



BGP Prefix Independent Convergence

Restrictions:

- Ensure that the Border Gateway Protocol (BGP) and the IP or Multiprotocol Label Switching (MPLS) network is up and running at the customer site that is connected to the provider site by more than one path (multihomed).
- BGP PIC does not support instances where the sum of number of primary paths and backup paths is greater than 2. Hence, only one primary path and one backup path are supported.
- BGP PIC does not work when Label Distribution Protocol (LDP) and Segment Routing (SR) are enabled in IGP.
- Ensure that the backup or alternate path has a unique next hop that is not the same as the next hop of the best path.
- As a best practice, enable the Bidirectional Forwarding Detection (BFD) protocol to quickly detect link failures.
- When paths of different technologies are resolved over ECMP, it results in *heterogeneous* ECMP, leading to severe network traffic issues. Don't use ECMP for any combination of the following technologies:
 - LDP.
 - BGP-LU, including services over BGP-LU loopback peering or recursive services at Level-3.
 - VPNv4.
 - 6PE and 6VPE.
 - EVPN.
 - Recursive static routing.

BGP PIC: Export of Backup Path Agnostic to its Multipath Eligibility

The BGP PIC: Export of Backup Path Agnostic to its Multipath Eligibility feature improves BGP convergence after a network failure. This convergence applies to both core and edge failures. The BGP PIC pre-programs a backup path so that when a failure is detected, the backup path can immediately take over, thus enabling fast failover. This feature enables BGP PIC on VPNv4 with additional paths or when the multiple paths that are ineligible to be multipath are received from the same neighbor. For backup paths to be multipath eligible, all the following attributes in the backup paths must be the same: weight, local preference, autonomous system

path, origin code, Multi Exit Discriminator (MED), and Interior Gateway Protocol (iGP) distance. Also, the next hop router for each multipath must be different. This feature introduces flexibility to allow the import of backup paths to the VRF even if the said attributes are not the same.

Configuration Example

```
Router# router bgp 10
Router(config-bgp) # address-family vpnv4 unicast
router(config-bgp-af) # export to vrf allow backup
```

Running Configuration

```
router bgp 10
  address-family vpnv4 unicast
    export to vrf allow backup
```

- [BGP PIC Implementation Considerations, on page 2](#)
- [Configure BGP PIC, on page 2](#)
- [BGP-LU Multipath PIC with Auto Protection, on page 4](#)

BGP PIC Implementation Considerations

- BVI as a core is not supported.
- For labelled BGP loopback peering, the system supports only one primary and one backup path. No support for BGP PIC multipath protect.
- PIC EDGE is supported for all services, such as IPv4, IPv6, VPNv4, VPNv6, 6PE, 6VPE, VPWS, VPLS, and EVPN over labelled unicast address-family.

Configure BGP PIC

Procedure

Step 1 **cef encap-sharing disable**

Example:

```
RP/0/RP0/CPU0:router(config)# cef encap-sharing disable
```

By default, without primary and backup path installation in the hardware, IPv4, IPv6, 6PE (per-vrf), 6VPE (per-vrf/per-ce), L3VPN (per-vrf/per-ce) has good convergence.

When the mode is a per-prefix by default, BGP-PIC does not give good convergence, hence you must do hardware-assisted PIC. For this, configure the **cef encap-sharing disable** command in XR Config mode.

With hardware-assisted BGP PIC that is configured using the **cef encap-sharing disable** command, separate hardware resources (FEC/EEDB) are allocated for every prefix. Cisco recommends you to make sure that the router has sufficient hardware resources for the resource allocation.

Caution

This CLI reprograms the CEF completely and impacts traffic. We recommend that you do it in the maintenance window.

Note

- The **cef encap-sharing disable** command does not take effect in the SRv6 core.
- Reload the router for the **cef encap-sharing disable** command to take effect.

Step 2 `router bgp as-number`**Example:**

```
RP/0/RP0/CPU0:router(config)# router bgp 100
```

Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.

Step 3 `address-family {vpnv4 unicast | vpnv6 unicast | ipv4 unicast | ipv6 unicast}`**Example:**

```
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
address-family ipv4 unicast
  additional-paths receive
  additional-paths selection route-policy backup 1
  allocate-label all
!
```

Step 4 `additional-paths selection route-policy route-policy-name`**Example:**

```
RP/0/RP0/CPU0:router(config-bgp-af)# additional-paths selection route-policy ap1
```

Configures extra paths selection mode for a prefix.

Note

Use the **additional-paths selection** command with an appropriate route-policy to calculate backup paths and to enable Prefix-Independent Convergence (PIC) functionality.

The route-policy configuration is a prerequisite for configuring the additional-paths selection mode for a prefix. This is an example route-policy configuration to use with additional-selection command:

```
route-policy ap1
  set path-selection backup 1 install
end-policy
```

Configure BGP PIC Multipath

For multipath bgp-pic, configure the route-policy as following:

```
route-policy ap1
  set path-select backup 1 install multipath-protect [advertise]
!
```

BGP-LU Multipath PIC with Auto Protection

Table 1: Feature History Table

Feature Name	Release Information	Feature Description
BGP-LU Multipath PIC with Auto Protection	Release 7.5.2	<p>BGP-LU multipath prefix independent convergence (PIC) supports auto protection. Each active path has a backup path, ensuring almost immediate restoration of multicast traffic when a path fails.</p> <p>In earlier releases, multipath configuration supports primary and backup path with a limitation that backup path can support only one failed path at a time.</p>


From Cisco IOS XR Release 7.5.2, the BGP-LU multipath PIC with auto protection is supported only with multipath protect configuration. When you configure multipath, every active path is assigned with a backup path to improve the convergence.

Configure BGP-LU Multipath

Consider the topology BGP-LU multipath with PIC, where PE2 (R4) peering with R1, R2, and R5 by eBGP interfaces. R1, R2 and R5 peering with R6 by eBGP interfaces, and R6 and R3 peering by iBGP loopback via IGP instance.

In the topology, consider the IP addresses of routers as, R1 (14.1.1.2), R2 (14.2.1.2), R5 (14.5.1.2), R6 (15.4.1.2) and R4 (192.0.2.1).

Figure 1: BGP-LU Multipath PIC Edge

**Without multipath protect:**

The following example shows how to configure BGP-LU multipath with route policy, without multipath protect.

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router (config)#route-policy INSTALL_BACKUP
RP/0/RP0/CPU0:router (config-rpl)#set path-selection multipath advertise
RP/0/RP0/CPU0:router (config-rpl)#end-policy
RP/0/RP0/CPU0:router (config)#router bgp 300
RP/0/RP0/CPU0:router (config-bgp)#address-family ipv4 unicast
RP/0/RP0/CPU0:router (config-bgp-af)#additional-paths ebgp 4
RP/0/RP0/CPU0:router (config-bgp-af)#additional-paths selection route-policy INSTALL_BACKUP
```

With multipath protect:

The following example shows how to configure BGP-LU multipath with route policy, with multipath protect.

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router (config)#route-policy INSTALL_BACKUP
RP/0/RP0/CPU0:router (config-rpl)#set path-selection multipath advertise multipath-protect
RP/0/RP0/CPU0:router (config-rpl)#end-policy
```

Verification**Without multipath protect:**

The following example shows the maximum-paths configured under BGP address-family and BGP-LU multipath with route policy, without multipath protect. Here there's no PIC path. Perform the following from the R4 (PE2) router.

```
RP/0/RP0/CPU0:R4#show run route-policy INSTALL_BACKUP
Wed Apr 13 05:41:59.455 UTC
route-policy INSTALL_BACKUP
```

```

    set path-selection multipath advertise
end-policy
!

/* To display current routes for BGP in the Routing Information Base (RIB) */

RP/0/RP0/CPU0:R4#show route 192.0.2.1
Wed Apr 13 05:42:04.515 UTC

Routing entry for 192.0.2.0/24
  Known via "bgp 300", distance 20, metric 0, [ei]-bgp, labeled unicast (3107)
  Tag 400, type external
  Installed Apr 13 05:40:45.811 for 00:01:18
  Routing Descriptor Blocks
    14.1.1.2, from 14.1.1.2, BGP external, BGP multi path
      Route metric is 0
    14.2.1.2, from 14.2.1.2, BGP external, BGP multi path
      Route metric is 0
    14.5.1.2, from 14.5.1.2, BGP external, BGP multi path
      Route metric is 0
    15.4.1.2, from 15.4.1.2, BGP external, BGP multi path
      Route metric is 0
  No advertising protos.

/* To display information about packets forwarded by Cisco Express Forwarding (CEF) */

RP/0/RP0/CPU0:R4#show cef 192.0.2.1
Wed Apr 13 05:42:10.781 UTC
192.0.2.0/24, version 204023, internal 0x50000001 0x40 (ptr 0xae10fbf0) [1], 0x600
(0xa753d558), 0xa08 (0xa57515c8)
Updated Apr 13 05:40:46.013
Prefix Len 24, traffic index 0, precedence n/a, priority 4
  gateway array (0xa6819bb8) reference count 2856, flags 0x78, source rib (7), 0 backups
    [953 type 5 flags 0x8441 (0x912908c8) ext 0x0 (0x0)]
  LW-LDI[type=5, refc=3, ptr=0xa753d558, sh-ldi=0x912908c8]
  gateway array update type-time 1 Apr 9 22:03:22.505
  LDI Update time Apr 9 22:03:22.590
  LW-LDI-TS Apr 13 05:40:46.013
    via 14.1.1.2/32, 7 dependencies, recursive, bgp-ext, bgp-multipath [flags 0x60a0]
      path-idx 0 auto-bkup-idx 1 NHID 0x0 [0xe3b2b398 0x0], Internal 0xb62a8370
      recursion-via-/32
      next hop 14.1.1.2/32 via 24008/0/21
        local label 25420
        next hop 14.1.1.2/32 BE500 labels imposed {ImplNull 24210}
    via 14.2.1.2/32, 7 dependencies, recursive, bgp-ext, bgp-multipath [flags 0x60a0]
      path-idx 1 auto-bkup-idx 2 NHID 0x0 [0x905fcdd8 0x0], Internal 0xb62a6750
      recursion-via-/32
      next hop 14.2.1.2/32 via 24006/0/21
        local label 25420
        next hop 14.2.1.2/32 BE510 labels imposed {ImplNull 24211}
    via 14.5.1.2/32, 7 dependencies, recursive, bgp-ext, bgp-multipath [flags 0x60a0]
      path-idx 2 auto-bkup-idx 3 NHID 0x0 [0x905fef08 0x0], Internal 0xb62a6c00
      recursion-via-/32
      next hop 14.5.1.2/32 via 24005/0/21
        local label 25420
        next hop 14.5.1.2/32 Hu0/7/0/11 labels imposed {ImplNull 24211}
    via 15.4.1.2/32, 5 dependencies, recursive, bgp-ext, bgp-multipath [flags 0x60a0]
      path-idx 3 auto-bkup-idx 0 NHID 0x0 [0xe3b2b530 0x0], Internal 0xb62a7740
      recursion-via-/32
      next hop 15.4.1.2/32 via 24007/0/21
        local label 25420
        next hop 15.4.1.2/32 Te0/7/0/29/3.100 labels imposed {ImplNull 24210}

```

With multipath protect:

The following example shows the maximum-paths configured under BGP address-family and BGP-LU multipath PIC with multipath protect using route-policy. Here, an additional backup path created with PIC auto protection. Perform the following from the R4 (PE2) router.

By default, BGP-LU multipath PIC with auto protection is enabled when you configure multipath using route-policy.

```
RP/0/RP0/CPU0:R4#show run route-policy INSTALL_BACKUP
Wed Apr 13 06:00:08.384 UTC
route-policy INSTALL_BACKUP
  set path-selection multipath advertise multipath-protect
end-policy
!
```

```
/* To display current routes for BGP in the Routing Information Base (RIB) */

RP/0/RP0/CPU0:R4#show route 192.0.2.1
Wed Apr 13 06:00:23.365 UTC

Routing entry for 192.0.2.0/24
  Known via "bgp 300", distance 20, metric 0, [ei]-bgp, labeled unicast (3107)
  Tag 400
  Number of pic paths 1 , type external
  Installed Apr 13 06:00:01.885 for 00:00:21
  Routing Descriptor Blocks
    14.1.1.2, from 14.1.1.2, BGP external, BGP multi path
    Route metric is 0
    14.2.1.2, from 14.2.1.2, BGP external, BGP multi path
    Route metric is 0
    14.5.1.2, from 14.5.1.2, BGP external, BGP multi path
    Route metric is 0
    14.6.1.2, from 14.6.1.2, BGP external, BGP backup path
    Route metric is 0
    15.4.1.2, from 15.4.1.2, BGP external, BGP multi path
    Route metric is 0
  No advertising protos.
```

```
/* To display information about packets forwarded by Cisco Express Forwarding (CEF) */

RP/0/RP0/CPU0:R4#show cef 192.0.2.1
Wed Apr 13 06:00:33.884 UTC
192.0.2.1/24, version 208786, internal 0x5000001 0x40 (ptr 0xae10fbf0) [1], 0x600
(0xa753d558), 0xa08 (0xae2dc188)
Updated Apr 13 06:00:02.034
Prefix Len 24, traffic index 0, precedence n/a, priority 4
gateway array (0xa6828928) reference count 2856, flags 0x100078, source rib (7), 0 backups

[953 type 5 flags 0x8441 (0x912990e8) ext 0x0 (0x0)]
LW-LDI[type=5, refc=3, ptr=0xa753d558, sh-ldi=0x912990e8]
gateway array update type-time 1 Apr 13 06:00:01.954
LDI Update time Apr 13 06:00:01.991
LW-LDI-TS Apr 13 06:00:02.034
  via 14.1.1.2/32, 7 dependencies, recursive, bgp-ext, bgp-multipath [flags 0x60a0]
    path-idx 0 auto-bkup-idx 1 NHID 0x0 [0xe3b2b398 0x0], Internal 0xb62a8370
    recursion-via-/32
    next hop 14.1.1.2/32 via 24008/0/21
    local label 25420
    next hop 14.1.1.2/32 BE500 labels imposed {ImplNull 24210}
  via 14.2.1.2/32, 7 dependencies, recursive, bgp-ext, bgp-multipath [flags 0x60a0]
    path-idx 1 auto-bkup-idx 2 NHID 0x0 [0x905fcd8 0x0], Internal 0xb62a6750
    recursion-via-/32
    next hop 14.2.1.2/32 via 24006/0/21
```

```

    local label 25420
    next hop 14.2.1.2/32 BE510          labels imposed {ImplNull 24211}
via 14.5.1.2/32, 7 dependencies, recursive, bgp-ext, bgp-multipath [flags 0x60a0]
  path-idx 2 auto-bkup-idx 4 NHID 0x0 [0x905fef08 0x0], Internal 0xb62a6c00
  recursion-via-/32
  next hop 14.5.1.2/32 via 24005/0/21
    local label 25420
    next hop 14.5.1.2/32 Hu0/7/0/11    labels imposed {ImplNull 24211}
via 14.6.1.2/32, 6 dependencies, recursive, bgp-ext, backup [flags 0x6120]
  path-idx 3 NHID 0x0 [0xe3b2b968 0x0]
  recursion-via-/32
  next hop 14.6.1.2/32 via 24003/0/21
    local label 25420
    next hop 14.6.1.2/32 Te0/7/0/15/3 labels imposed {ImplNull 24362}
via 15.4.1.2/32, 6 dependencies, recursive, bgp-ext, bgp-multipath [flags 0x60a0]
  path-idx 4 auto-bkup-idx 0 NHID 0x0 [0xe3b2b530 0x0], Internal 0xb62a7740
  recursion-via-/32
  next hop 15.4.1.2/32 via 24007/0/21
    local label 25420
    next hop 15.4.1.2/32 Te0/7/0/29/3.100 labels imposed {ImplNull 24210}

```

Usage Guidelines and Limitations

- Multipath PIC for L3VPN over labeled unicast, over Interior Gateway Protocol (IGP) with sub-second convergence may not be applicable.
- BGP-LU multipath PIC with auto protection is supported for both interface and loopback peering.