



BGP Prefix Independent Convergence

Table 1: Feature History Table

Feature Name	Release Information	Feature Description
BGP PIC Edge for Unlabeled Transport (IPv4/v6)	Release 7.3.1	This feature is now supported on routers that have Cisco NC57 line cards installed and operate in the native mode.

Table 2: Feature History Table

Feature Name	Release Information	Feature Description
BGP PIC: Export of Backup Path Agnostic to its Multipath Eligibility	Release 7.3.2	Prior to this release, you could only import the backup path of a prefix to the respective VRFs only when the backup path is multipath eligible. For a backup path to be multipath eligible, all the following attributes in the backup path must the same as the best path: 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 import of a backup path to the VRF even if the said attributes are not the same as that of the best path.

Table 3: Feature History Table

Feature Name	Release Information	Feature Description
BGP PIC Backup Path when the Primary Path is a Static Route with the next hop as an IP Address.	Release 7.5.1	This feature is now supported on routers that have Cisco NC57 line cards installed and operate in native and compatibility mode. This feature enables BGP PIC backup path when the primary path is a static route with the next hop as an IP Address.

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
```

- [Convergence for BGP Labeled Unicast PIC Edge, on page 3](#)
- [BGP PIC Implementation Considerations, on page 6](#)
- [Configure BGP PIC, on page 6](#)
- [BGP-LU Multipath PIC with Auto Protection, on page 8](#)
- [TCAM Enhancement to Improve L3VPN Routing Capability, on page 13](#)

Convergence for BGP Labeled Unicast PIC Edge

Table 4: Feature History Table

Feature Name	Release Information	Feature Description
Convergence for BGP Labeled Unicast PIC Edge	Release 7.7.1	This feature improves the convergence time of BGP labeled unicast (LU) routes to subseconds when an ingress provider edge router fails or loses PE router connectivity, and another PE router needs to be connected. This feature minimizes traffic drops when the primary paths fail for the BGP LU routes. This feature is now supported on routers that have Cisco NC57 line cards installed and operate in native and compatibility mode`

BGP Labeled Unicast (LU) PIC Edge feature enables you to create and store both the primary and backup path in the Routing Information Base (RIB), Forwarding Information Base (FIB), and Cisco Express Forwarding. When the router detects a failure, the backup or alternate path immediately takes over, thus this feature enables fast failover and convergence in subseconds.


For BGP LU PIC Edge to work, the edge iBGP devices, such as ingress PEs and Autonomous System Border Router (ASBR), must support BGP PIC and must receive backup BGP next hop.

The topology diagram given below illustrates the Convergence for BGP Labeled Unicast PIC Edge feature. The topology is explained as follows:

- The BGP LU PIC Edge feature is enabled on a provider edge router, PE1.
- PE1 learns the BGP LU prefix from the remote PE router, PE2.
- PE1 routes traffic through the Area Border Routers, ABR1, ABR2 and ABR3. If one of them fails, the preprogrammed backup of the failed ABR routes the traffic.
- PE1 routes traffic through the Area Border Routers, ABR1, ABR2 and ABR3.
- PE2 is marked as the backup or alternate next hop and is programmed into the FIB of PE1.
- When PE1 learns PE2 is not reachable through ABR1, it immediately changes the BGP next hop for the PE1's prefix to ABR2.
- The switchover occurs in less than a second regardless of the number of prefixes.
- Subsecond convergence occurs although updates to multiple BGP prefixes are pending.

Topology

Figure 1: BGP LU PIC Edge



Guidelines and Limitations

This feature supports BGP multipaths that allows the router to install multiple internal BGP paths and multiple external BGP paths to the forwarding table. The multiple paths enable BGP to load balance traffic across multiple links.

The convergence time is independent of the BGP LU route scale.

Configure Convergence for BGP Labeled Unicast PIC Edge

Perform the following steps to configure Convergence for BGP Labeled Unicast PIC Edge:

- Configure BGP labeled unicast and attach route-policy to BGP address families.
- Configure BGP labeled unicast multipath and attach route-policy to BGP address families

```
Router(config)# route-policy BGP-PIC-EDGE
Router(config-rpl)# set path-selection backup 1 install
Router(config-rpl)# end-policy
Router(config)# end
Router(config)# router bgp 200
```

```

Router(config-bgp)# bgp router-id 10.0.0.1
Router(config-bgp)# address-family ipv4 unicast
Router(config-bgp-af)# additional-paths receive
Router(config-bgp-af)# additional-paths send
Router(config-bgp-af)# additional-paths selection route-policy BGP-PIC-EDGE

/*Perform the following steps to configure BGP labeled unicast multipath and attach
route-policy to BGP address families: */
Router(config)# route-policy BGP-PIC-EDGE-MULTIPATH
Router(config-rpl)# set path-selection backup 1 install multipath-protect
Router(config)# end-policy
Router(config)# router bgp 200
Router(config)# bgp router-id 192.168.1.0
Router(config)# address-family ipv4 unicast
Router(config)# maximum-paths ibgp 2
Router(config)# additional-paths receive
Router(config)# additional-paths send
Router(config)# additional-paths selection route-policy BGP-PIC-EDGE-MULTIPATH

```

Running Configuration

```

route-policy BGP-PIC-EDGE
  set path-selection backup 1 install
  end-policy
router bgp 200
  bgp router-id 192.168.1.0
  address-family ipv4 unicast
    additional-paths receive
    additional-paths send
    additional-paths selection route-policy BGP-PIC-EDGE

route-policy BGP-PIC-EDGE-MULTIPATH
  set path-selection backup 1 install multipath-protect
  end-policy
router bgp 200
  bgp router-id 192.168.1.0
  address-family ipv4 unicast
    maximum-paths ibgp 2
    additional-paths receive
    additional-paths send
    additional-paths selection route-policy BGP-PIC-EDGE-MULTIPATH

```

Verification

Verify that the backup path is established.

```

Router# show cef 192.0.2.1/32
192.168.0.0/32, version 31, internal 0x5000001 0x40 (ptr 0x901d2370) [1], 0x0 (0x90d2beb8),
0xa08 (0x91c74378)
Prefix Len 32, traffic index 0, precedence n/a, priority 4
  via 203.0.113.1/32, 3 dependencies, recursive [flags 0x6000] << Primary Path
    path-idx 0 NHID 0x0 [0x90319650 0x0]
    recursion-via-/32
    next hop 192.51.100.1/32 via 24006/0/21
    next hop 209.165.200.225/32 Hu0/0/0/25 labels imposed {24002 24000}
    next hop 10.0.0.1/32 Hu0/0/0/26 labels imposed {24002 24000}
  via 203.0.113.2/32, 2 dependencies, recursive, backup [flags 0x6100] << Backup Path
    path-idx 1 NHID 0x0 [0x903197b8 0x0]
    recursion-via-/32
    next hop 209.165.200.225/32 via 24005/0/21

```

```
next hop 192.51.100.1/32 Hu0/0/0/25 labels imposed {24001 24000}
next hop 10.0.0.1/32 Hu0/0/0/26 labels imposed {24001 24000}
```

BGP PIC Implementation Considerations

- TE, SR-TE, SRv6, flex-LSP in the core are not supported.
- BGP PIC over BVI (core or edge) 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.
- The system supports BGP PIC multipath protect only for unlabelled BGP IPv4 and IPv6 address family. The support is for both loopback peering and interface peering mode.
- The system supports BGP PIC multipath protect only for labelled BGP in interface peering mode.
- Labelled BGP over IPv6 core BGP PIC is not supported.

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 5: Feature History Table

Feature Name	Release Information	Feature Description
BGP-LU Multipath PIC with Auto Protection	Release 7.5.2	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. 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.

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).



Note BGP-Labeled Unicast (LU) Prefix-Independent Convergence (PIC) auto-protection feature may cause equal cost multipath (ECMP) FEC NPU resource exhaustion on BGP peering devices for IPv4/IPv6 addresses. From Cisco IOS XR Release 7.5.4 Release 7.8.2 Release 7.9.1 onwards, the auto-protection feature for BGP-LU multipath PIC is disabled by default. To enable this feature, use the **hw-module fib bgp-mp-pic auto-protect enable** command. After executing the command, you must reload the router. For more information, see *BGP-LU Multipath PIC with Auto Protection* section in *BGP Prefix Independent Convergence* chapter in *BGP Configuration Guide for Cisco NCS 5500 Series Routers*.

Figure 2: 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 [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 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.

TCAM Enhancement to Improve L3VPN Routing Capability

Table 6: Feature History Table

Feature Name	Release Name	Description
Ternary Content-Addressable Memory Enhancement to Improve L3VPN Routing Capability	Release 7.11.1	Introduced in this release on: NCS 5700 fixed port routers; NCS 5500 modular routers (NCS 5700 line cards [Mode: Native]) You can improve your router's performance by introducing two additional Ternary Content-Addressable Memory (TCAM) labels. This enhancement removes the dependence on the Forwarding Equivalence Class (FEC) scale for handling more routes, which means that the router's capacity is no longer limited by FEC capacity. As a result, we have significantly increased the number of routes for L3VPN. The feature introduces these changes: CLI: • <code>remote-bgp-label-mode-per-prefix</code> • <code>hw-module fib mpls two-label-tcam-optimized</code>

Before this feature, the per-prefix label allocation mode for L3VPN prefixes at FIB for BGP PIC backup paths consumed two Forwarding Equivalence Class (FEC) hierarchy two resources per prefix. As a result, when considering a unidimensional scale, the NCS 5700 fixed port router in native mode equipped with a second-generation external TCAM (OP2) supported a limited scale for BGP Prefix-Independent Convergence (PIC) limited by hierarchy two FEC resources present in the system.

This feature allocates space for two labels from the OP2 TCAM into the router that uses BGP PIC, optimizing the usage of FEC hardware resources. This feature enables the router to significantly enhance unidimensional scaling by increasing the number of L3VPN routes with BGP PIC. The number of routes that can be supported is entirely dependent on the capacity of the Knowledge-based Processor (KBP) TCAM to hold route entries with 64-bit (two label) result.

This feature exclusively addresses BGP PIC Unipath scenarios. It is important to note that this feature is available on the NCS 5700 fixed port router and NCS 5500 modular router in native mode.

Run the **hw-module fib mpls two-label-tcam-optimized** to enable this feature.

With this feature, the router is no longer dependent on Forwarding Equivalence Class (FEC) scale for L3VPN routes and improves performance and removes any limitations previously imposed by the router's FEC capacity for L3VPN routes.

Removing the dependency on FEC resources, which restrict L3VPN route capacity, enables the router to efficiently manage a significantly larger number of routes, overcoming previous limitations

This feature allows NCS 5700 fixed port routers to configure both primary and backup Multiprotocol Label Switching (MPLS) labels in the KBP TCAM for a VPN route. Hence the routers can now effectively manage MPLS labels for VPN routes by setting up both primary and backup label entries in the KBP TCAM.

This feature is disabled by default. The feature should be enabled only for the VRFs which receive prefixes in the per-prefix mode from the remote edge router to optimize TCAM resources.

Restrictions for TCAM Enhancement to Improve L3VPN Routing Capability

- The Quality of Service Policy Propagation for BGP (QPPB) functionality is not compatible with prefixes from a Layer 3 Virtual Private Network (L3VPN) configured with Border Gateway Protocol (BGP) Prefix-Independent Convergence (PIC). However, other types of prefixes, such as IPv6 Virtual Private Network (6VPE), Global unlabeled routes, and prefixes without PIC, can still utilize QPPB configuration.
- This feature is disabled for L3VPN prefixes that possess an associated local label, such as L3VPN Option-B routes.
- L3VPN prefixes utilizing BGP-PIC with both primary and backup labels stored in KBP TCAM implement the TTL propagation. As a result, the **mpls ip-ttl-propagate** command does not affect these specific prefixes. However, it continues to apply to all other prefixes this particular feature does not affect.
- This feature is enabled only for VPNv4 and VPNv6 BGP PIC unipath prefixes. This feature does not apply to multipath prefixes.
- This feature can be enabled on NCS 5700 fixed port TCAM routers by applying the MDB SE profile for scaling purposes.

This feature can be enabled on NCS 5500 and NCS 5700 fixed port routers with external TCAM by applying the MDB SE profile for scaling purposes.

- This feature supports L3VPN prefixes in recursive level 2.

Configure TCAM Enhancement to Improve L3VPN Routing Capability

Configuration Example

Perform the following steps to configure TCAM Enhancement to Improve L3VPN Routing Capability. Depending on whether the scale optimization is applied universally to all L3VPN prefixes or selectively to specific ones, two options can be realized using the keywords: **selective** or **all**.

If you want to apply the scale optimization on all the L3VPN prefixes in a designated address family:

1. Use the **hw-module fib mpls two-label-tcam-optimized** command with the **selective** keyword to enable the two push labels to encapsulate traffic for a VPN in the knowledge-based processor (KBP) ternary content access memory (TCAM).
2. Configure the designated VRF.

3. Configure the **remote-bgp-label-mode-per-prefix** command with either IPv4 or IPv6 keyword, or both, as required, under the global VRF configuration. You must configure the **remote-bgp-label-mode-per-prefix** command after you configure the **hw-module fib mpls two-label-tcam-optimized selective** command, to optimize all prefixes in a designated VRF within a specific address family. All other prefixes outside this configuration adhere to the default behavior.
4. Specify the route target values to control the distribution of routes among VPNs.
5. Reload the router for the configuration to take effect.

```

Router(config)# hw-module fib mpls two-label-tcam-optimized selective
Router(config)# vrf vrf1
Router(config-vrf)# ipv4 remote-bgp-label-mode-per-prefix
Router(config-vrf)# ipv6 remote-bgp-label-mode-per-prefix
Router(config-vrf)# address-family ipv4 unicast
Router(config-vrf-af)# import route-target
Router(config-vrf-af)# 1:1
Router(config-vrf-af)# export route-target
Router(config-vrf-af)# 1:1
Router(config-vrf-af)# exit
Router(config-vrf)# address-family ipv6 unicast
Router(config-vrf-af)# import route-target
Router(config-vrf-af)# 1:1
Router(config-vrf-af)# export route-target
Router(config-vrf-af)# 1:1
Router(config-vrf-af)# exit
Router(config-vrf)# exit
Router(config)# exit

```

If you want to apply the scale optimization on all the L3VPN prefixes, regardless of the VRF instances, use the **hw-module fib mpls two-label-tcam-optimized** command with **all** keyword. Reload the router for the configuration to take effect.

Running Configuration

```

hw-module fib mpls two-label-tcam-optimized selective
vrf vrf1
    remote-bgp-label-mode-per-prefix
    remote-bgp-label-mode-per-prefix
address-family ipv4 unicast
    import route-target
    1:1
    !
    export route-target
    1:1
    !
    !
address-family ipv6 unicast
    import route-target
    1:1
    !
    export route-target
    1:1
    !
    !
!

```

