



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 Nexus 3600 series Switches.

- [Information About MPLS L3VPN Label Allocation, on page 1](#)
- [Prerequisites for MPLS L3VPN Label Allocation, on page 3](#)
- [Guidelines and Limitations for MPLS L3VPN Label Allocation, on page 3](#)
- [Default Settings for MPLS L3VPN Label Allocation, on page 4](#)
- [Configuring MPLS L3VPN Label Allocation, on page 4](#)
- [Advertisement and Withdraw Rules, on page 9](#)
- [Enabling Local Label Allocation, on page 11](#)
- [Verifying MPLS L3VPN Label Allocation Configuration, on page 13](#)
- [Configuration Examples for MPLS L3VPN Label Allocation, on page 14](#)

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.

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

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.

About Labeled and Unlabeled Unicast Paths

Subsequent Address Family Identifier (SAFI) is an indication of the BGP route. Example 1 is for an unlabeled route and 4 for a labeled route.

- Unlabeled unicast (U) for IPv4 is SAFI 1.
- Labeled unicast (LU) for IPv4 is SAFI 4.
- Unlabeled unicast (U) for IPv6 is AFI 2 and SAFI 1.
- Labeled unicast (LU) for IPv6 is AFI 2 and SAFI 4.

Cisco NX-OS Release 9.2(2) supports both, IPv4 and IPv6 unlabeled and labeled unicast on one BGP session. This behavior is the same irrespective of whether one or both SAFI-1 and SAFI-4 are enabled on the same session or not.

This behavior is applicable for all eBGP, iBGP, and redistributed paths and the eBGP and iBGP neighbors.

Prerequisites for MPLS L3VPN Label Allocation

L3VPN label allocation has the following prerequisites:

- Ensure that you have configured MPLS, and LDP 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:

- Layer 3 VPN label allocation is also supported on the Cisco Nexus 3600 platform switches.
- 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.

- 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 1: Default L3VPN Label Allocation Parameters

Parameters	Default
L3VPN feature	Disabled
Label allocation mode	Per prefix

Configuring MPLS L3VPN Label Allocation

Configuring Per-VRF L3VPN Label Allocation Mode

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

SUMMARY STEPS

1. **configure terminal**
2. **feature bgp**
3. **feature-set mpls**
4. **feature-set mpls l3vpn**
5. **router bgp** *as - number*
6. **vrf** *vrf-name*
7. **address-family** { *ipv4* | *ipv6* } **unicast** | **multicast** }
8. **label-allocation-mode per-vrf**
9. **show bgp l3vpn detail vrf** *vrf-name*
10. **copy running-config startup-config**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	feature bgp Example:	Enables the BGP feature.

	Command or Action	Purpose
	<pre>switch(config)# feature bgp switch(config)#</pre>	
Step 3	feature-set mpls Example: <pre>switch(config)# feature-set mpls switch(config)#</pre>	Enables the MPLS feature-set.
Step 4	feature-set mpls l3vpn Example: <pre>switch(config)# feature-set mpls l3vpn switch(config)#</pre>	Enables the MPLS Layer 3 VPN feature.
Step 5	router bgp as - number Example: <pre>switch(config)# router bgp 1.1</pre>	Configures a BGP routing process and enters router configuration mode. The as-number 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	vrf vrf-name Example: <pre>switch(config-router)# vrf vpn1</pre>	Enters router VRF configuration mode. The vrf-name can be any case-sensitive, alphanumeric string up to 32 characters..
Step 7	address-family { ipv4 ipv6 } unicast multicast } Example: <pre>switch(config-router-vrf)# address-family ipv6 unicast</pre>	Specifies the IP address family type and enters address family configuration mode.
Step 8	label-allocation-mode per-vrf Example: <pre>switch(config-router-vrf-af)# label-allocation-mode per-vrf</pre>	Allocates labels on a per-VRF basis.
Step 9	show bgp l3vpn detail vrf vrf-name Example: <pre>switch(config-router-vrf-af)# show bgp l3vpn detail vrf vpn1</pre>	(Optional) Displays information about Layer 3 VPN configuration on BGP for this VRF. The vrf-name can be any case-sensitive, alphanumeric string up to 32 characters.
Step 10	copy running-config startup-config Example: <pre>switch(config-router-vrf)# copy running-config startup-config</pre>	(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.

SUMMARY STEPS

1. **configure terminal**
2. **feature bgp**
3. **feature-set mpls**
4. **feature-set mpls l3vpn**
5. **router bgp** *as - number*
6. **address-family { ipv4 | ipv6 } unicast | multicast }**
7. **allocate-label { all | route-map route-map }**
8. **show running-config bgp**
9. **copy running-config startup-config**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	feature bgp Example: <pre>switch(config)# feature bgp switch(config)#</pre>	Enables the BGP feature.
Step 3	feature-set mpls Example: <pre>switch(config)# feature-set mpls switch(config)#</pre>	Enables the MPLS feature-set.
Step 4	feature-set mpls l3vpn Example: <pre>switch(config)# feature-set mpls l3vpn switch(config)#</pre>	Enables the MPLS Layer 3 VPN feature.

	Command or Action	Purpose
Step 5	router bgp <i>as - number</i> Example: <pre>switch(config)# router bgp 1.1</pre>	Configures a BGP routing process and enters router configuration mode. The as-number 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 { ipv4 ipv6 } unicast multicast } Example: <pre>switch(config-router-vrf)# address-family ipv6 unicast</pre>	Specifies the IP address family type and enters address family configuration mode.
Step 7	allocate-label { all route-map route-map } Example: <pre>switch(config-router-af)# allocate-label all</pre>	Allocates labels for IPv6 prefixes in the default VRF. • 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: <pre>switch(config-router-af)# show running-config bgp</pre>	(Optional) Displays information about the BGP configuration.
Step 9	copy running-config startup-config Example: <pre>switch(config-router-vrf)# copy running-config startup-config</pre>	(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

6PE advertises IPv6 prefixes in global VRF over IPv4 based MPLS network with the allocated label to iBGP peers that have the labeled-unicast address-family enabled. 6PE requires LDP enabled on core facing interfaces to transport IPv6 traffic over IPv4 based MPLS network and “address-family ipv6 labeled-unicast” under BGP to exchange label for IPv6 prefixes between PEs.



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.

SUMMARY STEPS

1. configure terminal

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

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	feature bgp Example: <pre>switch(config)# feature bgp switch(config)#</pre>	Enables the BGP feature.
Step 3	feature-set mpls Example: <pre>switch(config)# feature-set mpls switch(config)#</pre>	Enables the MPLS feature-set.
Step 4	feature-set mpls l3vpn Example: <pre>switch(config)# feature-set mpls l3vpn switch(config)#</pre>	Enables the MPLS Layer 3 VPN feature.
Step 5	router bgp <i>as - number</i> Example: <pre>switch(config)# router bgp 1.1</pre>	Configures a BGP routing process and enters router configuration mode. The as-number 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: <pre>switch(config-router)# neighbor 209.165.201.1 switch(config-router-neighbor)#</pre>	Adds an entry to the BGP or multiprotocol BGP neighbor table. The ip-address argument specifies the IP address of the neighbor in dotted decimal notation.

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

Example

What to do next

•

Advertisement and Withdraw Rules

The following table shows the advertisement and withdraw behavior for different scenarios.

Table 2: Advertisement and Withdraw Rules

Case	Bestpath/ Addpath Type	Local Label Present?	NHS or NHU	Update-group SAFI	Advertise or withdraw?
1	Unlabeled path. For example, no RX label.	Yes	NHS	SAFI-1	Advertise by default.
2				SAFI-4	Advertise
3			NHU	SAFI-1	Advertise
4				SAFI-4	Withdraw
5		No	NHS	SAFI-1	Advertise
6				SAFI-4	Withdraw
7			NHU	SAFI-1	Advertise
8				SAFI-4	Withdraw
9	Labeled path. For example, with an RX label.	Yes	NHS	SAFI-1	Advertise by default. Withdraw with NbrKnob.
10				SAFI-4	Advertise

Case	Bestpath/ Addpath Type	Local Label Present?	NHS or NHU	Update-group SAFI	Advertise or withdraw?
11			NHU	SAFI-1	Withdraw
12				SAFI-4	Advertise
13		No	NHS	SAFI-1	Advertise
14				SAFI-4	Withdraw
15			NHU	SAFI-1	Withdraw
				SAFI-4	Advertise

Enabling Local Label Allocation

SUMMARY STEPS

1. **configure terminal**
2. **feature bgp**
3. **feature-set mpls**
4. **router bgp** *as - number*
5. **address-family { ipv4 | ipv6 } unicast | multicast }**
6. **allocate-label { all | route-map *route-map* }**
7. **neighbor *ip-address***
8. **[no] advertise local-labeled-route**
9. **address-family { ipv4 | ipv6 } unicast | multicast }**
10. **[no] advertise local-labeled-route**
11. **route-map *label_routemap* permit 10**
12. **show running-config bgp**

13. copy running-config startup-config**DETAILED STEPS****Procedure**

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	feature bgp Example: <pre>switch(config)# feature bgp switch(config)#</pre>	Enables the BGP feature.
Step 3	feature-set mpls Example: <pre>switch(config)# feature-set mpls switch(config)#</pre>	Enables the MPLS feature-set.
Step 4	router bgp as - number Example: <pre>switch(config)# router bgp 1.1</pre>	Configures a BGP routing process and enters router configuration mode. The as-number 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 5	address-family { ipv4 ipv6 } unicast multicast } Example: <pre>switch(config-router-vrf)# address-family ipv4 unicast</pre>	Specifies the IP address family type and enters the address family configuration mode.
Step 6	allocate-label { all route-map route-map } Example: <pre>switch(config-router-af)# allocate-label all</pre>	Allocates labels for IPv6 prefixes in the default VRF. • 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.

	Command or Action	Purpose
Step 7	neighbor ip-address Example: <pre>switch(config-router)# neighbor 209.165.201.1 switch(config-router-neighbor)#</pre>	Adds an entry to the BGP or multiprotocol BGP neighbor table. The ip-address argument specifies the IP address of the neighbor in dotted decimal notation.
Step 8	[no] advertise local-labeled-route Example: <pre>switch(config-router-neighbor)# advertise local-labeled-route</pre>	Indicates whether to advertise an IPv4 or IPv6 route with a local label to the BGP neighbor via the IPv4 or IPv6 unicast SAFI (SAFI-1). The default is enabled so that it can be advertised to the BGP neighbor.
Step 9	address-family { ipv4 ipv6 } unicast multicast } Example: <pre>switch(config-router-vrf)# address-family ipv6 unicast</pre>	Specifies the IP address family type and enters the address family configuration mode.
Step 10	[no] advertise local-labeled-route Example: <pre>switch(config-router-neighbor)# advertise local-labeled-route</pre>	Indicates whether to advertise an IPv4 or IPv6 route with a local label to the BGP neighbor via the IPv4 or IPv6 unicast SAFI (SAFI-1). The default is enabled so that it can be advertised to the BGP neighbor.
Step 11	route-map label_routemap permit 10 Example: <pre>switch(config-router-vrf)# route-map label_routemap permit 10</pre>	
Step 12	show running-config bgp Example: <pre>switch(config-router-af)# show running-config bgp</pre>	(Optional) Displays information about the BGP configuration.
Step 13	copy running-config startup-config Example: <pre>switch(config-router-vrf)# copy running-config startup-config</pre>	(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:

Table 3: Verifying MPLS L3VPN Label Allocation Configuration

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.

Command	Purpose
<code>show ip route [vrf vrf-name]</code>	Displays label information for routes.

Configuration Examples for MPLS L3VPN Label Allocation

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

```
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
```