

# 在IOS-XR上的一个BGP邻居上标记和未标记

## 目录

[先决条件](#)

[要求](#)

[使用的组件](#)

[简介](#)

[功能交换](#)

[IOS-XR实施细节](#)

[配置U或LU](#)

[BGP表](#)

[启用本地标签分配](#)

[IOS-XR上的新实施](#)

[更新组](#)

[路由传播](#)

[分配本地标签](#)

[为未标记路径分配标签](#)

[最佳路径计算](#)

[检验BGP中的行为](#)

[iBGP和下一跳自身 — 统一MPLS](#)

[解决方法](#)

[IOS-XR的配置示例。](#)

本文档介绍在Cisco IOS® XR中通过一个BGP会话接收和通告未标记和已标记的路径的行为。

## 先决条件

### 要求

本文档没有任何特定的要求。

### 使用的组件

本文档特定于Cisco IOS® XR，但不限于特定软件版本或硬件。

本文档中的信息都是基于特定实验室环境中的设备编写的。本文档中使用的所有设备最初均采用原始（默认）配置。如果您使用的是真实网络，请确保您已经了解所有命令的潜在影响。

## 简介

地址系列标识符(AFI)表示BGP路由的类型。例如，IPv4为1,IPv6为2。

后续地址系列标识符(SAFI)是路由类型的进一步指示。示例为1表示未标记的路由，4表示已标记的路由。

IPv4的未标记单播是AFI 1和SAFI 1。

IPv4的单播标记为AFI 1和SAFI 4。

IPv6的未标记单播是AFI 2和SAFI 1。

IPv6的单播标记为AFI 2和SAFI 4。

标记单播(LU)通常称为RFC 3107“在BGP-4中传输标签信息”。

下面，U表示未标记的单播，因此SAFI 1和LU表示已标记的单播，因此SAFI 4。

请注意，Cisco IOS® XR需要“allocate-label all| route-policy ...”，否则路由不会作为SAFI 4发起或传播到下一个BGP发言者。

Cisco IOS® XR不支持一个BGP会话上的IPv4/v6单播和标记单播。

在Cisco IOS® XR上，在6.2.1中完成了对在一个BGP会话上同时进行未标记和标记的支持。

当不支持在一个会话上同时运行时，会出现问题，因为上次收到的更新/撤消会覆盖前一个会话，即使它们是在另一个SAFI上收到的。

在运行IOS-XR 6.2.1之前的IOS-XR代码的路由器上，在同一BGP会话上配置SAFI 1和SAFI 4时，路由器会发出以下警告：

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

IOS-XR 5.3.0和IOS-XR 5.2.2中引入了此警告消息。

## 功能交换

BGP对等体之间交换的功能必须匹配。否则，BGP会话不会启动。

这是BGP打开消息中为AFI 1/SAFI 1和AFI 1/SAFI 4交换的功能的Wireshark捕获：

```

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

```

图1

以下是仅在与仅配置了U的IOS的会话上配置了LU的IOS-XR的示例。

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

IOS路由器会针对此配置错误打印系统日志消息：

```
*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实施细节

### 配置U或LU

在BGP neighbor命令下配置“address-family ipv4 unicast”，以为BGP会话启用ipv4单播。

在BGP neighbor命令下配置“address-family ipv6 unicast”，为BGP会话启用ipv6 unicast。

在BGP neighbor命令下配置“address-family ipv4 labeled-unicast”，以启用BGP会话的ipv4 labeled unicast。

在BGP neighbor命令下配置“address-family ipv6 labeled-unicast”，为BGP会话启用ipv6 labeled unicast。

在IOS-XR中，AFI/SAFI组合是按BGP对等体配置的。

以下是同时具有SAFI 1和4的BGP会话的示例：

```
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
```

请注意，路由器BGP下仍然只有“地址系列单播”，而不是“地址系列标签单播”。SAFI 1和4路径都存储在此BGP表中。

### BGP表

无论IOS-XR是早于还是高于6.2.1，只有一个BGP表用于存储U和LU路由。这一点显而易见，因为您只能在路由器bgp下配置（启用）“地址系列ipv4单播”或“地址系列ipv6单播”。您不能在路由器bgp下配置“address-family ipv4 labeled-unicast”或“address-family ipv6 labeled-unicast”。

U和LU路径可以相同。在IOS-XR 6.2.1之前，再次接收同一路径，但这一次，无论是否带有标签，都将覆盖之前收到的路径。在IOS-XR 6.2.1之后，如果两条相同的路径仅因标签而不同，则它们会被视为不同。路径添加、删除或修改由不同的SAFI执行。

以下是AFI 1/SAFI 4的BGP表中路由的示例。由于为所有前缀启用了标签分配，因此此路径将与本地标签一起存储。由于只有一个BGP表用于存储U和LU路由，因此前缀会同时显示命令“show bgp

## ipv4 unicast”和“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
```

请注意，路径标有“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
```

请注意，路径标有“labeled-unicast”。

如果路径同时作为U和LU存在，则本地路径ID不同。

```
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

```

## 启用本地标签分配

您必须配置命令“allocate-label”，以便BGP中的接收或源路径具有本地MPLS标签。如果没有此命令，路由将没有本地标签。

```

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

```

标签分配适用于所有路由或根据已配置的路由策略进行。

## IOS-XR上的新实施

在IOS-XR的旧实施中，在同一BGP会话上配置U和LU时会发出警告。警告在IOS-XR版本5.3.0和5.2.2中引入。警告在IOS-XR版本6.2.1中删除，因为同一BGP会话支持标记和未标记。

示例：

```

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.
```

此错误消息的说明：

此消息表明用户已在同一邻居下同时配置了IPv4单播和IPv4 Labeled-unicast或IPv6 Unicast和IPv6 Labeled-unicast地址系列。不支持此特定配置。

建议的操作：配置到路由器的两个邻居会话。在第一个邻居会话下配置单播地址系列，在第二个邻居会话下配置标记单播地址系列。

一对IOS-XR路由器之间两个BGP会话的配置示例。对每个BGP会话使用不同的（环回）地址。

```
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
```

```
!
```

在IOS-XR 6.2.1中，U和LU在默认VRF上的同一BGP会话上都受支持!

无论BGP会话是内部BGP还是外部BGP都无关紧要。

在任何非默认VRF中，BGP扬声器不支持同一会话上的U和LU。

## 更新组

在IOS-XR 6.2.1之前，所有U、LU和U + LU BGP扬声器都保留在单独的更新组中。在IOS-XR版本6.2.1之后，情况不再如此。一个更新组中的某些BGP扬声器只能是U或LU，或同时是U和LU。

## 路由传播

下表显示了不同场景的广告和撤消行为。有16种方案。

除非注释列中另有说明，否则所有内容均适用于IOS-XR版本6.2.1及更高版本。

案例	最佳路径/地址 路径类型	存在本地 地址标签 ?	NHS 或 NHU	更新组 SAFI	广告还是退出 ?	备注
1	未标记路径 ，即无接收器 标签	Yes	NHS	SAFI-1	默认情况下通告 退回 <i>advertise local-labeled- route(safi- unicast)disabl e命令</i>	仅在6.5.1之后可能。
2				SAFI-4	通告	仅在6.5.1之后可能。
3			NHU	SAFI-1	通告	IPv4/v6 redist路由和6PE:隐式NHS始终 仅在6.5.1之后可能。
4				SAFI-4	退回	仅在6.5.1之后可能。 IPv4/v6 redist路由和6PE:NHU被忽略 ；隐式NHS始终
5		无	NHS	SAFI-1	通告	
6				SAFI-4	退回	
7			NHU	SAFI-1	通告	
8				SAFI-4	退回	
9	标记路径，即 带有rx标签	Yes	NHS	SAFI-1	默认情况下通告 使用advertise local-labeled- route(safi- unicast)disabl e命令撤消	6.2.1之前：默认行为是通告。
10				SAFI-4	通告	
11			NHU	SAFI-1	退回	6.2.1之前：行为是通告。
12				SAFI-4	通告	
13 个		无	NHS	SAFI-1	通告	
14				SAFI-4	退回	
15			NHU	SAFI-1	退回	6.2.1之前：行为是通告。



表1 iBGP和eBGP会话的通告行为

NHS =下一跳自己

NHU =下一跳未更改。

如果NHU生效，这意味着未为iBGP会话配置下一跳自己。

请注意，BGP扬声器向eBGP对等体发送消息时，NHS始终是这种情况。

NHS或NHU可能针对iBGP扬声器，具体取决于下一跳自身的配置。对iBGP对等体的默认行为是NHU。

第二列：请注意，仅当最佳路径或标记有add-path的路径之一未标记或标记时，路径才被视为未标记或标记。

对于路由传播，最佳路径的特征至关重要。根据特征（列2到4），它确定路径是通告为U还是LU，还是同时通告为两者。

如果启用了功能Additional Paths(ADD-PATH)，并且路径标有“add-path”，则该路径的特征也将扮演如何通告该路径的角色。

“存在本地标签？”：“否”表示：可能收到带有已接收更新的标签，但标签未安装。如果“allocate-label”命令不存在，则不会安装本地标签。

您可以通过详细查看前缀来验证本地标签是否存在。使用“show bgp <prefix> detail”或“show route <prefix> detail”。

在以下示例中，接收前缀时不带标签（通过SAFI 1对等），并且不分配本地标签：

```
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.
```

## 分配本地标签

### 为未标记路径分配标签

默认情况下，无标记路径(SAFI 1)从不进行标记，即使配置了“allocate-label”命令也是如此。

自IOS-XR版本6.5.1起，“allocate-label”命令使用关键字“unlabeled-path”，因此也可以为未标记的路径分配标签。

```
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
```

该路径是SAFI 1路径，因此没有收到标签。

由于使用了“unlabeled-path”命令，现在有一个本地标签。

```
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
...
```

这将允许表1中的问题1到4。

要了解删除“allocate-label”命令时仍分配本地标签的原因，请运行debug bgp label。

以下是一个示例：

```
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 '#####'
```

最好为特定前缀或前缀组启用此调试。示例如下：

```
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"
```

您可以看到，此路由器收到前缀10.100.1.1/32的标签，是ASBR，不是RR，并且至少有一个BGP会话具有下一跳自己。这会导致此前缀需要本地标签。

本地标签仍保留：

```
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: 0x40000000d060001, 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.
```

当不需要本地标签时，调试显示以下消息：

```
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
```

如果前缀在LFIB中，则取决于前缀是否被标记，以及allocate-label是否适用于该前缀。

收到的标签为24002，表示以下前缀。由于BGP没有allocate-label命令，因此LFIB中未安装该命令。

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

如果allocate-label命令存在，则本地标签在LFIB中存在：

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

```
RP/0/0/CPU0:R4#show mpls forwarding
Local  Outgoing  Prefix          Outgoing  Next Hop  Bytes
Label  Label      or ID           Interface  Next Hop  Switched
-----
24000  Aggregate  10.1.24.0/24    default   default   0
24001  Aggregate  10.1.45.0/24    default   default   0
24002 24002      10.100.1.1/32    default   10.1.24.2 0
```

即使BGP前缀通过LU会话接收，但未分配本地标签，在NHS完成后，路由也不会通过另一个LU会话通告。表1中的情况为14。如果传出BGP会话为eBGP，则情况为14。

示例：

```
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
```

```
...
```

这可能是因为在路由器BGP下没有地址系列单播的命令“allocate-label”。

**删除“allocate-label”命令时，您需要重新启动进程BGP，以便路由器删除BGP路由的本地标签。**

表1中的新advertise local-labeled-route命令是一个新命令，用于指示具有本地标签的路由不应通过SAFI-1通告为无标签路由。

此命令如下：

```
advertise local-labeled-route [disable]
```

此命令在neighbor address-family下配置。此命令的功能是指示是否应通过IPv4/v6单播(SAFI 1)将带有本地标签的IPv4/v6路由通告给BGP邻居。

默认行为是通告带有本地标签的路由。

新命令也可配置为：

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

此命令在BGP部分的af-group下配置。其功能与上述功能相同，适用于所有BGP邻居。

默认行为是通告带有本地标签的路由。

地址系列IPv4单播下的“Advertise routes with local-label via Unicast SAFI”或“Do not advertise routes with local-label via Unicast SAFI”一行显示在命令上，表示BGP扬声器允许通告带有本地标签的路由。

默认行为示例：

```
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
...
或
```

```
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

## 最佳路径计算

最佳路径计算过程中没有更改。如果路径是SAFI 1或SAFI 4，或者路径是否带有标签，则最佳路径计算过程不会有任何差异。因此，SAFI 1或SAFI 4路径之间没有首选项。无论同一BGP会话或不同会话上是否存在SAFI 1/SAFI 4。因此，如果一个BGP会话是SAFI 1和4，并且在两个地址系列上都收到前缀，则最佳路径计算将选择一个作为最佳路径，因为所有属性都相同。如果U和LU路径之间的所有BGP属性相同，则最后收到的路径成为最佳路径。

如果从不同的BGP对等体接收SAFI 1和SAFI 4路径，则通向BGP的路径始终存在差异，始终从两条路径中选择相同的最佳路径。即使在这种情况下，所有属性都相同，邻居地址也不同。查看BGP[最佳路径选择算法](#)，从具有最低邻居地址的邻居选择路径（最后步骤13）作为最佳路径。

使用命令“show bgp <AFI> <SAFI> <prefix> bestpath-compare”验证最佳路径最佳的原因。

此首选项可由用户使用RPL进行。

下面是此类RPL的示例。

```
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
```

LU路径是最佳路径。

带权重的RPL用于首选U路径。

```
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)
```

U路径现在是最佳路径。

没有新命令可优先选择标记路径而非未标记路径。您只需在地址系列单播或BGP邻居下的标记单播下配置RPL。

## 检验BGP中的行为

要调试IOS-XR中的BGP更新传播，可以打开以下debug命令：*debug bgp update <BGP neighbor> in* /输出。

这将显示从该BGP扬声器传入或传出的BGP更新。对于未标记的IPv4未播(AFI 1/SAFI 1)，地址系列显示为(ip4u)；对于已标记的IPv4单播(AFI 1/SAFI 4)，显示为(ipv4lu)。IPv6也会出现这种情况。

有一个新字段“labeled-unicast”，表示路径是通过SAFI 4获知的。

示例：

```
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
```

要验证前缀是否通告，可以在末尾使用“show bgp ... neighbors”命令和关键字“advertised-routes”。

示例：

R4两次向邻居10.100.1.7通告10.100.1.1/32，因为已启用add-path（两条路径不同）。

```
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	?
<b>10.100.1.1/32</b>	<b>10.100.1.4</b>	<b>10.100.1.9</b>	<b>65001i</b>
		<b>10.100.1.10</b>	<b>65001i</b>
10.100.1.4/32	10.100.1.4	Local	?

```
Processed 11 prefixes, 12 paths
```

## iBGP和下一跳自身 — 统一MPLS

表1中的规则适用。使用统一MPLS或无缝MPLS时，区域边界路由器(ABR)既充当路由反射器，也是iBGP路由的下一跳。ABR位于标记流量的转发路径中。ABR必须具有next-hop-self的显式配置。

## 解决方法

- 如果在一个BGP会话上不支持U和LU，则解决方法是在一对路由器之间有两个BGP会话。每台路由器上必须使用两个（环回）接口。

## IOS-XR的配置示例。

```
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
!

```

U和LU路径通过两个不同的BGP会话发送/接收。

```

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

```