L3VPN and L2VPN Services over BGP-LU over LDP over RSVP-TE

L3VPN and L2VPN services traditionally rely on a single MPLS transport layer for end-to-end reachability. In large-scale service provider networks, transport domains are often segmented and require inter-domain MPLS reachability while maintaining traffic engineering, resiliency, and operational scalability.

The BGP-LU over LDP over RSVP-TE solution enables scalable transport of L2VPN and L3VPN services across multiple MPLS domains. Recursive MPLS forwarding allows VPN traffic to leverage BGP-LU reachability while maintaining RSVP-TE traffic engineering, fast reroute, path protection, and BGP-PIC capabilities. The feature is supported on Cisco 8000 platforms based on the K100 and P100 fixed ASIC and provides transport feature parity between L2VPN and L3VPN services.

Table 1: Feature History Table

Feature Name	Description
L3VPN and L2VPN Services over BGP-LU over LDP over RSVP-TE	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100, K100]) This feature supports three-level MPLS label recursion using BGP Labeled Unicast (BGP-LU) over LDP over RSVP-TE transport solution and enables Layer 2 VPN (L2VPN) and Layer 3 VPN (L3VPN) transport across multiple MPLS domains using recursive MPLS label resolution. VPN service routes are resolved through BGP-LU, then through LDP transport paths carried over RSVP-TE tunnels.

Hierarchical forwarding model

This feature architecture allows VPN service routes to recursively resolve through these layers:

VPN Route -> BGP-LU Route -> LDP Route -> RSVP-TE Tunnel

This hierarchical forwarding model combines the strengths of each of these technologies:

- VPN labels: Provide service identification and isolation.
- BGP-LU labels: Enable inter-domain MPLS reachability.

- LDP labels: Distribute transport labels within each IGP domain.
- RSVP-TE tunnels: Provide traffic-engineered transport and protection.

Feature architecture benefits

The feature architecture delivers:

- Traffic Engineering (TE)
- Fast Reroute (FRR)
- Path Protection
- BGP-LU PIC
- Multi-domain MPLS transport
- Service transport parity between L2VPN and L3VPN

This chapter contains the following sections:

Prerequisites for L3VPN and L2VPN over BGP-LU over LDP over RSVP-TE

To successfully implement L3VPN and L2VPN services over BGP-LU, LDP, and RSVP-TE, ensure that you have met all of the following requirements:

- Confirm OSPF or IS-IS reachability within each domain.
- Enable MPLS forwarding on all relevant devices.
- Establish RSVP-TE tunnels within all relevant IGP domains.
- Set up LDP sessions along the intended transport path.
- Configure BGP Labeled Unicast between PE and Aggregation (AGG) routers.
- Configure VPNv4 or VPNv6 services configured on PE devices.

Restrictions for L3VPN and L2VPN over BGP-LU over LDP over RSVP-TE

Follow these requirements:

- Not supported on Q100 and Q200 ASIC platforms.
- Supported only on K100-based and P100-based fixed Cisco 8000 systems.

Supported L3VPN and L2VPN services over BGP-LU over LDP over RSVP-TE

The following table summarizes the technologies and the services they support in this context:

Table 2: Supported L3VPN and L2VPN services

Technology	Supported services
L3VPN	• IPv4 L3VPN • IPv6 L3VPN • Single-hop and multi-hop tunnels • BGP-LU PIC • RSVP-TE FRR • RSVP-TE Path Protection
L2VPN	• VPWS • VPLS • Active/Standby Pseudowires • Multi-Segment Pseudowires (MSPW) • RSVP-TE FRR • RSVP-TE Path Protection

These services are supported when deploying L3VPN and L2VPN solutions over the BGP-LU, LDP, and RSVP-TE protocol stack. This enables flexible VPN connectivity and resilience features for network operators.

Supported features for L3VPN and L2VPN over BGP-LU over LDP over RSVP-TE

We recommend implementing the following features to ensure full compatibility and optimal performance when deploying L3VPN and L2VPN services over BGP-LU, LDP, and RSVP-TE:

- L2VPN and L3VPNoBGPLUoLDPoTE with single or multi-hop tunnels.
- L2VPN and L3VPNoBGPLUoLDPoTE with TE protection.
- L2VPN and L3VPNoBGPLUoLDPoTE with BGP-LU PIC.
- Migration capabilities for legacy BGP-LU flows to this new approach.
- L2VPNoBGPLUoLDPoTE with VPLS and VPWS with active or standby pseudowire (PW) and multi-segment pseudowire (MSPW).
- Transport feature parity maintenance between L2VPN and L3VPN services.

These features help optimize interoperability and resilience for L3VPN and L2VPN deployments across supported protocols.

Network topology for L3VPN and L2VPN services over BGP-LU over LDP over RSVP-TE

The topology contains multiple IGP domains interconnected through aggregation routers that act as BGP-LU border nodes and exchange labeled reachability.

These are the main components in the topology:

- PE1 and PE2: Provide VPN services.
- AGG1 and AGG2: Act as BGP-LU border nodes between domains.
- OSPF1, OSPF2, and OSPF3: Represent independent IGP domains.
- RSVP-TE tunnels: Provide traffic-engineered transport.
- LDP: Distributes transport labels inside each IGP domain.
- BGP-LU: Distributes labeled reachability across domains.

Figure 1: L3VPN and L2VPN Services over BGP-LU over LDP over RSVP-TE

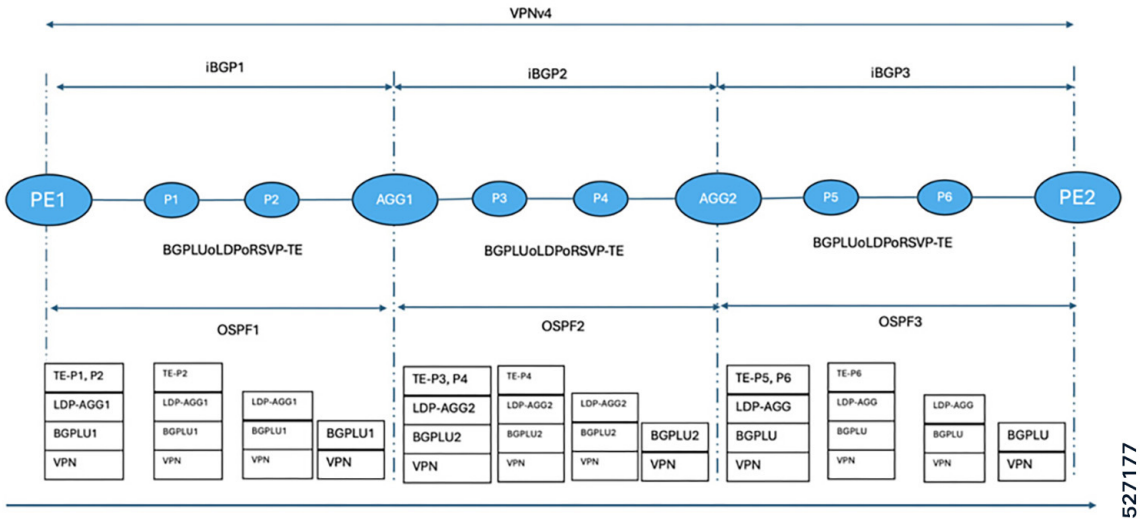

Layered forwarding architecture

Layer	Function
VPN	Service identification
BGP-LU	Inter-domain MPLS reachability
LDP	MPLS transport within IGP domain
RSVP-TE	Traffic engineered forwarding path

How the routers forward VPN packets

The forwarding stack installed on P1 or AGG1 consists of the VPN service label, BGP-LU label, LDP transport label, and RSVP-TE tunnel label. This recursive label resolution enables VPN traffic to be transported across multiple MPLS domains while leveraging RSVP-TE traffic engineering and protection capabilities.

For the sample prefix `99.0.0.1/28`, the label stack used consists of:

- VPN label: `30001`
- BGP-LU label: `40710`
- LDP label: implicit null
- RSVP-TE tunnel label: `36576`

This stack enables VPN traffic to be encapsulated and router from the ingress PE to the destination VPN route, across multiple transport and engineering domains.

The **show cef** output confirms labels imposed {ImplNull 40710 30001}. The show mpls forwarding prefix 1.1.1.1/32 detail output confirms that the LDP resolution is using tunnel-te5462 with TE label 36576 and Imp-Null.

Summary

Routers forward VPN packets across a network using a recursive label stack that enables traffic crossing multiple transport domains. This process involves assigning protocol-specific labels and encapsulating packets to ensure correct routing through layer 3 VPNs, leveraging both traffic engineering and transport mechanisms.

The key components involved in the process are:

- Provider Edge (PE) routers: Push VPN and BGP-LU labels to encapsulate packet for inter-domain transport.

  Provider (P or AGG) routers: Exchange BGP-LU reachability, resolve label stacks, and provide the transport path across each OSPF domain.

- Transport protocols: LDP provides transport within each OSPF domain, and RSVP-TE provides protected transport only between P1 or AGG1 and P3 or AGG2 for traffic engineering and resilience.

Workflow

Figure 2: Forwarding Label Stack Visibility on P1/AGG1 for VPN Services over BGP-LU over LDP over RSVP-TE

This following figure illustrates the forwarding state programmed on P1/AGG1.

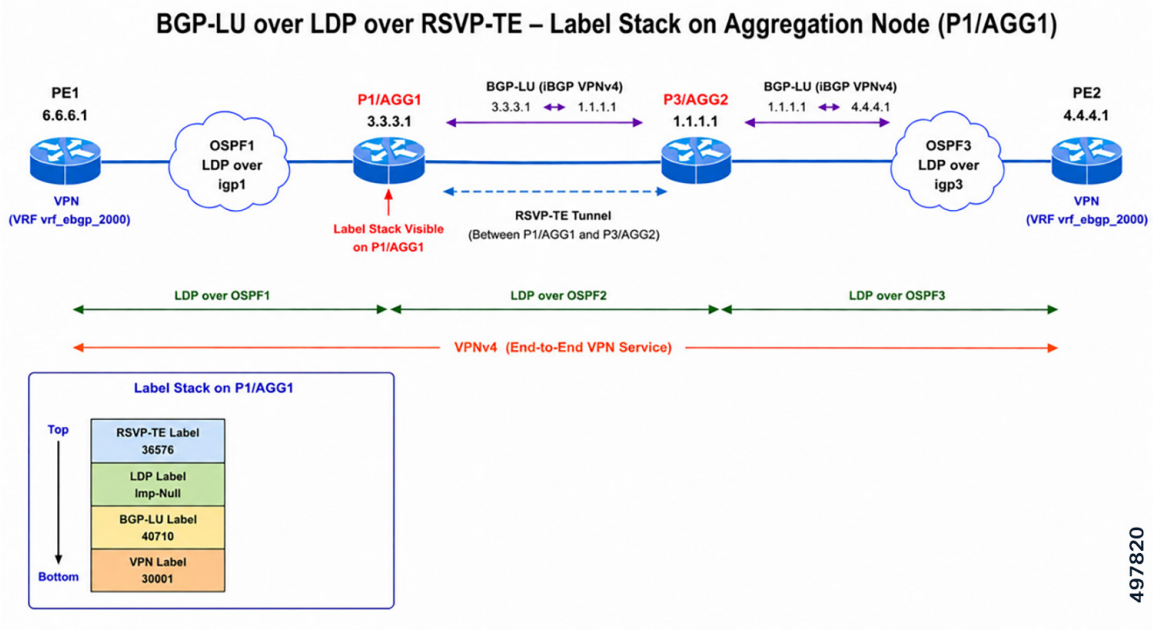

These stages describe how the router forwards VPN packets across the recursive transport stack:

1. The ingress PE pushes the VPN service label for the destination VPN route.
2. The same ingress PE pushes the BGP-LU label for the inter-domain next hop.
3. The transport path resolves through LDP over the RSVP-TE tunnel, ensuring correct routing through each OSPF domain.

4. The packet is forwarded across the RSVP-TE tunnel toward the AGG next hop.
5. The penultimate router removes the outer transport label where implicit-null applies.
6. The AGG router receives the packet, swaps the BGP-LU label for the next transport label, and forwards the packet toward the remote side and eventual VPN destination.

Result

This forwarding process enables VPN packets to traverse multiple MPLS domains seamlessly, using recursive label stacking to leverage traffic engineering, protection, and optimal routing across the provider network.

Operational information for L3VPN and L2VPN over BGP-LU over LDP over RSVP-TE

L3VPN and L2VPN services over BGP-LU over LDP over RSVP-TE uses recursive MPLS label resolution to transport L3VPN and L2VPN services across multiple MPLS domains. Provide Edge (PE) routers resolve VPN routes through BGP-LU routes, LDP transport paths, and RSVP-TE tunnels.

Control plane operation

1. AGG-1 allocates a local MPLS label to its loopback route and advertises the labeled route to PE1 through BGP-LU.
2. Other AGG routers in the path receive the labeled route, install it, allocate their own local labels, and advertise the route toward the upstream AGG router.
3. PE1 learns BGP-LU reachability for AGG-1 and for remote destinations beyond AGG-1.
4. In each IGP domain, LDP provides MPLS transport reachability to the next BGP-LU hop.
5. Where traffic engineering is required, the LDP path resolves over RSVP-TE tunnels.
6. PE and AGG routers can install IP and MPLS forwarding entries for BGP-LU destinations. A PE that does not reflect BGP-LU routes might install only the IP route and might not allocate a local label for those routes.

Data plane forwarding

1. PE1 receives a packet for a remote VPN destination.
2. PE1 pushes the VPN service label learned from the remote PE.
3. PE1 pushes the BGP-LU label learned for the BGP next hop, which is the AGG-1 loopback.
4. PE1 pushes the transport label stack that resolves the AGG-1 loopback through LDP over RSVP-TE.
5. The transport label carries the packet to AGG-1.
6. After penultimate hop popping, AGG-1 receives the packet with the BGP-LU label and the VPN label.
7. AGG-1 swaps the BGP-LU label with the transport label for the next domain and forwards the packet toward the remote node.

8. The downstream transport path carries the packet to the remote side, where the VPN label identifies the destination service.

Configure an aggregation router for L3VPN and L2VPN services over BGP-LU over LDP over RSVP-TE

Set up your aggregation (AGG) node for transporting L3VPN and L2VPN services using BGP label unicast (BGP-LU), LDP, and RSVP-TE tunneling.

Configure your router to support advanced VPN transport mechanisms for scalable and resilient service delivery over an MPLS network.

Follow these steps to configure your aggregation router for recursive VPN transport over BGP-LU, LDP, and RSVP-TE:

Before you begin

- Ensure all physical connections, hardware, and interface allocations required for L3VPN and L2VPN operations are in place.
- Confirm that IP addressing and connectivity for all relevant nodes are validated.
- Review and complete prerequisites for L3VPN and L2VPN services over BGP-LU, LDP, and RSVP-TE.
- Obtain necessary access rights and prepare configuration examples for each service.

Procedure

Step 1 Configure the VRF and route targets for the VPN service.

Step 2 Configure global BGP parameters, including the router ID, cluster ID, graceful restart, and labeled-unicast label allocation.

Step 3 Configure the IPv4 and IPv6 labeled-unicast address families under BGP neighbors.

Step 4 Configure the VPNv4 and VPNv6 address families.

Step 5 Configure the PE and AGG neighbors.

Step 6 Configure the VRF address families and enable per-VRF label mode.

Step 7 Enable MPLS LDP, configure the LDP router ID, and enable host-route label allocation.

Step 8 Enable MPLS traffic engineering and BFD on the transport interfaces.

Step 9 Enable RSVP on the transport interfaces.

Step 10 Configure the TE tunnel with destination, autoroute announce, fast reroute, path protection, and explicit path option.

Example:

Overall running configuration

```
#VRF
vrf vrf_ebgp_2000
 address-family ipv4 unicast
  import route-target
   101:2000
```

```
   201:2000
  !
  export route-target
   101:2000
   201:2000
  !
 !
 address-family ipv6 unicast
  import route-target
   101:2000
   201:2000
  !
  export route-target
   101:2000
   201:2000
  !

#BGP
router bgp 1
 nsr
 bfd minimum-interval 150
 bfd multiplier 3
 bgp router-id 3.3.3.1
 bgp cluster-id 2000
 bgp graceful-restart
 ibgp policy out enforce-modifications
 address-family ipv4 unicast
  additional-paths receive
  additional-paths send
  additional-paths selection route-policy INSTALL_BACKUP
  nexthop trigger-delay critical 0
  nexthop trigger-delay non-critical 5
  network 3.3.3.1/32 route-policy PASS_ALL
  allocate-label route-policy PASS_ALL unlabeled-path
 !
 address-family vpnv4 unicast
  label mode per-nexthop-received-label
  additional-paths selection route-policy INSTALL_BACKUP
  retain route-target all
 !
 address-family ipv6 unicast
  additional-paths receive
  additional-paths send
  additional-paths selection route-policy INSTALL_BACKUP
  nexthop trigger-delay critical 0
  nexthop trigger-delay non-critical 5
  allocate-label route-policy PASS_ALL
 !
 address-family vpnv6 unicast
  label mode per-nexthop-received-label
  additional-paths selection route-policy INSTALL_BACKUP
  retain route-target all
 !
 neighbor 1.1.1.1
  remote-as 1
  bfd fast-detect
  password encrypted 04590C163035495D1D48
  update-source Loopback1
  address-family ipv4 labeled-unicast
   maximum-prefix 4294967295 75 warning-only
   next-hop-self
   route-policy PASS_ALL in
   route-policy PASS_ALL out
  !
```

```
 address-family vpnv4 unicast
  route-policy PASS_ALL in
  route-reflector-client
  route-policy block_default_route_v4 out
 !
 address-family ipv6 labeled-unicast
  maximum-prefix 4294967295 75 warning-only
  next-hop-self
  route-policy PASS_ALL in
  route-reflector-client
  route-policy PASS_ALL out
 !
 address-family vpnv6 unicast
  route-policy PASS_ALL in
  route-reflector-client
  route-policy block_default_route_v6 out
 !
!
!
neighbor 4.4.4.1
 remote-as 1
 password encrypted 110B1E1528060E1F107B
 update-source Loopback1
 address-family vpnv4 unicast
  route-policy PASS_ALL in
  route-reflector-client
  route-policy default_route_v4 out
 !
 address-family ipv6 labeled-unicast
  maximum-prefix 4294967295 75 warning-only
  next-hop-self
  route-policy PASS_ALL in
  route-policy PASS_ALL out
 !
 address-family vpnv6 unicast
  route-policy PASS_ALL in
  route-reflector-client
  route-policy default_route_v6 out
 !
!
 !
!
neighbor 6.6.6.1
 remote-as 1
 bfd fast-detect
 password encrypted 06040831735A0C0A1146
 update-source Loopback1
 address-family vpnv4 unicast
  route-policy PASS_ALL in
  route-reflector-client
  route-policy default_route_v4 out
 !
 address-family ipv6 labeled-unicast
  maximum-prefix 4294967295 75 warning-only
  next-hop-self
  route-policy PASS_ALL in
  route-reflector-client
  route-policy PASS_ALL out
 !
 address-family vpnv6 unicast
  route-policy PASS_ALL in
  route-reflector-client
  route-policy default_route_v6 out
 !
```

```
 !
 vrf vrf_ebgp_2000
  rd auto
  bgp router-id 3.3.3.50
  address-family ipv4 unicast
   label mode per-vrf
   network 0.0.0.0/0
   redistribute connected
   redistribute static
  !
  address-family ipv6 unicast
   label mode per-vrf
   network ::/0
   redistribute connected
   redistribute static
  !

#MPLS LDP
mpls ldp
 nsr
 graceful-restart
 router-id 3.3.3.1
 address-family ipv4
  discovery targeted-hello accept
  label
   local
    allocate for host-routes
   !
  !
 !
 interface tunnel-te5462
  address-family ipv4
  !
 !
 interface tunnel-te5463
  address-family ipv4
  !
 !
 interface tunnel-te5464
  address-family ipv4
  !
 !
 interface tunnel-te5465
  address-family ipv4

#MPLS TE
mpls traffic-eng
 interface Bundle-Ether110
  bfd fast-detect
 !
 interface Bundle-Ether310
  bfd fast-detect

interface HundredGigE0/0/0/28
bfd fast-detect
 !
 interface Bundle-Ether300.100
  bfd fast-detect

#RSVP
interface Bundle-Ether110
interface Bundle-Ether310
interface HundredGigE0/0/0/28
```

```
interface Bundle-Ether300.100

#TE tunnel
interface tunnel-te5462
 description P1_P3
 ipv4 unnumbered Loopback1
 ipv6 enable
 logging events lsp-status state
 autoroute announce
 !
 destination 1.1.1.1
 fast-reroute
 path-protection
 !
 path-option 1 explicit name pri-tunnel-5462
```

Your aggregation router is now configured for L3VPN and L2VPN services using recursive VPN transport over BGP-LU, LDP, and RSVP-TE. All key address families, neighbors, tunneling, and label allocation options are established for robust service delivery.

What to do next

- Validate the configuration by checking VPN traffic flow and route propagation between nodes.
- Test failover scenarios and tunnel recovery for resilience.
- Monitor relevant interfaces for status or BFD or RSVP events.

Verify label recursion for L3VPN and L2VPN services on an aggregation router

Use this task to confirm recursive label resolution for VPN traffic on the aggregation (AGG) node.

This task is applicable when you need to ensure that label resolutions for VPN, BGP-LU, LDP, and RSVP-TE are functioning correctly for L3VPN and L2VPN services on the AGG router.

Follow these steps to verify label recursion for L3VPN and L2VPN services:

Before you begin

- Complete the feature configuration and ensure that the control plane has converged.
- Ensure that the control plane has converged.

Procedure

Step 1 Verify the VPN prefix recursion.

Example:

```
Router# show cef vrf vrf_ebgp_2000 99.0.0.1
99.0.0.0/28, version 141085, internal 0x5000001 0xf0 (ptr 0xaf490688) [1], 0x0 (0x0), 0x0 (0x0)
Prefix Len 28, traffic index 0, precedence n/a, priority 15, encap-id 0x5a288000003a1, Encap-Sharing
```

```
  gateway array (0xd7a8a2b8) reference count 2003, flags 0x20002010, source rib (7), 0 backups
                [1 type 3 flags 0x88401 (0xaf17a458) ext 0x0 (0x0)]
  LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]
  gateway array update type-time 1 Jun  3 08:21:31.494
 LDI Update time Jun  3 08:21:31.494
 LW-LDI-TS Jun  3 08:21:31.494
   via 4.4.4.1/32, 15 dependencies, recursive [flags 0x6000]
    path-idx 0 NHID 0x0 [0x3d377c48 0x0]
    next hop VRF - 'default', table - 0xe0000000
    next hop 4.4.4.1/32 via 50006/0/21
     next hop 1.1.1.1/32 tunnel-te5462 labels imposed {ImplNull 40710 30001}
     next hop 1.1.1.1/32 tunnel-te5463 labels imposed {ImplNull 40710 30001}
     next hop 1.1.1.1/32 tunnel-te5464 labels imposed {ImplNull 40710 30001}
     next hop 1.1.1.1/32 tunnel-te5465 labels imposed {ImplNull 40710 30001}
```

Confirm that the recursive path resolves through the BGP next hop and the TE tunnel. In the sample output, the imposed labels are {`ImplNull 40710 30001`}.

Here `30001` is VPN label, `40710` is BGP-LU label, `ImplNull` is for LDP label, and next hop is through the tunnel.

Step 2 Verify the VPN service label.

Example:

```
Router# show bgp vrf vrf_ebgp_2000 99.0.0.1
Mon Jun  8 12:01:08.754 UTC
BGP routing table entry for 99.0.0.0/28, Route Distinguisher: 3.3.3.1:0
Versions:
  Process           bRIB/RIB   SendTblVer
  Speaker            4534511      4534511
Last Modified: Jun  3 08:21:21.814 for 5d03h
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
  100, (Received from a RR-client)
    4.4.4.1 (metric 4000) from 4.4.4.1 (4.4.4.1)
      Received Label 30001
      Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate, imported
      Received Path ID 0, Local Path ID 1, version 4534511
      Extended community: RT:101:2000 RT:201:2000
      Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 4.4.4.1:0
```

Confirm that the prefix has the expected received VPN label. In the sample output, the received label is `30001`.

Step 3 Verify the BGP-LU label.

Example:

```
Router# show bgp ipv4 labeled-unicast 4.4.4.1
Mon Jun  8 12:02:35.438 UTC
BGP routing table entry for 4.4.4.1/32
Versions:
  Process           bRIB/RIB   SendTblVer
  Speaker                  6            6
    Local Label: 50006
Last Modified: Jun  3 08:19:12.814 for 5d03h
Paths: (2 available, best #1)
  Advertised IPv4 Labeled-unicast paths to update-groups (with more than one peer):
    0.3
  Path #1: Received by speaker 0
  Advertised IPv4 Labeled-unicast paths to update-groups (with more than one peer):
    0.3
  Local
```

```
    1.1.1.1 (metric 4000) from 1.1.1.1 (4.4.4.1)
      Received Label 40710
      Origin IGP, metric 0, localpref 100, valid, internal, best, group-best, labeled-unicast
      Received Path ID 1, Local Path ID 1, version 6
      Originator: 4.4.4.1, Cluster list: 0.0.11.184
  Path #2: Received by speaker 0
  Not advertised to any peer
  Local
    2.2.2.1 (metric 5000) from 2.2.2.1 (4.4.4.1)
      Received Label 16108
      Origin IGP, metric 0, localpref 100, valid, internal, backup, add-path, labeled-unicast
      Received Path ID 1, Local Path ID 2, version 6
      Originator: 4.4.4.1, Cluster list: 0.0.11.184
```

Confirm that the route has the expected received BGP-LU label and backup path information. In the sample output, the primary received label is `40710`.

Step 4 Verify route installation and PIC state.

Example:

```
RP/0/RP0/CPU0:AGG1#show route 4.4.4.1 detail
Mon Jun  8 11:57:55.533 UTC

Routing entry for 4.4.4.1/32
  Known via "bgp 1", distance 200, metric 0, [ei]-bgp
  Number of pic paths 1 , type internal
  Installed Jun  3 08:19:12.919 for 5d03h
  Routing Descriptor Blocks
    1.1.1.1, from 1.1.1.1
      Route metric is 0, Wt is 1
      Label: 0x9f06 (40710)
      Tunnel ID: None
      Binding Label: None
      Extended communities count: 0
      NHID: 0x0 (Ref: 0)
      Path Grouping ID: 1
    2.2.2.1, from 2.2.2.1, BGP backup path
      Route metric is 0, Wt is 1
      Label: 0x3eec (16108)
      Tunnel ID: None
      Binding Label: None
      Extended communities count: 0
      NHID: 0x0 (Ref: 0)
      Path Grouping ID: 1
  Route version is 0x4 (4)
  Local Label: 0xc356 (50006)
  IP Precedence: Not Set
  QoS Group ID: Not Set
  Flow-tag: Not Set
  Fwd-class: Not Set
  Route Priority: RIB_PRIORITY_RECURSIVE_SVD_LCL (14) SVD Type RIB_SVD_TYPE_LOCAL
  Download Priority 4, Download Version 1184278
```

Confirm that the route is installed with the expected primary and backup labels.

Step 5 Verify MPLS forwarding for the BGP-LU prefix.

Example:

```
Router# show mpls  forwarding  prefix 4.4.4.1/32 det
Mon Jun  8 12:04:06.208 UTC
Local  Outgoing    Prefix             Outgoing     Next Hop        Bytes
Label  Label       or ID              Interface                    Switched
------ ----------- ------------------ ------------ --------------- ------------
```

```
50006  40710       4.4.4.1/32                     1.1.1.1         0
     Updated: Jun  3 08:19:12.945
     Path Flags: 0x6000 [  ]
     Version: 1184278, Priority: 4
     Label Stack (Top -> Bottom): { 40710 }
     NHID: 0x0, Encap-ID: N/A, Path idx: 0, Backup path idx: 1, Weight: 1
     MAC/Encaps: 0/4, MTU: 0
     Packets Switched: 0

       16108       4.4.4.1/32                     2.2.2.1         0            (!)
     Updated: Jun  3 08:19:12.945
     Path Flags: 0x6100 [  BKUP, NoFwd ]
     Version: 1184278, Priority: 4
     Label Stack (Top -> Bottom): { 16108 }
     NHID: 0x0, Encap-ID: N/A, Path idx: 1, Backup path idx: 0, Weight: 1
     MAC/Encaps: 0/4, MTU: 0
     Packets Switched: 0
     (!): FRR pure backup
```

Confirm that the local label maps to the expected outgoing BGP-LU label and that the backup path is programmed.

Step 6 Verify LDP resolution for the transport next hop.

Example:

```
Router# show mpls forwarding  prefix 1.1.1.1/32 detail
Mon Jun  8 12:05:33.295 UTC
Local  Outgoing    Prefix             Outgoing     Next Hop        Bytes
Label  Label       or ID              Interface                    Switched
------ ----------- ------------------ ------------ --------------- ------------
71473  Pop         1.1.1.1/32         tt5462       1.1.1.1         0
     Updated: Jun  3 08:19:19.655
     Version: 4101844, Priority: 1
     Label Stack (Top -> Bottom): { 36576 Imp-Null }
     NHID: 0x0, Encap-ID: N/A, Path idx: 0, Backup path idx: 0, Weight: 0
     MAC/Encaps: 14/18, MTU: 1500
     Outgoing Interface: tunnel-te5462 (ifhandle 0x78004d7c)
     Packets Switched: 0

       Pop         1.1.1.1/32         tt5463       1.1.1.1         0
     Updated: Jun  3 08:19:19.655
     Version: 4101844, Priority: 1
     Label Stack (Top -> Bottom): { 36577 Imp-Null }
     NHID: 0x0, Encap-ID: N/A, Path idx: 1, Backup path idx: 0, Weight: 0
     MAC/Encaps: 14/18, MTU: 1500
     Outgoing Interface: tunnel-te5463 (ifhandle 0x78004d84)
     Packets Switched: 0

       Pop         1.1.1.1/32         tt5464       1.1.1.1         0
     Updated: Jun  3 08:19:19.655
     Version: 4101844, Priority: 1
     Label Stack (Top -> Bottom): { 36578 Imp-Null }
     NHID: 0x0, Encap-ID: N/A, Path idx: 2, Backup path idx: 0, Weight: 0
     MAC/Encaps: 14/18, MTU: 1500
     Outgoing Interface: tunnel-te5464 (ifhandle 0x78004d8c)
     Packets Switched: 0

       Pop         1.1.1.1/32         tt5465       1.1.1.1         0
     Updated: Jun  3 08:19:19.655
     Version: 4101844, Priority: 1
     Label Stack (Top -> Bottom): { 36579 Imp-Null }
     NHID: 0x0, Encap-ID: N/A, Path idx: 3, Backup path idx: 0, Weight: 0
     MAC/Encaps: 14/18, MTU: 1500
     Outgoing Interface: tunnel-te5465 (ifhandle 0x78004d94)
     Packets Switched: 0
```

Confirm that the LDP resolution uses the TE tunnel interfaces and implicit-null behavior.

Step 7 Verify the TE tunnel label.

Example:

```
Router# show mpls  forwarding  tunnels 5462 det
Mon Jun  8 12:10:11.266 UTC
Tunnel                    Outgoing    Outgoing     Next Hop        Bytes
Name                      Label       Interface                    Switched
------------------------- ----------- ------------ --------------- ------------
tunnel-te5462             36576       Hu0/0/0/28   12.23.245.2     0
     Updated: Jun  3 08:19:19.031
     Version: 4101799, Priority: 1
     Label Stack (Top -> Bottom): { 36576 }
     NHID: 0x0, Encap-ID: N/A, Path idx: 0, Backup path idx: 0, Weight: 0
     MAC/Encaps: 14/18, MTU: 1500
     Outgoing Interface: HundredGigE0/0/0/28 (ifhandle 0x78000160)
     Packets Switched: 0

  Interface:
    Name: tunnel-te5462 (ifhandle 0x78004d7c)
    Local Label: 71079, Forwarding Class: 0, Weight: 0
    Packets/Bytes Switched: 14650529304/2766291687736
```

Confirm that the tunnel has the expected outgoing label. In the sample output, the outgoing tunnel label is `36576`.

Step 8 Verify TE tunnel state.

Example:

```
Router# show mpls  traffic-eng  tunnels 5462 detail
Mon Jun  8 12:13:01.559 UTC

Name: tunnel-te5462  Destination: 1.1.1.1  Ifhandle:0x78004d7c
  Signalled-Name: P1-AGG-Saturn_t5462
  Status:
    Admin:    up Oper:   up (Uptime 5d03h)

    path option 1,  type explicit pri-tunnel-5462 (Basis for Setup, path weight 6000)
      Accumulative metrics: TE 6000 IGP 6000 Delay 600000
      Accumulative biased metrics: TE 6000 IGP 6000 Delay 600000
    G-PID: 0x0800 (derived from egress interface properties)
    Bandwidth Requested: 0 kbps  CT0
    Creation Time: Thu May 28 20:10:42 2026 (1w3d ago)
  Config Parameters:
    Source: router ID (default)
    Bandwidth:        0 kbps (CT0) Priority:  7  7 Affinity: 0x0/0xffff
    Metric Type: TE (global)
    Path Selection:
      Tiebreaker: Min-fill (default)
    Hop-limit: disabled
    Cost-limit: disabled
    Delay-limit: disabled
    Delay-measurement: disabled
    Path-invalidation timeout: 10000 msec (default), Action: Tear (default)
    AutoRoute:     IPv4  LockDown: disabled   Policy class: not set
    Forward class: 0 (not enabled)
    Forwarding-Adjacency: disabled
    Autoroute Destinations: 0
    Loadshare:          0 equal loadshares
    Load-interval: 300 seconds
    Auto-bw: disabled
    Auto-Capacity: Disabled:
    Self-ping: Disabled
```

```
  Fast Reroute: Enabled, Protection Desired: Any
  Path Protection: Enabled
  BFD Fast Detection: Disabled
  Reoptimization after affinity failure: Enabled
  Soft Preemption: Disabled
SNMP Index: 2590
Binding SID: None
Path Protection Info:
  No valid path-option for Path Protection
  Number of Switchovers 0, Standby Ready 0 times, Standby Reopt 0 times
  LSP Wrap Protection: Disabled
History:
  Reopt. LSP:
    Last Failure:
      LSP not signalled, identical to the [CURRENT] LSP
      Date/Time: Mon Jun 08 12:10:16 UTC 2026 [00:02:45 ago]
  Prior LSP:
    ID: 2 Path Option: 1
    Removal Trigger: resv delete
Current LSP Info:
  Instance: 4, Signaling Area: IS-IS core level-2
  In-use path-option: 1
  Uptime: 5d03h (since Wed Jun 03 08:19:18 UTC 2026)
  Outgoing Interface: HundredGigE0/0/0/28, Outgoing Label: 36576
```

Confirm that the tunnel is operational and that FRR and path protection are enabled.

Successful verification confirms that VPN traffic resolves through the VPN service label, the BGP-LU label, the LDP transport label, and the RSVP-TE tunnel label, ensuring end-to-end label recursion on the AGG router.

Troubleshooting checks and commands for L3VPN and L2VPN over BGP-LU over LDP over RSVP-TE

Use these checks when you troubleshoot the feature:

- Verify VPN label allocation Confirm that the VPN route has the expected received label.
- Verify BGP-LU route installation: Confirm that the BGP-LU next hop is installed with primary and backup labels.
- Verify LDP label distribution: Confirm that the transport prefix has the expected LDP behavior.
- Verify RSVP-TE tunnel state: Confirm that the TE tunnel is operationally up and programmed in forwarding.
- Verify recursive CEF resolution: Confirm that the VPN route resolves through the expected next hop and label stack.
- Verify PIC backup path installation: Confirm that the backup path is present in routing and forwarding output.

Supported commands

Use the following commands to verify and troubleshoot L3VPN and L2VPN services operating with BGP-LU over LDP and RSVP-TE protocols:

- **show route detail**
- **show cef detail**
- **show bgp ipv4 labeled-unicast**
- **show bgp vpnv4 unicast**
- **show mpls forwarding**
- **show mpls ldp bindings**
- **show mpls traffic-eng tunnels detail**