



Configuring MPLS Layer 3 VPN Label Allocation

This chapter describes how to configure label allocation for Multiprotocol Label Switching (MPLS) Layer 3 virtual private networks (L3VPNs) on Cisco NX-OS devices.

This chapter includes the following sections:

- [Finding Feature Information, page 23-412](#)
- [Information About MPLS L3VPN Label Allocation, page 23-412](#)
- [Licensing Requirements for MPLS L3VPN Label Allocation, page 23-414](#)
- [Prerequisites for MPLS L3VPN Label Allocation, page 23-414](#)
- [Guidelines and Limitations for MPLS L3VPN Label Allocation, page 23-414](#)
- [Default Settings for MPLS L3VPN Label Allocation, page 23-415](#)
- [Configuring MPLS L3VPN Label Allocation, page 23-415](#)
- [Verifying MPLS L3VPN Label Allocation Configuration, page 23-420](#)
- [Configuration Examples for MPLS L3VPN Label Allocation, page 23-420](#)
- [Additional References for MPLS L3VPN Label Allocation, page 23-421](#)
- [Feature History for MPLS L3VPN Label Allocation, page 23-422](#)

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your 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 “New and Changed Information” chapter or the Feature History table below.

Information About MPLS L3VPN Label Allocation

The MPLS provider edge (PE) router stores both local and remote routes and includes a label entry for each route. By default, Cisco NX-OS uses per-prefix label allocation which means that each prefix is assigned a label. For distributed platforms, the per-prefix labels consume memory. When there are many VPN routing and forwarding instances (VRFs) and routes, the amount of memory that the per-prefix labels consume can become an issue.

You can enable per-VRF label allocation to advertise a single VPN label for local routes throughout the entire VRF. The router uses a new VPN label for the VRF decoding and IP-based lookup to learn where to forward packets for the PE or customer edge (CE) interfaces.

You can enable different label allocation modes for Border Gateway Protocol (BGP) Layer 3 VPN routes to meet different requirements and to achieve trade-offs between scalability and performance. All labels are allocated within the global label space. Cisco NX-OS supports the following label allocation modes:

- **Per-prefix**—A label is allocated for each VPN prefix. VPN packets received from remote PEs can be directly forwarded to the connected CE that advertised the prefix, based on the label forwarding table. However, this mode also uses many labels. This mode is the only mode available when VPN packets sent from PE to CE are label switched. This is the default label allocation mode.
- **Per-VRF**—A single label is assigned to all local VPN routes in a VRF. This mode requires an IPv4 or IPv6 lookup in the VRF forwarding table once the VPN label is removed at the egress PE. This mode is the most efficient in terms of label space as well as BGP advertisements, and the lookup does not result in any performance degradation. Cisco NX-OS uses the same per-VRF label for both IPv4 and IPv6 prefixes.



Note EIBGP load balancing is not supported for a VRF that uses per-VRF label mode.

- **Aggregate Labels**—BGP can allocate and advertise a local label for an aggregate prefix. Forwarding requires an IPv4 or IPv6 lookup that is similar to the per-VRF scenario. A single per-VRF label is allocated and used for all prefixes that need a lookup.
- **VRF connected routes**—When directly connected routes are redistributed and exported, an aggregate label is allocated for each route. The packets that come in from the core are decapsulated and a lookup is done in the VRF IPv4 or IPv6 table to determine whether the packet is for the local router or for another router or host that is directly connected. A single per-VRF label is allocated for all such routes.
- **Label hold down**—When a local label is no longer associated with a prefix, to allow time for updates to be sent to other PEs, the local label is not released immediately. A ten minute hold down timer is started per label. Within this hold down period, the label can be reclaimed for the prefix. When the timer expires, BGP releases the label.

Per-VRF Label Allocation Mode

The following conditions apply when you configure per-VRF label allocation:

- The VRF uses one label for all local routes.
- When you enable per-VRF label allocation, any existing per-VRF aggregate label is used. If no per-VRF aggregate label is present, the software creates a new per-VRF label.

The CE does not lose data when you disable per-VRF label allocation because the configuration reverts to the default per-prefix labeling configuration.

- A per-VRF label forwarding entry is deleted only if the VRF, BGP, or address family configuration is removed.

IPv6 Label Allocation

IPv6 prefixes are advertised with the allocated label to iBGP peers that have the labeled-unicast address-family enabled. The received eBGP next hop is not propagated to such peers; instead, the local IPv4 session address is sent as an IPv4-mapped IPv6 next hop. The remote peer resolves this next hop through one or more IPv4 MPLS LSPs in the core network.

You can use a route reflector to advertise the labeled 6PE prefixes between PEs. You must enable the labeled-unicast address-family between the route reflector and all such peers. The route reflector does not need to be in the forwarding path and propagates the received next hop as is to iBGP peers and route reflector clients.



Note

6PE also supports both per-prefix and per-VRF label allocation modes, as in 6VPE.

Licensing Requirements for MPLS L3VPN Label Allocation

Product	License Requirement
Cisco NX-OS	L3VPN label allocation requires an MPLS license. For a complete explanation of the Cisco NX-OS licensing scheme and how to obtain and apply licenses, see the <i>Cisco NX-OS Licensing Guide</i> .

Prerequisites for MPLS L3VPN Label Allocation

L3VPN label allocation has the following prerequisites:

- Ensure that you have configured MPLS, and LDP or RSVP TE in your network. All routers in the core, including the PE routers, must be able to support MPLS forwarding.
- Ensure that you have installed the correct license for MPLS and any other features you will be using with MPLS.
- Ensure that you disable the external/internal Border Gateway Protocol (BGP) multipath feature if it is enabled before you configure per-VRF label allocation mode.
- Before configuring a 6VPE per VRF label, ensure that the IPv6 address family is configured on that VRF.

Guidelines and Limitations for MPLS L3VPN Label Allocation

L3VPN label allocation has the following configuration guidelines and limitations:

- F Series modules do not natively support label switching. They can leverage M Series modules for label switching using proxy forwarding. For more information on proxy forwarding, see the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide*.
- Enabling per-VRF label allocation causes BGP reconvergence, which can result in data loss for traffic coming from the MPLS VPN core.



Note You can minimize network disruption by enabling per-VRF label allocation during a scheduled MPLS maintenance window. Also, if possible, avoid enabling this feature on a live router.

- Per-prefix MPLS counters for VPN prefixes are lost when you enable per-VRF label allocation.
- Aggregate labels and per-VRF labels are global across all virtual device contexts (VDCs) and are in a separate, dedicated label range.
- Aggregate prefixes for per-prefix label allocation share the same label in a given VRF.

Default Settings for MPLS L3VPN Label Allocation

Table 23-1 lists the default settings for L3VPN label allocation parameters.

Table 23-1 Default L3VPN Label Allocation Parameters

Parameters	Default
L3VPN feature	Disabled
Label allocation mode	Per prefix

Configuring MPLS L3VPN Label Allocation

This section includes the following topics:

- [Configuring Per-VRF L3VPN Label Allocation Mode, page 23-415](#)
- [Allocating Labels for IPv6 Prefixes in the Default VRF, page 23-417](#)
- [Enabling Sending MPLS Labels in IPv6 over an IPv4 MPLS Core Network \(6PE\) for iBGP Neighbors, page 23-418](#)

Configuring Per-VRF L3VPN Label Allocation Mode

You can configure per-VRF L3VPN label allocation mode for Layer 3 VPNs.

Prerequisites

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

SUMMARY STEPS

1. **configure terminal**
2. **feature bgp**
3. **feature-set mpls**
4. **feature mpls l3vpn**
5. **router bgp as-number**

6. `vrf vrf-name`
7. `address-family {ipv6 | ipv4} {unicast | multicast}`
8. `label-allocation-mode per-vrf`
9. (Optional) `show bgp l3vpn detail vrf vrf-name`
10. (Optional) `copy running-config startup-config`

DETAILED STEPS

	Command	Purpose
Step 1	<code>configure terminal</code> Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	<code>feature bgp</code> Example: switch(config)# feature bgp	Enables the BGP feature.
Step 3	<code>feature-set mpls</code> Example: switch(config)# feature-set mpls	Enables the MPLS feature-set.
Step 4	<code>feature mpls l3vpn</code> Example: switch(config)# feature mpls l3vpn	Enables the MPLS Layer 3 VPN feature.
Step 5	<code>router bgp as-number</code> Example: switch(config)# router bgp 1.1 switch(config-router)#	Configures a BGP routing process and enters router configuration mode. The <i>as-number</i> argument indicates the number of an autonomous system that identifies the router to other BGP routers and tags the routing information. The AS number can be a 16-bit integer or a 32-bit integer in the form of a higher 16-bit decimal number and a lower 16-bit decimal number in xx.xx format.
Step 6	<code>vrf vrf-name</code> Example: switch(config-router)# vrf vpn1 switch(config-router-vrf)#	Enters router VRF configuration mode. The <i>vrf-name</i> can be any case-sensitive, alphanumeric string up to 32 characters.
Step 7	<code>address-family {ipv4 ipv6} unicast multicast}</code> Example: switch(config-router-vrf)# address-family ipv6 unicast switch(config-router-vrf-af)#	Specifies the IP address family type and enters address family configuration mode.
Step 8	<code>label-allocation-mode per-vrf</code> Example: switch(config-router-vrf-af)# label-allocation-mode per-vrf	Allocates labels on a per-VRF basis.

	Command	Purpose
Step 9	<pre>show bgp l3vpn detail vrf vrf-name</pre> <p>Example: switch(config-router-vrf-af)# show bgp l3vpn detail vrf vpn1</p>	(Optional) Displays information about Layer 3 VPN configuration on BGP for this VRF. The <i>vrf-name</i> can be any case-sensitive, alphanumeric string up to 32 characters.
Step 10	<pre>copy running-config startup-config</pre> <p>Example: switch(config-router-vrf-af)# copy running-config startup-config</p>	(Optional) Copies the running configuration to the startup configuration.

Allocating Labels for IPv6 Prefixes in the Default VRF

If you are running IPv6 over an IPv4 MPLS core network (6PE), you can allocate labels for the IPv6 prefixes in the default VRF.



Note

By default, labels are not allocated for IPv6 prefixes in the default VRF.

Prerequisites

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

SUMMARY STEPS

1. **configure terminal**
2. **feature bgp**
3. **feature-set mpls**
4. **feature mpls l3vpn**
5. **router bgp as-number**
6. **address-family ipv6 {unicast | multicast}**
7. **allocate-label {all | route-map route-map}**
8. (Optional) **show running-config bgp**
9. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	<pre>configure terminal</pre> <p>Example: switch# configure terminal switch(config)#</p>	Enters global configuration mode.
Step 2	<pre>feature bgp</pre> <p>Example: switch(config)# feature bgp</p>	Enables the BGP feature.

	Command	Purpose
Step 3	feature-set mpls Example: switch(config)# feature-set mpls	Enables the MPLS feature-set.
Step 4	feature mpls l3vpn Example: switch(config)# feature mpls l3vpn	Enables the MPLS Layer 3 VPN feature.
Step 5	router bgp as-number Example: switch(config)# router bgp 1.1 switch(config-router)#	Configures a BGP routing process and enters router configuration mode. The <i>as-number</i> argument indicates the number of an autonomous system that identifies the router to other BGP routers and tags the routing information. The AS number can be a 16-bit integer or a 32-bit integer in the form of a higher 16-bit decimal number and a lower 16-bit decimal number in xx.xx format.
Step 6	address-family ipv6 {unicast multicast} Example: switch(config-router)# address-family ipv6 unicast switch(config-router-af)#	Specifies the IPv6 address family type and enters address family configuration mode.
Step 7	allocate-label {all route-map route-map} Example: switch(config-router-af)# allocate-label all	Allocates labels for IPv6 prefixes in the default VRF. <ul style="list-style-type: none"> The all keyword allocates labels for all IPv6 prefixes. The route-map keyword allocates labels for IPv6 prefixes matched in the specified route map. The route-map can be any case-sensitive alphanumeric string up to 63 characters.
Step 8	show running-config bgp Example: switch(config-router-af)# show running-config bgp	(Optional) Displays information about the BGP configuration.
Step 9	copy running-config startup-config Example: switch(config-router-af)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Enabling Sending MPLS Labels in IPv6 over an IPv4 MPLS Core Network (6PE) for iBGP Neighbors

You can enable sending MPLS labels to iBGP neighbors.



Note

The **address-family ipv6 labeled-unicast** command is supported only for iBGP neighbors. You cannot use this command with the **address-family ipv6 unicast** command.

Prerequisites

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

SUMMARY STEPS

1. **configure terminal**
2. **feature bgp**
3. **feature-set mpls**
4. **feature mpls l3vpn**
5. **router bgp *as-number***
6. **neighbor *ip-address***
7. **address-family ipv6 labeled-unicast**
8. (Optional) **show running-config bgp**
9. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	feature bgp Example: switch(config)# feature bgp	Enables the BGP feature.
Step 3	feature-set mpls Example: switch(config)# feature-set mpls	Enables the MPLS feature-set.
Step 4	feature mpls l3vpn Example: switch(config)# feature mpls l3vpn	Enables the MPLS Layer 3 VPN feature.
Step 5	router bgp <i>as-number</i> Example: switch(config)# router bgp 1.1 switch(config-router)#	Configures a BGP routing process and enters router configuration mode. The <i>as-number</i> argument indicates the number of an autonomous system that identifies the router to other BGP routers and tags the routing information. The AS number can be a 16-bit integer or a 32-bit integer in the form of a higher 16-bit decimal number and a lower 16-bit decimal number in xx.xx format.
Step 6	neighbor <i>ip-address</i> Example: switch(config-router)# neighbor 209.165.201.1 switch(config-router-neighbor)#	Adds an entry to the BGP or multiprotocol BGP neighbor table. The <i>ip-address</i> argument specifies the IP address of the neighbor in dotted decimal notation.

	Command	Purpose
Step 7	address-family ipv6 labeled-unicast Example: switch(config-router-neighbor)# address-family ipv6 labeled-unicast switch(config-router-neighbor-af)#	Specifies IPv6 labeled unicast address prefixes. This command is accepted only for iBGP neighbors.
Step 8	show running-config bgp Example: switch(config-router-neighbor-af)# show running-config bgp	(Optional) Displays information about the BGP configuration.
Step 9	copy running-config startup-config Example: switch(config-router-neighbor-af)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Verifying MPLS L3VPN Label Allocation Configuration

To display the L3VPN label allocation configuration, perform one of the following tasks:


Command	Purpose
show bgp l3vpn [detail] [vrf vrf-name]	Displays Layer 3 VPN information for BGP in a VRF.
show bgp vpnv4 unicast labels [vrf vrf-name]	Displays label information for BGP.
show ip route [vrf vrf-name]	Displays label information for routes.

For detailed information about the fields in the output from these commands, see the *Cisco NX-OS MPLS Command Reference*.

Configuration Examples for MPLS L3VPN Label Allocation

This section uses the following sample MPLS network shown in [Figure 23-1](#).

Figure 23-1 Sample MPLS Layer3 Network



The following example shows how to configure per-VRF label allocation for an IPv4 MPLS network.

PE Configuration

```

PE1
-----
vrf context vpn1
  rd 100:1
  address-family ipv4 unicast
    route-target export 200:1
router bgp 100
  neighbor 10.1.1.2 remote-as 100
  address-family vpnv4 unicast
    send-community extended
  update-source loopback10
vrf vpn1
  address-family ipv4 unicast
    label-allocation-mode per-vrf
  neighbor 36.0.0.2 remote-as 300
  address-family ipv4 unicast

```

Additional References for MPLS L3VPN Label Allocation

For additional information related to implementing L3VPN Label Allocation, see the following sections:

- [Related Documents, page 23-422](#)
- [MIBs, page 23-422](#)

Related Documents

Related Topic	Document Title
CLI commands	<i>Cisco Nexus 7000 Series NX-OS MPLS Command Reference</i>

MIBs

MIBs	MIBs Link
MPLS-L3VPN-STD-MIB	To locate and download Cisco MIBs, go to the following URL: http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

Feature History for MPLS L3VPN Label Allocation

Table 23-2 lists the release history for this feature.

Table 23-2 Feature History for L3VPN Label Allocation

Feature Name	Releases	Feature Information
Per-VRF label allocation	5.2(1)	This feature was introduced.

