0, traffic share count is 1
    AS Hops 0
    MPLS label: none
  [RPR]10.0.101.1, from 10.0.101.1, 00:00:56 ago
    Route metric is 0, traffic share count is 1
    AS Hops 0
    MPLS label: none

```

The command output in the following example shows the Cisco Express Forwarding/forwarding information base entries for the backup and alternate paths:

```
Router# show ip cef 10.0.0.0 255.255.0.0 detail
```

```

10.0.0.0/16, epoch 0, flags rib only nolabel, rib defined all labels
  recursive via 10.0.101.2
    attached to GigabitEthernet0/2
  recursive via 10.0.101.1, repair
    attached to GigabitEthernet0/2

```

Additional References

Related Documents

Related Topic	Document Title
Cisco IOS commands	Cisco IOS Master Commands List, All Releases
BGP commands: complete command syntax, command mode, command history, defaults, usage guidelines, and examples	Cisco IOS IP Routing: BGP Command Reference
Basic MPLS VPNs	Configuring MPLS Layer 3 VPNs
A failover feature that creates a new path after a link or node failure	MPLS VPN—BGP Local Convergence

Related Topic	Document Title
Configuring multiprotocol VRFs	<i>MPLS VPN—VRF CLI for IPv4 and IPv6 VPNs</i>
BGP routing	<i>BGP Feature Roadmap</i>

Standards

Standard	Title
draft-walton-bgp-add-paths-04.txt	<i>Advertisement of Multiple Paths in BGP</i>

MIBs

MIB	MIBs Link
None	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	<i>BGP/MPLS VPNs</i>
RFC 1771	<i>A Border Gateway Protocol 4 (BGP-4)</i>

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 BGP PIC

Table 1 lists the release history for this feature.

Use Cisco Feature Navigator to find information about platform support and software image support. Cisco Feature Navigator enables you to determine which software images support a specific software release, feature set, or platform. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>. An account on Cisco.com is not required.



Note

Table 1 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.

Table 1 Feature Information for BGP PIC

Feature Name	Releases	Feature Information
BGP PIC Edge for IP and MPLS-VPN	12.2(33)SRE 12.2(33)XNE 15.0(1)S	<p>The BGP PIC Edge for IP and MPLS-VPN feature improves BGP convergence after a network failure. This convergence is applicable to both core and edge failures and can be used in both IP and MPLS networks. The BGP PIC Edge for IP and MPLS-VPN feature creates and stores a backup/alternate path in the routing information base (RIB), forwarding information base (FIB), and Cisco Express Forwarding so that when a failure is detected, the backup/alternate path can immediately take over, thus enabling fast failover.</p> <p>In 12.2(33)SRE, this feature was introduced on the Cisco 7200 and Cisco 7600 routers.</p> <p>In 12.2(33)XNE, support was added for the Cisco 10000 router.</p> <p>The following sections provide information about this feature:</p> <ul style="list-style-type: none"> • Prerequisites for BGP PIC, page 2 • Restrictions for BGP PIC, page 2 • Information About BGP PIC, page 3 • How to Configure BGP PIC, page 12 <p>The following commands were introduced or modified: bgp additional-paths install, bgp recursion host, show ip bgp, show ip cef, show ip route, show vrf.</p>

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at 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. (1005R)

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.

© 2009–2011 Cisco Systems, Inc. All rights reserved.

