

Labeled and Unlabeled Together on one BGP Neighbor on IOS-XR

Contents

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Introduction](#)

[Capability Exchange](#)

[IOS-XR Implementation Specifics](#)

[Configuring U or LU](#)

[BGP Table](#)

[Enabling the Local Label Allocation](#)

[New Implementation on IOS-XR](#)

[Update-groups](#)

[Route Propagation](#)

[Allocating a Local Label](#)

[Allocate Label for Unlabeled Paths](#)

[Bestpath Calculation](#)

[Verifying the Behavior in BGP](#)

[iBGP and next-hop-self – Unified MPLS](#)

[Workarounds](#)

[Example configuration for IOS-XR.](#)

This document describes the behavior of receiving and advertising unlabeled and labeled paths over one BGP session in Cisco IOS[®] XR.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is specific to Cisco IOS[®] XR, but it is not restricted to a specific software release or hardware.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Introduction

Address Family Identifier (AFI) is an indication of the kind of BGP route. Examples are 1 for IPv4 and 2 for IPv6.

Subsequent Address Family Identifier (SAFI) is a further indication of the kind of route. Examples are 1 for an unlabeled route and 4 for a labeled route.

Unlabeled unicast for IPv4 is AFI 1 and SAFI 1.

Labeled unicast for IPv4 is AFI 1 and SAFI 4.

Unlabeled unicast for IPv6 is AFI 2 and SAFI 1.

Labeled unicast for IPv6 is AFI 2 and SAFI 4.

Labeled unicast (LU) is often referred to as RFC 3107 "Carrying Label Information in BGP-4."

Hereafter, U refers to unlabeled unicast, so SAFI 1 and LU refers to labeled unicast, so SAFI 4.

Note that Cisco IOS[®] XR needs "allocate-label all| route-policy ..." or the route will not be originated or propagated to the next BGP speaker as SAFI 4.

IPv4/v6 unicast and labeled-unicast both on one BGP session was not supported on Cisco IOS[®] XR.

On Cisco IOS[®] XR, the support to have both unlabeled and labeled on one BGP session was done in 6.2.1.

When running both on one session is not supported, it is problematic as the last received update/withdraw overrides the previous one, even if they were received on another SAFI.

When you configure both SAFI 1 and 4 on the same BGP session on a router running IOS-XR code prior to IOS-XR 6.2.1, the router gives the following warning:

```
bgp[1051]: %ROUTING-BGP-4-INCOMPATIBLE_AFI : IPv4 Unicast and IPv4 Labeled-unicast Address families together are not supported under the same neighbor.
```

This warning message was introduced in IOS-XR 5.3.0 and IOS-XR 5.2.2.

Capability Exchange

The capabilities exchanged between BGP peers must match. If not, the BGP session does not come up.

This is a wireshark capture of the capability exchanged for AFI 1/SAFI 1 and AFI 1/SAFI 4 in the BGP Open message:

```

Marker: ffffffffffffffffffffffffffffffffff
Length: 95
Type: OPEN Message (1)
Version: 4
My AS: 65003
Hold Time: 180
BGP Identifier: 10.100.1.4
Optional Parameters Length: 66
v Optional Parameters
  v Optional Parameter: Capability
    Parameter Type: Capability (2)
    Parameter Length: 6
    v Capability: Multiprotocol extensions capability
      Type: Multiprotocol extensions capability (1)
      Length: 4
      AFI: IPv4 (1)
      Reserved: 00
      SAFI: Unicast (1)
  v Optional Parameter: Capability
    Parameter Type: Capability (2)
    Parameter Length: 6
    v Capability: Multiprotocol extensions capability
      Type: Multiprotocol extensions capability (1)
      Length: 4
      AFI: IPv4 (1)
      Reserved: 00
      SAFI: Labeled Unicast (4)
  > Optional Parameter: Capability
  > Optional Parameter: Capability

```

Image 1

Here's an example for IOS-XR configured with LU only on a session to IOS configured with U only.

IOS-XR:

```

RP/0/0/CPU0:R4#show bgp neighbor 10.100.1.8

BGP neighbor is 10.100.1.8
  Remote AS 65003, local AS 65003, internal link
  Remote router ID 0.0.0.0
  BGP state = Idle
...

Connections established 0; dropped 0
Local host: 10.100.1.4, Local port: 179, IF Handle: 0x00000000
Foreign host: 10.100.1.8, Foreign port: 33396
Last reset 00:00:14, due to BGP Notification sent: unsupported/disjoint capability
Time since last notification sent to neighbor: 00:00:14
Error Code: unsupported/disjoint capability
Notification data sent:

```

None

The IOS router prints a syslog message for this configuration mistake:

```
*Aug  8 12:40:44.719: %BGP-3-NOTIFICATION: received from neighbor 10.100.1.4 active 2/7  
(unsupported/disjoint capability) 0 by
```

IOS-XR Implementation Specifics

Configuring U or LU

You configure “address-family ipv4 unicast” under the BGP neighbor command to enable ipv4 unicast for the BGP session.

You configure “address-family ipv6 unicast” under the BGP neighbor command to enable ipv6 unicast for the BGP session.

You configure “address-family ipv4 labeled-unicast” under the BGP neighbor command to enable ipv4 labeled unicast for the BGP session.

You configure “address-family ipv6 labeled-unicast” under the BGP neighbor command to enable ipv6 labeled unicast for the BGP session.

In IOS-XR the AFI/SAFI combination is configured per BGP peer.

This is an example for a BGP session having both SAFI 1 and 4:

```
router bgp 65003  
  address-family ipv4 unicast  
  redistribute connected  
  allocate-label all unlabeled-path  
...  
neighbor 10.100.1.7  
  remote-as 65003  
  update-source Loopback0  
  address-family ipv4 unicast  
  route-reflector-client  
!  
  address-family ipv4 labeled-unicast  
  route-reflector-client
```

Notice that there is still only the "address family unicast" and not "address-family labeled-unicast" under router BGP. Both SAFI 1 and 4 paths are stored in this one BGP table.

BGP Table

Regardless if IOS-XR is older or newer than 6.2.1, there is only one BGP table to store the U and LU routes. This is evident by the fact that you can only configure (enable) "address-family ipv4 unicast" or "address-family ipv6 unicast" under router bgp. You cannot configure "address-family ipv4 labeled-unicast" or "address-family ipv6 labeled-unicast" under router bgp.

The U and LU path can be identical. Before IOS-XR 6.2.1, receiving the same path again but this time with or without a label, would override the previously received path. After IOS-XR 6.2.1, the two identical paths will be seen as different if they only differ by the label. The path additions,

deletions, or modifications are performed by the different SAFIs.

Here is an example of a route in the BGP table with AFI 1/SAFI 4. Because label allocation is enabled for all prefixes, this path will be stored with a local label. Because there is only one BGP table to store U and LU routes, the prefix shows up with both the commands “show bgp ipv4 unicast” and “show bgp ipv4 labeled-unicast”!

```
RP/0/0/CPU0:R4#show bgp ipv4 unicast 10.100.1.1/32
BGP routing table entry for 10.100.1.1/32
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          5         5
  Local Label: 24000
Last Modified: Aug  6 15:03:59.574 for 16:06:13
Paths: (1 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.3 0.4
  Advertised to peers (in unique update groups):
    10.1.45.5
  Path #1: Received by speaker 0
  Advertised to update-groups (with more than one peer):
    0.3 0.4
  Advertised to peers (in unique update groups):
    10.1.45.5
65002 65001
  10.1.24.2 from 10.1.24.2 (10.100.1.2)
    Received Label 24003
    Origin IGP, localpref 100, valid, external, best, group-best, labeled-unicast
    Received Path ID 0, Local Path ID 0, version 5
    Origin-AS validity: not-found
```

Notice that the path is marked with “labeled-unicast”.

```
RP/0/0/CPU0:R4#show bgp ipv4 labeled-unicast 10.100.1.1/32
BGP routing table entry for 10.100.1.1/32
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          5         5
  Local Label: 24000
Last Modified: Aug  6 15:03:59.574 for 16:08:41
Paths: (1 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.3 0.4
  Advertised to peers (in unique update groups):
    10.1.45.5
  Path #1: Received by speaker 0
  Advertised to update-groups (with more than one peer):
    0.3 0.4
  Advertised to peers (in unique update groups):
    10.1.45.5
65002 65001
  10.1.24.2 from 10.1.24.2 (10.100.1.2)
    Received Label 24003
    Origin IGP, localpref 100, valid, external, best, group-best, labeled-unicast
    Received Path ID 0, Local Path ID 0, version 5
    Origin-AS validity: not-found
```

Notice that the path is marked with “labeled-unicast”.

If the path is present as both U and LU, then the local Path ID is different.

```

RP/0/0/CPU0:R4#show bgp ipv4 labeled-unicast 10.100.1.1/32 detail
BGP routing table entry for 10.100.1.1/32
Versions:
Process          bRIB/RIB  SendTblVer
Speaker          30        30
  Local Label: 24003 (no rewrite);
  Flags: 0x00003028+0x00010000;
Last Modified: Aug 30 10:45:50.502 for 00:01:59
Paths: (2 available, best #1)
  Advertised IPv4 Unicast paths to peers (in unique update groups):
    10.100.1.8    10.100.1.9
  Advertised IPv4 Labeled-unicast paths to update-groups (with more than one peer):
    0.8
Path #1: Received by speaker 0
  Flags: 0x4000000009060205, import: 0x20
  Advertised IPv4 Unicast paths to peers (in unique update groups):
    10.100.1.8    10.100.1.9
  Advertised IPv4 Labeled-unicast paths to update-groups (with more than one peer):
    0.8
65001, (Received from a RR-client)
  10.100.1.9 (metric 2) from 10.100.1.9 (10.100.1.9)
    Origin IGP, metric 0, localpref 100, valid, internal, best, group-best
    Received Path ID 0, Local Path ID 1, version 29
Path #2: Received by speaker 0
  Flags: 0x4080000008020205, import: 0x20
  Not advertised to any peer
65001, (Received from a RR-client)
  10.100.1.9 (metric 2) from 10.100.1.9 (10.100.1.9)
    Received Label 24001
    Origin IGP, metric 0, localpref 100, valid, internal, labeled-unicast
    Received Path ID 0, Local Path ID 0, version 0

```

Enabling the Local Label Allocation

You must configure the command “allocate-label” in order for the received or sourced paths in BGP to have a local MPLS label. Without this command, the routes will not have a local label.

```

RP/0/0/CPU0:R4#conf t
RP/0/0/CPU0:R4(config)#router bgp 65003
RP/0/0/CPU0:R4(config-bgp)# address-family ipv4 unicast
RP/0/0/CPU0:R4(config-bgp-af)#allocate-label ?
all          Allocate labels for all prefixes
route-policy Use a route policy to select prefixes for label allocation

```

The label allocation occurs for all routes or per the configured route-policy.

New Implementation on IOS-XR

In the old implementation on IOS-XR, a warning is given when configuring both the U and LU on the same BGP session. The warning is introduced in IOS-XR releases 5.3.0, and 5.2.2. The warning is removed in IOS-XR release 6.2.1, because labeled and unlabeled are supported on the same BGP session.

Example:

```

RP/0/0/CPU0:ios#conf t
RP/0/0/CPU0:ios(config)#router bgp 65001

```

```

RP/0/0/CPU0:ios(config-bgp)#add ipv4 unicast
RP/0/0/CPU0:ios(config-bgp-af)#exit
RP/0/0/CPU0:ios(config-bgp)#neighbor 10.0.0.1
RP/0/0/CPU0:ios(config-bgp-nbr)#remote-as 65001
RP/0/0/CPU0:ios(config-bgp-nbr)#exit
RP/0/0/CPU0:ios(config-bgp)#neighbor 10.0.0.1
RP/0/0/CPU0:ios(config-bgp-nbr)#address-family ipv4 unicast
RP/0/0/CPU0:ios(config-bgp-nbr-af)#exit
RP/0/0/CPU0:ios(config-bgp-nbr)#address-family ipv4 labeled-unicast
RP/0/0/CPU0:ios(config-bgp-nbr-af)#commit

```

```

RP/0/0/CPU0:Aug 21 14:14:22.222 : bgp[1052]: %ROUTING-BGP-4-INCOMPATIBLE_AFI : IPv4 Unicast and
IPv4 Labeled-unicast Address families together are not supported under the same neighbor.

```

The explanation for this error message:

This message indicates that the user has configured both IPv4 Unicast and IPv4 Labeled-unicast or IPv6 Unicast and IPv6 Labeled-unicast address families together under the same neighbor. This particular configuration is not supported.

Recommended action: Configure two neighbor sessions to the router. Configure the unicast address-family under the first neighbor session and configure the labeled-unicast address-family under the second neighbor session.

Configuration example for two BGP sessions between a pair of IOS-XR routers. Use a different (loopback) address for each BGP session.

```

hostname R1

interface Loopback0
  ipv4 address 10.100.1.1 255.255.255.255
!
interface Loopback1
  ipv4 address 10.100.1.101 255.255.255.255
!
router bgp 65001
  address-family ipv4 unicast
  !
  neighbor 10.100.1.2
  remote-as 65001
  update-source Loopback0
  address-family ipv4 unicast
  !
  !
  neighbor 10.100.1.102
  remote-as 65001
  update-source Loopback1
  address-family ipv4 labeled-unicast
  !
  !

hostname R2

interface Loopback0
  ipv4 address 10.100.1.2 255.255.255.255
!
interface Loopback1
  ipv4 address 10.100.1.102 255.255.255.255
!
router bgp 65001
  address-family ipv4 unicast

```

```

!
neighbor 10.100.1.1
  remote-as 65001
  update-source Loopback0
  address-family ipv4 unicast
!
!
neighbor 10.100.1.101
  remote-as 65001
  update-source Loopback1
  address-family ipv4 labeled-unicast
!
!

```

In IOS-XR 6.2.1, both U and LU are supported on the same BGP session **on the default VRF!**

It does not matter if the BGP session is internal or external BGP.

U and LU on the same session is not supported for BGP speakers in any non-default VRF.

Update-groups

Before IOS-XR 6.2.1, all U, LU, and U + LU BGP speakers were kept in separate update groups. After IOS-XR release 6.2.1, this is no longer true. Some BGP speakers in one update group can be only U, or only LU, or both U and LU.

Route Propagation

The following table shows the advertisement and withdraw behavior for different scenarios. There are 16 scenarios.

All applies to IOS-XR version 6.2.1 and later, unless otherwise mentioned in the comments column.

Case	Bestpath/Ad dpath type	Local label presen t?	NHS or NHU	Update- group SAFI	Advertise or withdraw?	Comments
1	Unlabeled path i.e. no rx label	Yes	NHS	SAFI-1	Advertise by default Withdraw with advertise local-labeled- route (safi- unicast) disable command	Only possible after 6.5.1.
2				SAFI-4	Advertise	Only possible after 6.5.1. IPv4/v6 redist routes and 6PE: implicit NHS always
3			NHU	SAFI-1	Advertise	Only possible after 6.5.1.
4				SAFI-4	Withdraw	Only possible after 6.5.1. IPv4/v6 redist routes and 6PE: NHU ignored; implicit NHS always
5		No	NHS	SAFI-1	Advertise	

6				SAFI-4	Withdraw	
7			NHU	SAFI-1	Advertise	
8				SAFI-4	Withdraw	
					Advertise by default	
					Withdraw	
9	Labeled path i.e. with rx label	Yes	NHS	SAFI-1	Withdraw with advertise local-labeled-route (safi-unicast) disable command	Before 6.2.1: default behavior is to Advertise.
10				SAFI-4	Advertise	
11			NHU	SAFI-1	Withdraw	Before 6.2.1: behavior is to advertise.
12				SAFI-4	Advertise	
13		No	NHS	SAFI-1	Advertise	
14				SAFI-4	Withdraw	
15			NHU	SAFI-1	Withdraw	Before 6.2.1: behavior is to advertise.
16				SAFI-4	Advertise	

Table 1 Advertisement behavior for iBGP and eBGP Sessions

NHS = Next Hop Self

NHU = Next Hop Unchanged.

If NHU is in effect, this means that next-hop self is not configured for an iBGP session.

Note that NHS is always the case when the BGP speaker sends to an eBGP peer.

There can be NHS or NHU towards an iBGP speaker, depending on the configuration of next-hop-self. The default behavior towards iBGP peers is NHU.

For the second column: note that the path is considered to be unlabeled or labeled only if the best-path or one of the paths marked with add-path is unlabeled or labeled.

For the route propagation it matters what characteristics the best path has. Depending on the characteristics (columns 2 to 4), it determines if the path is advertised as U, or LU, or both.

If the feature Additional Paths (ADD-PATH) is enabled and a path is marked with "add-path", the characteristics of that path also play a role how that path will be advertised.

“Local Label Present?: No” means the following: it is possible that a label is received with the received update, but the label is not installed. The local label is not installed if the command “allocate-label” is not there.

You can verify if the local label is present by looking at the prefix in detail. Use “show bgp <prefix> detail” or “show route <prefix> detail”.

In the following example, the prefix is received without a label (so over a SAFI 1 peering) and no local label is assigned:

```
RP/0/0/CPU0:R2#show bgp ipv4 labeled-unicast 10.100.1.5/32 detail
```

```
BGP routing table entry for 10.100.1.5/32
```

```
Versions:
```

```
Process          bRIB/RIB  SendTblVer
Speaker          3         3
```

```
Flags: 0x04001001+0x00000000;
```

```
Last Modified: Sep  5 03:44:45.647 for 01:01:27
```

```
Paths: (1 available, best #1)
```

```
Advertised to update-groups (with more than one peer):
```

```
0.3
```

```
Path #1: Received by speaker 0
```

```
Flags: 0x4000000001040207, import: 0x00
```

```
Advertised to update-groups (with more than one peer):
```

```
0.3
```

```
Local, (Received from a RR-client)
```

```
10.100.1.1 (metric 2) from 10.100.1.1 (10.100.1.1)
```

```
Origin incomplete, metric 0, localpref 100, valid, internal, best, group-best
```

```
Received Path ID 0, Local Path ID 1, version 3
```

```
RP/0/0/CPU0:R2#show route 10.100.1.5/32 detail
```

```
Routing entry for 10.100.1.5/32
```

```
Known via "bgp 65001", distance 200, metric 0, type internal
```

```
Installed Sep  5 03:44:45.480 for 01:01:37
```

```
Routing Descriptor Blocks
```

```
10.100.1.1, from 10.100.1.1
```

```
Route metric is 0
```

```
Label: None
```

```
Tunnel ID: None
```

```
Extended communities count: 0
```

```
NHID:0x0(Ref:0)
```

```
Route version is 0x23 (35)
```

```
No local label
```

```
IP Precedence: Not Set
```

```
QoS Group ID: Not Set
```

```
Flow-tag: Not Set
```

```
Route Priority: RIB_PRIORITY_RECURSIVE (12) SVD Type RIB_SVD_TYPE_LOCAL
```

```
Download Priority 4, Download Version 52
```

```
No advertising protos.
```

Allocating a Local Label

Allocate Label for Unlabeled Paths

By default, unlabeled paths (SAFI 1) are never labeled, even when the “allocate-label” command is configured.

As of IOS-XR release 6.5.1., there is the keyword “unlabeled-path” for the “allocate-label” command, so that also the unlabeled paths can be allocated a label.

```
RP/0/0/CPU0:R4#conf t
```

```
RP/0/0/CPU0:R4(config)#router bgp 65003
```

```
RP/0/0/CPU0:R4(config-bgp)# address-family ipv4 unicast
```

```
RP/0/0/CPU0:R4(config-bgp-af)#allocate-label all ?
```

```
unlabeled-path Allocate label for unlabeled paths too
```

```
<cr>
```

```
RP/0/0/CPU0:R4(config-bgp-af)#allocate-label all unlabeled-path ?
```

```
<cr>
```

```
RP/0/0/CPU0:R4(config-bgp-af)#allocate-label all unlabeled-path
```

```
RP/0/0/CPU0:R4(config-bgp-af)#commit
```

The path is a SAFI 1 path, so there is no received label.

Because of the “unlabeled-path” command there is now a local label.

```
RP/0/0/CPU0:R4#show bgp ipv4 labeled-unicast 10.100.1.1/32 detail
```

```
BGP routing table entry for 10.100.1.1/32
```

```
Versions:
```

```
Process          bRIB/RIB  SendTblVer
Speaker          16        16
```

```
Local Label: 24003 (no rewrite);
```

```
Flags: 0x01303028+0x00000000;
```

```
Last Modified: Aug 27 19:08:47.502 for 00:00:59
```

```
Paths: (1 available, best #1)
```

```
Advertised IPv4 Unicast paths to update-groups (with more than one peer):
```

```
0.3
```

```
Advertised IPv4 Labeled-unicast paths to update-groups (with more than one peer):
```

```
0.7
```

```
Advertised IPv4 Labeled-unicast paths to peers (in unique update groups):
```

```
10.1.45.5
```

```
Path #1: Received by speaker 0
```

```
Flags: 0x4000000009040207, import: 0x20
```

```
Advertised IPv4 Unicast paths to update-groups (with more than one peer):
```

```
0.3
```

```
Advertised IPv4 Labeled-unicast paths to update-groups (with more than one peer):
```

```
0.7
```

```
Advertised IPv4 Labeled-unicast paths to peers (in unique update groups):
```

```
10.1.45.5
```

```
65001, (Received from a RR-client)
```

```
10.100.1.10 (metric 2) from 10.100.1.10 (10.100.1.10)
```

```
Origin IGP, metric 0, localpref 100, valid, internal, best, group-best
```

```
Received Path ID 0, Local Path ID 1, version 16
```

```
RP/0/0/CPU0:R4#show route 10.100.1.1/32 detail
```

```
Routing entry for 10.100.1.1/32
```

```
Known via "bgp 65003", distance 200, metric 0
```

```
...
```

```
Route version is 0x4 (4)
```

```
Local Label: 0x5dc3 (24003)
```

```
IP Precedence: Not Set
```

```
...
```

This will allow the cases 1 to 4 in table 1.

To figure out why a local label is still assigned when the “allocate-label” command is removed, run the debug “debug bgp label”.

Here’s an example:

```
RP/0/0/CPU0:R4#debug bgp label
```

```
RP/0/0/CPU0:R4#show debug
```

```
#### debug flags set from tty 'con0_0_CPU0' ####
```

```
ip-bgp default label flag is ON with value '#####'
```

It is better to enable this debug for a specific prefix or group of prefixes. Here is an example:

```
RP/0/0/CPU0:R4#sh running-config route-policy match-prefix
route-policy match-prefix
  if destination in (10.100.1.1/32) then
    pass
  else
    drop
  endif
end-policy
!
```

```
RP/0/0/CPU0:R4#debug bgp label route-policy match-prefix
```

```
RP/0/0/CPU0:R4#show debug
```

```
#### debug flags set from tty 'con0_0_CPU0' ####
```

```
ip-bgp default label flag is ON with value '#####match-prefix####'
```

```
RP/0/0/CPU0:R4#con t
```

```
RP/0/0/CPU0:R4(config)#router bgp 65003
```

```
RP/0/0/CPU0:R4(config-bgp)# address-family ipv4 unicast
```

```
RP/0/0/CPU0:R4(config-bgp-af)#no allocate-label all
```

```
RP/0/0/CPU0:R4(config-bgp-af)#commit
```

```
RP/0/0/CPU0:Aug 23 12:43:02.786 : bgp[1048]: [default-lbl] (ip4u): Label computation done:
table=TBL:default (1/1), net=10.100.1.1/32: netfl=0x05043001,
path=0x1073ed5c(10.1.24.2/32,10.1.24.2,0,0x400000000d060001), pathrcvdlab=24002: asbr=1,
rr=0/1, nhselfcount=1: result="label required"
```

You can see that this router received a label for the prefix 10.100.1.1/32, is an ASBR, is not an RR, and has next-hop-self for at least one BGP session. This results in this prefix needing a local label.

The local label remains:

```
RP/0/0/CPU0:R4#show bgp ipv4 unicast 10.100.1.1/32 detail
```

```
BGP routing table entry for 10.100.1.1/32
```

```
Versions:
```

```
Process          bRIB/RIB  SendTblVer
Speaker          13        13
```

```
Local Label: 16002 (no rewrite);
```

```
Flags: 0x05043001+0x00000200;
```

```
Last Modified: Aug 23 12:37:11.133 for 00:05:53
```

```
Paths: (1 available, best #1)
```

```
Advertised to update-groups (with more than one peer):
0.6
```

```
Advertised to peers (in unique update groups):
```

```
10.1.46.6      10.100.1.8    10.100.1.7
```

```
Path #1: Received by speaker 0
```

```
Flags: 0x400000000d060001, import: 0x1f
```

```
Advertised to update-groups (with more than one peer):
0.6
```

```
Advertised to peers (in unique update groups):
```

```
10.1.46.6      10.100.1.8    10.100.1.7
```

```
65002 65001
```

```
10.1.24.2 from 10.1.24.2 (10.100.1.2)
```

```
Received Label 24002
```

```
Origin IGP, localpref 100, valid, external, best, group-best, import-candidate
```

```
Received Path ID 0, Local Path ID 1, version 13
```

```
Origin-AS validity: not-found
```

```
RP/0/0/CPU0:R4#show route 10.100.1.1/32 detail
```

```
Routing entry for 10.100.1.1/32
```

```
Known via "bgp 65003", distance 20, metric 0, [ei]-bgp, labeled unicast (3107)
```

```
Tag 65002, type external
```

```
Installed Aug 23 12:37:11.440 for 00:06:02
```

```
Routing Descriptor Blocks
```

```
10.1.24.2, from 10.1.24.2, BGP external
```

```
Route metric is 0
```

```
Label: 0x5dc2 (24002)
```

```
Tunnel ID: None
```

```
Extended communities count: 0
```

```
NHID:0x0(Ref:0)
```

```
Route version is 0x4 (4)
```

```
Local Label: 0x3e82 (16002)
```

```
IP Precedence: Not Set
```

```
QoS Group ID: Not Set
```

```
Route Priority: RIB_PRIORITY_NON_RECURSIVE_LOW (11) SVD Type RIB_SVD_TYPE_LOCAL
```

```
Download Priority 4, Download Version 28
```

```
No advertising protos.
```

The debug shows the following message when a local label is not required:

```
RP/0/0/CPU0:Aug 23 13:01:15.801 : bgp[1048]: [default-lbl]: Prefix 10.100.1.1/32:() doesn't require label, releasing
```

```
RP/0/0/CPU0:Aug 23 13:01:15.801 : bgp[1048]: [default-lbl]: bgp_label_release_label: perform label release onnet 10.100.1.1/32net retain 0 label_retain 0
```

If the prefix is in the LFIB depends if the prefix was received labeled or not and if allocate-label applies to that prefix.

The received label is 24002 for the following prefix. It is not installed in the LFIB, because BGP does not have the allocate-label command.

```
RP/0/0/CPU0:R4#show bgp ipv4 unicast 10.100.1.1/32
```

```
BGP routing table entry for 10.100.1.1/32
```

```
Versions:
```

```
Process          bRIB/RIB  SendTblVer
```

```
Speaker          4         4
```

```
Local Label: 24002
```

```
Last Modified: Aug 8 13:52:57.276 for 00:00:36
```

```
Paths: (1 available, best #1)
```

```
Advertised to update-groups (with more than one peer):
```

```
0.6
```

```
Advertised to peers (in unique update groups):
```

```
10.100.1.7
```

```
Path #1: Received by speaker 0
```

```
Advertised to update-groups (with more than one peer):
```

```
0.6
```

```
Advertised to peers (in unique update groups):
```

```
10.100.1.7
```

```
65002 65001
```

```
10.1.24.2 from 10.1.24.2 (10.100.1.2)
```

```
Received Label 24002
```

```
Origin IGP, localpref 100, valid, external, best, group-best, labeled-unicast
```

```
Received Path ID 0, Local Path ID 0, version 4
```

```
Origin-AS validity: not-found
```

```
router bgp 65003
```

```
bgp unsafe-ebgp-policy
```

```
address-family ipv4 unicast
```

!

RP/0/0/CPU0:R4# **show mpls forwarding**

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched
24000	Aggregate	10.1.24.0/24	default		0
24001	Aggregate	10.1.45.0/24	default		0

If the allocate-label command is present, then the local label is present in the LFIB:

```

router bgp 65003
  bgp unsafe-ebgp-policy
  address-family ipv4 unicast
  allocate-label all
!

```

RP/0/0/CPU0:R4#**show mpls forwarding**

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched
24000	Aggregate	10.1.24.0/24	default		0
24001	Aggregate	10.1.45.0/24	default		0
24002	24002	10.100.1.1/32		10.1.24.2	0

Even if the BGP prefix is received over an LU session, but no local label is assigned, then the route is not advertised over another LU session were NHS is done. This is case 14 in table 1. This is the case if the outgoing BGP session is eBGP.

Example:

RP/0/0/CPU0:R2#**show bgp ipv4 unicast 10.100.1.1/32 detail**

BGP routing table entry for 10.100.1.1/32

Versions:

```

Process          bRIB/RIB  SendTblVer
Speaker          3         3

```

Flags: 0x00001001+0x00000000;

Last Modified: Aug 22 09:00:20.646 for 00:10:56

Paths: (1 available, best #1)

Not advertised to any peer

Path #1: Received by speaker 0

Flags: 0x4080000001060001, import: 0x20

Not advertised to any peer

65001

10.1.12.1 from 10.1.12.1 (10.100.1.1)

Received Label 3

Origin IGP, metric 0, localpref 100, valid, external, best, group-best, labeled-unicast

Received Path ID 0, Local Path ID 0, version 3

Origin-AS validity: not-found

RP/0/0/CPU0:R2#**show route 10.100.1.1/32 detail**

Routing entry for 10.100.1.1/32

Known via "bgp 65002", distance 20, metric 0, labeled unicast (3107)

Tag 65001, type external

Installed Aug 22 09:00:20.416 for 00:10:59

Routing Descriptor Blocks

10.1.12.1, from 10.1.12.1, BGP external

Route metric is 0

Label: 0x100004 (1048580)

Tunnel ID: None

```
Binding Label: None
Extended communities count: 0
NHID:0x0(Ref:0)
Route version is 0x1 (1)
No local label
IP Precedence: Not Set
...
```

This is likely caused by not having the command “allocate-label” for address-family unicast under router BGP.

You will need to restart the process BGP when removing the “allocate-label” command in order for the router to remove the local labels for the BGP routes.

The new advertise local-labeled-route command in table 1 is a new command to indicate that a route with local label should not be advertised as an unlabeled route via SAFI-1.

This command is the following:

advertise local-labeled-route [disable]

This command is configured under the neighbor address-family. The function of this command is to indicate whether an IPv4/v6 route with a local label should be advertised or not to the BGP neighbor via IPv4/v6 unicast (SAFI 1).

The default behavior is to advertise routes with a local label.

The new command could also be configured as:

advertise local-labeled-route safi-unicast [disable]

This command is configured under the af-group of the BGP section. Its function is the same as the one above and applies to all BGP neighbors.

The default behavior is to advertise routes with a local label.

The line “Advertise routes with local-label via Unicast SAFI” or “Do not advertise routes with local-label via Unicast SAFI” is present on the command "show bgp neighbor" under the address family IPv4 Unicast to indicate that the BGP speaker permits the advertisement of routes with local label or not.

Example for the default behavior:

```
RP/0/0/CPU0:R4#show bgp neighbor 10.1.45.5
...
For Address Family: IPv4 Unicast
BGP neighbor version 5
Update group: 0.1 Filter-group: 0.5 No Refresh request being processed
Extended Nexthop Encoding: advertised and received
Route refresh request: received 0, sent 0
0 accepted prefixes, 0 are bestpaths
Exact no. of prefixes denied : 0.
Cumulative no. of prefixes denied: 0.
Prefix advertised 2, suppressed 0, withdrawn 0
Maximum prefixes allowed 1048576
Threshold for warning message 75%, restart interval 0 min
```

```
An EoR was not received during read-only mode
Last ack version 5, Last synced ack version 0
Outstanding version objects: current 0, max 1, refresh 0
Additional-paths operation: None
```

```
Advertise routes with local-label via Unicast SAFI
```

```
...
```

```
or
```

```
RP/0/0/CPU0:R4# conf t
RP/0/0/CPU0:R4(config)#router bgp 65003
RP/0/0/CPU0:R4(config-bgp)# neighbor 10.1.45.5
RP/0/0/CPU0:R4(config-bgp-nbr)# address-family ipv4 unicast
RP/0/0/CPU0:R4(config-bgp-nbr-af)#advertise local-labeled-route disable
RP/0/0/CPU0:R4(config-bgp-nbr-af)#commit
```

```
RP/0/0/CPU0:R4#show bgp neighbor 10.1.45.5
```

```
BGP neighbor is 10.1.45.5
```

```
...
```

```
For Address Family: IPv4 Unicast
  BGP neighbor version 5
  Update group: 0.1 Filter-group: 0.5 (Update-group Change
  pending)
  No Refresh request being processed
  Extended Nexthop Encoding: advertised and received
  Route refresh request: received 0, sent 0
  0 accepted prefixes, 0 are bestpaths
  Exact no. of prefixes denied : 0.
  Cumulative no. of prefixes denied: 0.
  Prefix advertised 2, suppressed 0, withdrawn 0
  Maximum prefixes allowed 1048576
  Threshold for warning message 75%, restart interval 0 min
  An EoR was not received during read-only mode
  Last ack version 5, Last synced ack version 0
  Outstanding version objects: current 0, max 1, refresh 0
  Additional-paths operation: None
  Do not advertise routes with local-label via Unicast SAFI
```

Bestpath Calculation

There are no changes in the bestpath calculation process. If a path is SAFI 1 or SAFI 4 or if the path has a label or not, makes no difference in the bestpath calculation process. Hence, there is no preference between a SAFI 1 or SAFI 4 path. This is regardless if there is SAFI 1/SAFI 4 on the same BGP session or on different sessions. So, if one BGP session is SAFI 1 and 4, and a prefix is received over both address families, then the best path calculation will pick one as the best path, since all attributes are the same. **If all BGP attributes are the same between the U and LU path, the path received last, becomes the best path.**

If the SAFI 1 and SAFI 4 paths are received from different BGP peers, there is always a difference in the paths leading to BGP always picking the same best path from the two paths. Even if in this case all attributes are the same, the neighbor address is different. Looking at the [BGP Best Path Selection Algorithm](#), the path from the neighbor with the lowest neighbor address (final step 13) is picked as the best path.

Use the command “show bgp <AFI> <SAFI> <prefix> bestpath-compare” to verify the reason why the best path is best.

This preference can be made by the user, by using RPL.

Here's an example of such RPL.

```
RP/0/0/CPU0:R7#show bgp ipv4 un 10.100.1.1/32 detail
BGP routing table entry for 10.100.1.1/32
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          682      682
  Flags: 0x00003001+0x00010000;
Last Modified: Aug 28 13:16:26.826 for 00:00:10
Paths: (2 available, best #2)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Flags: 0x4000000000020005, import: 0x20
  Not advertised to any peer
  65001
    10.100.1.4 (metric 2) from 10.100.1.4 (10.100.1.10)
      Origin IGP, metric 0, localpref 100, valid, internal
      Received Path ID 1, Local Path ID 0, version 0
      Originator: 10.100.1.10, Cluster list: 10.100.1.4
  Path #2: Received by speaker 0
  Flags: 0x4080000001060005, import: 0x20
  Not advertised to any peer
  65001
    10.100.1.4 (metric 2) from 10.100.1.4 (10.100.1.10)
      Received Label 24003
      Origin IGP, metric 0, localpref 100, valid, internal, best, group-best, labeled-unicast
      Received Path ID 1, Local Path ID 0, version 682
      Originator: 10.100.1.10, Cluster list: 10.100.1.4
```

The LU path is the best.

RPL with weight is used to prefer the U path.

```
route-policy weight
  if destination in (10.100.1.1/32) then
    set weight 60000
  endif
end-policy
```

```
router bgp 65003
  address-family ipv4 unicast
  additional-paths receive
  additional-paths send
  !
  neighbor 10.100.1.4
  remote-as 65003
  update-source Loopback0
  address-family ipv4 unicast
    route-policy weight in
  !
  address-family ipv4 labeled-unicast
  !
  !
```

```
RP/0/0/CPU0:R7#show bgp ipv4 un 10.100.1.1/32 bestpath-compare
BGP routing table entry for 10.100.1.1/32
Versions:
  Process          bRIB/RIB  SendTblVer
```

```

Speaker                726                726
Last Modified: Aug 28 13:39:27.826 for 00:04:54
Paths: (2 available, best #1)
Not advertised to any peer
Path #1: Received by speaker 0
Not advertised to any peer
65001
  10.100.1.4 (metric 2) from 10.100.1.4 (10.100.1.10)
    Origin IGP, metric 0, localpref 100, weight 60000, valid, internal, best, group-best
    Received Path ID 1, Local Path ID 0, version 726
    Originator: 10.100.1.10, Cluster list: 10.100.1.4
    best of AS 65001, Overall best
Path #2: Received by speaker 0
Not advertised to any peer
65001
  10.100.1.4 (metric 2) from 10.100.1.4 (10.100.1.10)
    Received Label 24003
    Origin IGP, metric 0, localpref 100, valid, internal, labeled-unicast
    Received Path ID 1, Local Path ID 0, version 0
    Originator: 10.100.1.10, Cluster list: 10.100.1.4
    Lower weight than best path (path #1)

```

The U path is now the best.

There is no new command to prefer labeled paths over unlabeled paths. You can just configure the RPL either under address-family unicast or labeled-unicast under a BGP neighbor.

Verifying the Behavior in BGP

To debug the BGP update propagation in IOS-XR, you can turn on the following debug command:
debug bgp update <BGP neighbor> in | out.

This will show the incoming or outgoing BGP update from or to that BGP speaker. The address family is shown as (**ip4u**) for unlabeled IPv4 unicast (AFI 1/SAFI 1) or as (**ipv4lu**) for labeled IPv4 unicast (AFI 1/SAFI 4). The equivalent occurs for IPv6.

There is a new field “labeled-unicast” that indicates that the path is learnt via SAFI 4.

Example:

```

RP/0/0/CPU0:R1#show bgp ipv4 unicast 10.100.1.7/32
BGP routing table entry for 10.100.1.7/32
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          26        26
Last Modified: Sep  4 10:45:44.551 for 00:29:11
Paths: (1 available, best #1)
Not advertised to any peer
Path #1: Received by speaker 0
Not advertised to any peer
Local
  10.100.1.4 (metric 3) from 10.100.1.102 (10.100.1.4)
    Received Label 24000
    Origin incomplete, metric 0, localpref 100, valid, internal, best, group-best, labeled-unicast
    Received Path ID 0, Local Path ID 1, version 26
    Originator: 10.100.1.4, Cluster list: 10.100.1.2

```

To verify if the prefix is advertised, you can use the “show bgp ... neighbors” command with the

keyword "advertised-routes" at the end.

Example:

R4 advertises 10.100.1.1/32 to neighbor 10.100.1.7 twice because add-path is enabled (the two paths are different).

```
RP/0/0/CPU0:R4#show bgp ipv4 labeled-unicast neighbors 10.100.1.7 advertised-routes
```

Network	Next Hop	From	AS Path
10.1.24.0/24	10.100.1.4	Local	?
10.1.34.0/24	10.100.1.4	Local	?
10.1.45.0/24	10.100.1.4	Local	?
10.1.46.0/24	10.100.1.4	Local	?
10.1.47.0/24	10.100.1.4	Local	?
10.1.48.0/24	10.100.1.4	Local	?
10.1.49.0/24	10.100.1.4	Local	?
10.1.104.0/24	10.100.1.4	Local	?
10.1.114.0/24	10.100.1.4	Local	?
10.100.1.1/32	10.100.1.4	10.100.1.9	65001i
		10.100.1.10	65001i
10.100.1.4/32	10.100.1.4	Local	?

Processed 11 prefixes, 12 paths

iBGP and next-hop-self – Unified MPLS

The rules from table 1 apply. With Unified MPLS or Seamless MPLS, an Area Border Router (ABR) acts as route-reflector but is also the next-hop for the iBGP routes. The ABRs are in the forwarding path of the labeled traffic. The ABRs must have the explicit configuration for next-hop-self.

Workarounds

- If there is no support for U and LU over one BGP session, a workaround is to have two BGP sessions between a pair of routers. Two (loopback) interfaces must be used on each router.

Example configuration for IOS-XR.

```
interface Loopback0
  ipv4 address 10.100.1.7 255.255.255.255
!
interface Loopback1
  ipv4 address 10.100.1.107 255.255.255.255
!
router bgp 65003
  address-family ipv4 unicast
!
  neighbor 10.100.1.4 -> towards loopback0 on peer
  remote-as 65003
  update-source Loopback0
  address-family ipv4 unicast
!
!
  neighbor 10.100.1.104 -> towards loopback1 on peer
  remote-as 65003
```

update-source Loopback1

address-family ipv4 labeled-unicast

!

The U and LU paths are send/received over two different BGP sessions.

RP/0/0/CPU0:R7#**show bgp ipv4 unicast 10.100.1.1/32 detail**

BGP routing table entry for 10.100.1.1/32

Versions:

Process	bRIB/RIB	SendTblVer
Speaker	753	753

Flags: 0x00001001+0x00010000;

Last Modified: Aug 28 14:06:40.826 for 00:22:10

Paths: (2 available, best #1)

Not advertised to any peer

Path #1: Received by speaker 0

Flags: 0x4000000001060005, import: 0x20

Not advertised to any peer

65001

10.100.1.4 (metric 2) from 10.100.1.4 (10.100.1.10)

Origin IGP, metric 0, localpref 100, valid, internal, best, group-best

Received Path ID 1, Local Path ID 0, version 753

Originator: 10.100.1.10, Cluster list: 10.100.1.4

Path #2: Received by speaker 0

Flags: 0x4080000000020005, import: 0x20

Not advertised to any peer

65001

10.100.1.104 (metric 2) from 10.100.1.104 (10.100.1.10)

Received Label 24003

Origin IGP, metric 0, localpref 100, valid, internal, **labeled-unicast**

Received Path ID 1, Local Path ID 0, version 0

Originator: 10.100.1.10, Cluster list: 10.100.1.4