

# 在Cisco IOS XR中部署BGP软下一跳

## 目录

[简介](#)

[先决条件](#)

[使用的组件](#)

[背景信息](#)

[解决方案](#)

[BGP路由的下一跳不可访问](#)

[配置](#)

[BGP ODN AS NextHop软验证命令](#)

[BGP ODN AS下一跳硬验证命令](#)

[考虑SR策略路径度量的BGP最佳路径选择](#)

[BGP最佳路径选择优先选择SR策略路径](#)

[BGP最佳路径选择强制SR策略路径](#)

[概述](#)

[默认行为](#)

[使用SR策略度量进行RIB相关验证](#)

[SR策略相关验证与RIB度量](#)

[SR策略相关验证与SR策略度量](#)

[SR策略相关验证，RIB度量和SR策略不用于最佳路径计算](#)

[SR策略相关验证，使用RIB度量和SR策略进行最佳路径计算](#)

[确认](#)

[检查策略是打开还是关闭](#)

[验证是否使用策略](#)

[Show BGP Nexthops命令](#)

[BGP跟踪](#)

[管理距离和度量](#)

[管理员值](#)

[验证ODN中的管理和度量类型](#)

[有效指标](#)

[BGP路径比较](#)

[比较BGP路径是否带颜色](#)

[使用show命令更详细地了解场景](#)

[默认行为](#)

[使用SR策略度量进行RIB相关验证](#)

[SR策略相关验证与RIB度量](#)

[SR策略相关验证与SR策略度量](#)

[SR策略相关验证，RIB度量和SR策略不用于最佳路径计算](#)

[SR策略相关验证，使用RIB度量和SR策略进行最佳路径计算](#)

## 简介

本文档介绍Cisco IOS® XR中边界网关协议(BGP)的下一跳处理<sup>行为</sup>。BGP要求在路由信息库(RIB)中安装路径之前，路径的下一跳(NH)是可到达的。此规则适用于所有BGP扬声器。这是下一跳验证检查。BGP软下一跳功能可确保不再需要在RIB中访问BGP下一跳。

## 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。

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。如果您的网络处于活动状态，请确保您了解所有命令的潜在影响。

## 背景信息

在单个自治系统(AS)网络、多域网络或AS间场景中，如果NH未在域或自治系统之间重分布，则可能无法到达NH。

问题不仅限于前端提供商边缘(PE)路由器，还包括出口和入口PE之间的中间BGP扬声器(例如路由反射器(RR)和自治系统边界路由器(ASBR))。中间BGP扬声器在安装和传播路由之前必须具有到NH的连通性。

按需下一跳(ODN)是在路由器上安装SR策略的分段路由(SR)应用程序。连接到这些SR策略的服务路由可以是BGP路由。如果下一跳有效，则这些BGP路由只能安装在RIB和思科快速转发(CEF)表中。有无缝MPLS或AS间MPLS虚拟专用网(VPN)等设计，其中RIB中的路由不保证到网络另一部分(如另一个区域或另一个域)中BGP下一跳的可达性。如果可达性由控制器或SR路径计算元素(SR-PCE)保证，而SR路径计算元素提供对整个网络中网络元素的可达性，则这不是问题。

当前，BGP服务路由仅当BGP路由的下一跳在RIB中作为非默认路由时，才能使用SR策略。

如果具有SR策略的BGP发言者在RIB中没有路由(默认路由除外)用于BGP下一跳，则可以使用解决方法。解决方法是将特定(非默认)静态路由配置为null0，覆盖那些不可达的NH，通过BGP-LU注入路由，或在IGP域之间重分布路由。

这很麻烦，并且/或影响了可扩展性。

## 解决方案

PE(头端)接收有色BGP L3VPN前缀。它可以在本地学习SR策略，或为颜色和下一跳请求ODN SR策略。

如果配置了NH验证，BGP将对NH执行软验证，并在启用命令时应用NH AD/metric。对于彩色NH，AD/度量来自SR控制器。下一跳的软验证意味着不检查RIB可达性，但检查是对SR策略信息执行的。这包括SR策略路由类型、管理距离和该度量类型的度量值。

引入新命令，在头端路由器或路由反射器上执行此下一跳验证。

为RR引入了新命令，以跳过彩色扩展通信路径的下一跳可达性验证。

为RR引入了新命令，以便SR策略不用于BGP最佳路径计算。

此功能在Cisco IOS XR版本7.3.2和7.4.1中引入。

## BGP路由的下一跳不可访问

不通告具有不可访问下一跳的BGP路由。

此路由是路由反射器上的VPNv4路由。其下一跳（PE环回）不可访问，因为路由表中没有下一跳地址的路由。

```
RP/0/RP0/CPU0:RR#show bgp vpnv4 unicast rd 65001:2 10.0.0.9/32
BGP routing table entry for 10.0.0.9/32, Route Distinguisher: 65001:2
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          0         0
Last Modified: Oct 26 10:40:12.136 for 00:03:07
Paths: (1 available, no best path)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
  65002, (Received from a RR-client)
    10.0.0.5 (inaccessible) from 10.0.0.5 (10.0.0.5)
      Received Label 24002
      Origin IGP, metric 0, localpref 100, valid, internal, not-in-vrf
      Received Path ID 0, Local Path ID 0, version 0
      Extended community: Color:101 RT:65001:101
```

因此，BGP VPNv4路由不会通告。

```
RP/0/RP0/CPU0:RR#show route 10.0.0.5
Routing entry for 0.0.0.0/0
  Known via "isis 1", distance 115, metric 20, candidate default path, type level-1
  Installed Oct 25 09:35:07.256 for 1d01h
  Routing Descriptor Blocks
    10.2.7.2, from 10.0.0.3, via GigabitEthernet0/0/0/0
    Route metric is 20
  No advertising protos.
```

当前的解决方法是配置一条静态路由，该路由覆盖头端路由器上的PE环回地址。这是到Null0的此类静态路由的示例。

```
address-family ipv4 unicast
  10.0.0.0/24 Null0
!
```

此到Null0的静态路由在RIB中为所有远程PE环回地址（BGP下一跳地址）创建可达性。此静态路由涵盖10.0.0.0 - 10.0.0.255范围内的所有地址。

下一跳通过静态路由解析。您可以使用此命令查看此内容。

```
RP/0/RP0/CPU0:PE1#show bgp vpnv4 unicast nexthops 10.0.0.5 color 101
```

```
Nexthop: 10.0.0.5 C:101
```

```
VRF: default
```

```
Nexthop ID: 0x6000008, Version: 0x0
```

```
Nexthop Flags: 0x00480002
```

```
Nexthop Handle: 0x7fa734042e94
```

```
RIB Related Information:
```

```
Firsthop interface handle 0x0000000c
```

```
Gateway TBL Id: 0xe0000000 Gateway Flags: 0x00000080
```

```
Gateway Handle: 0x7fa7988c7ce8
```

```
Gateway: reachable, non-Connected route, prefix length 24
```

```
Resolving Route: 10.0.0.0/24 (static)
```

```
Paths: 0
```

```
RIB Nexthop ID: 0x0
```

```
Status: [Reachable][Connected][Not Local]
```

```
Metric: 0
```

```
ORR afi bits: 0x0
```

```
Registration: Synchronous, Completed: 01:22:27
```

```
Events: Critical (0)/Non-critical (0)
```

```
Last Received: 01:22:27 (Registration)
```

```
Last gw update: (Crit-sync) 01:22:27(rib)
```

```
Reference Count: 4
```

```
Prefix Related Information
```

```
Active Tables: [IPv4 Unicast][VPNv4 Unicast]
```

```
Metrics: [0x0][0x0]
```

```
Reference Counts: [0][4]
```

```
Interface Handle: 0x0
```

```
Attr ref-count: 7
```

```
SR policy color 101, State: [Up]
```

```
Not registered, bsid 24009
```

```
Skip Reg on restart [No]
```

```
First notif received [Yes]
```

```
SR Policy Flags [0x2]
```

```
BGP TE registered [No]
```

```
ODN registered [No]
```

```
IPv6 capability required/enabled: Yes/Yes
```

```
Last SR policy update: 01:22:35
```

如果SR策略用于验证下一跳，则您会看到以下输出：

```
RP/0/RP0/CPU0:PE1#show bgp vpnv4 unicast nexthops 10.0.0.5 color 101
```

```
Nexthop: 10.0.0.5 C:101
```

```
VRF: default
```

```
Nexthop ID: 0x6000008, Version: 0x0
```

```
Nexthop Flags: 0x00480000
```

```
Nexthop Handle: 0x7fa734042e94
```

```
RIB Related Information:
```

```
Firsthop interface handle 0x00000000
```

```
Gateway TBL Id: 0xe0000000 Gateway Flags: 0x00000080
```

```
Gateway Handle: 0x7fa7988c7ce8
```

```
Gateway: unreachable, non-Connected route, prefix length 8192
```

```
Resolving Route: 10.0.0.0/24 (static)
```

```
Paths: 0
```

```
RIB Nexthop ID: 0x0
```

```
Status: [Unreachable]
```

```
Metric: 4294967295
```

```
ORR afi bits: 0x0
```

```
Registration: Synchronous, Completed: 01:25:30
```

Events: Critical (1)/Non-critical (0)  
Last Received: 00:00:43 (Critical)  
Last gw update: (Crit-notif) 00:00:43(rib)  
Reference Count: 2

#### Prefix Related Information

Active Tables: [IPv4 Unicast][VPNv4 Unicast]  
Metrics: [0xffffffff][0xffffffff]  
Reference Counts: [0][2]

Interface Handle: 0x0

Attr ref-count: 5

SR policy color 101, State: [Up]

Not registered, bsid 24009

Skip Reg on restart [No]

First notif received [Yes]

SR Policy Flags [0x2]

BGP TE registered [No]

ODN registered [No]

IPv6 capability required/enabled: Yes/Yes

Last SR policy update: 01:25:38

## 配置

以下配置命令是新命令：

```
nexthop validation color-extcomm sr-policy  
nexthop validation color-extcomm disable  
bgp bestpath igp-metric sr-policy  
bgp bestpath sr-policy prefer  
bgp bestpath sr-policy force  
nexthop validation color-extcomm disable
```

## BGP ODN AS NextHop软验证命令

在PE(HE)上：

```
RP/0/RP0/CPU0:PE1(config)#router bgp 65001
```

```
RP/0/RP0/CPU0:PE1(config-bgp)#nexthop ?
```

```
mpls          Configure next-hop related items for mpls
```

```
resolution    Configure next-hop related items for resolution
```

```
validation    Configure next-hop reachability validation
```

```
RP/0/RP0/CPU0:PE1(config-bgp)#nexthop validation ?
```

```
color-extcomm Configure next-hop reachability validation for color-extcomm paths
```

```
RP/0/RP0/CPU0:PE1(config-bgp)#nexthop validation color-extcomm ?
```

```
disable       Disable next-hop reachability validation for color-extcomm paths
```

```
sr-policy     Enable BGP next-hop reachability validation by SR Policy for color-extcomm paths
```

```
RP/0/RP0/CPU0:PE1(config-bgp)#nexthop validation color-extcomm sr-policy
```

```
RP/0/RP0/CPU0:PE1(config-bgp)#commit
```

以下是主命令：它打开BGP软下一跳行为。如果SR策略为下一跳和颜色启用，则不执行RIB验证。

## BGP ODN AS下一跳硬验证命令

BGP硬下一跳是默认行为。

此命令是恢复为此行为的命令：`no nexthop validation color-extcomm`。

## 考虑SR策略路径度量的BGP最佳路径选择

当我们将到NH的内部网关协议(IGP)可达性，且算法在BGP最佳路径选择过程中达到步骤8时，首选BGP路径是到下一跳距离最低(IGP)的路径。这是默认行为。请参阅[BGP最佳路径选择算法](#)。

除非配置了**`bgp bestpath igp-metric ignore`**命令，否则此情况为真。在这种情况下，根本不考虑IGP开销。

目前，仅考虑到BGP NH的IGP度量；不是SR策略路径提供的度量。这仍是默认行为，但有命令指示BGP使用SR策略路径度量，而不是BGP最佳路径选择算法的IGP度量。

```
RP/0/RP0/CPU0:PE1(config)#router bgp 65001
RP/0/RP0/CPU0:PE1(config-bgp)#bgp bestpath igp-metric ?
  ignore      Ignore IGP metric during path comparison
  sr-policy   Use next-hop admin/metric from SR policy at Next Hop metric comparison stage
```

```
RP/0/RP0/CPU0:PE1(config-bgp)#bgp bestpath igp-metric sr-policy
RP/0/RP0/CPU0:PE1(config-bgp)#commit
```

此命令将启用PCE/路径管理和度量值的考虑。只有在SR策略启用时，这些管理员/度量值才能传递到BGP。此命令使BGP算法能够根据SR策略中下一跳的管理和度量来选择最佳路径。如果没有此命令，默认行为是仅考虑下一跳的IGP度量。这称为“下一跳的RIB验证”。

## BGP最佳路径选择优先选择SR策略路径

有些平台不支持具有本地下一跳或SR策略下一跳的路径混合。该平台可能不支持在两种路径类型上进行转发时混合使用这一功能。考虑到使用等价多路径(ECMP)或不等价多路径(UCMP)或备份路径，这一点非常重要。任何类型的路径都可以是BGP的最佳路径。默认行为是仅考虑与BGP最佳路径具有相同下一跳类型的BGP路径。

此命令指示BGP在路由器执行最佳路径计算时首选具有颜色/下一跳SR策略的路由。这意味着在最佳路径计算过程中不考虑SR策略关闭或没有SR策略的路径。

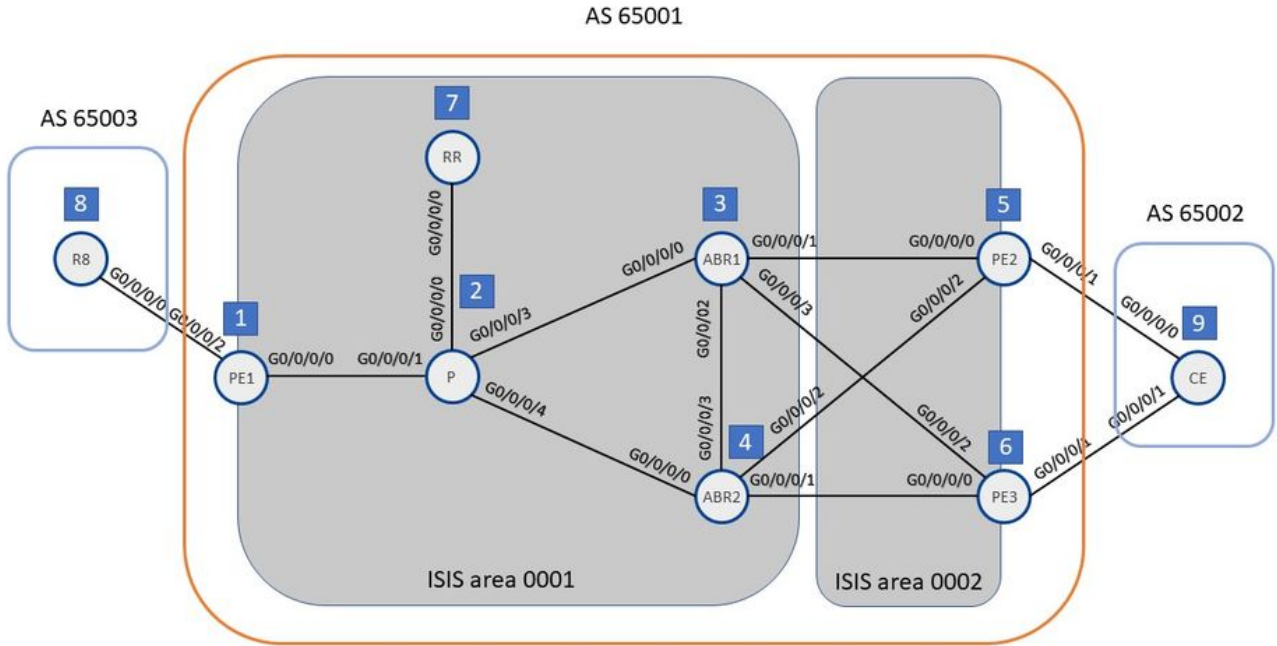
### `bgp bestpath sr-policy {force |首选}`

必须配置两个关键字之一。

```
RP/0/RP0/CPU0:PE1(config-bgp)#bgp bestpath sr-policy ?
  force    Consider only paths over SR Policy for bestpath selection, eBGP no-color ineligible
  prefer   Consider only paths over SR Policy for bestpath selection, eBGP no-color eligible
```

如果配置首选选项，则无颜色的eBGP路径将标记为符合条件（因此可以成为最佳路径的一部分）。如果不需要此行为，您可以向eBGP路径添加虚拟SR策略。否则，您可以配置此命令的force选项，以使没有颜色的eBGP路由不符合要求。

如图所示，参考网络。



从路由器PE1到网络10.99.99.99/32有三条可能的路径。前缀10.99.99.99/32由R8和CE路由器通告。

BGP有3条路由10.99.99.99/32的路径：2个iBGP（PE2和PE3是BGP下一跳路由器）和1个eBGP路径（来自R8）。

iBGP路径具有下一跳10.0.0.5和10.0.0.6。eBGP路径具有下一跳10.1.8.8。

配置没有此命令**bgp bestpath sr-policy**。

```
RP/0/RP0/CPU0:PE1#show bgp vrf one 10.9.9.9/32
BGP routing table entry for 10.9.9.9/32, Route Distinguisher: 65000:1
Versions:
  Process          bRIB/RIB   SendTblVer
  Speaker          474       474
  Local Label: 24005
Last Modified: Nov 29 09:04:07.948 for 00:00:49
Paths: (3 available, best #3)
Advertised to PE peers (in unique update groups):
  10.0.0.4         10.0.0.3
Path #1: Received by speaker 0
Not advertised to any peer
65002
  10.0.0.5 C:101 (bsid:24007) (admin 20) (metric 23) from 10.0.0.3 (10.0.0.5)
  Received Label 24018
  Origin IGP, metric 0, localpref 100, valid, internal, group-best, imported
  Received Path ID 0, Local Path ID 0, version 0
  Extended community: Color:101 RT:65001:101
  Originator: 10.0.0.5, Cluster list: 10.0.0.3
  SR policy color 101, up, not-registered, bsid 24007
  Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 65001:2
Path #2: Received by speaker 0
Not advertised to any peer
65002
  10.0.0.6 from 10.0.0.4 (10.0.0.6)
```

```

Received Label 24004
Origin IGP, metric 0, localpref 100, valid, internal, imported
Received Path ID 0, Local Path ID 0, version 0
Extended community: RT:65001:101
Originator: 10.0.0.6, Cluster list: 10.0.0.4
Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 65001:3
Path #3: Received by speaker 0
Advertised to PE peers (in unique update groups):
 10.0.0.4      10.0.0.3
65003
10.1.8.8 from 10.1.8.8 (10.0.0.8)
Origin IGP, metric 0, localpref 100, valid, external, best, group-best, import-candidate
Received Path ID 0, Local Path ID 1, version 474
Extended community: RT:65001:101
Origin-AS validity: (disabled)

```

eBGP路径没有颜色或SR策略。这是最佳路径。

如果eBGP路由确实有颜色，但没有SR策略，则仍会选择它作为最佳路径。

如果eBGP路由确实有颜色和SR策略，则选择它作为最佳路径。

这是另一个例子。eBGP路由没有颜色，并且没有配置SR策略和**bgp bestpath sr-policy prefer**命令。

**注意：**eBGP邻居在VRF内。这意味着您必须在VRF下配置**bgp bestpath sr-policy prefer**命令。

```

router bgp 65001
next-hop validation color-extcomm sr-policy
bgp unsafe-ebgp-policy
bgp bestpath igp-metric sr-policy
address-family vpnv4 unicast
!
neighbor 10.0.0.3
remote-as 65001
update-source Loopback0
address-family vpnv4 unicast
!
!
neighbor 10.0.0.4
remote-as 65001
update-source Loopback0
address-family vpnv4 unicast
!
!
neighbor 10.0.0.7
remote-as 65001
shutdown
update-source Loopback0
address-family vpnv4 unicast
!
!
vrf one
rd 65000:1
bgp unsafe-ebgp-policy
bgp bestpath sr-policy prefer
address-family ipv4 unicast
redistribute connected
!

```



```
neighbor 10.1.8.8
  remote-as 65003
  address-family ipv4 unicast
!
```

RP/0/RP0/CPU0:PE1#show bgp vrf one 10.9.9.9/32 bestpath-compare

BGP routing table entry for 10.9.9.9/32, Route Distinguisher: 65000:1

Versions:

```
Process          bRIB/RIB  SendTblVer
Speaker          579      579
```

Local Label: 24004 (no rewrite);

Flags: 0x01343001+0x00020000;

Last Modified: Nov 30 07:36:55.948 for 00:03:05

Paths: (3 available, best #3)

Advertised to PE peers (in unique update groups):

```
10.0.0.4      10.0.0.3
```

Path #1: Received by speaker 0

Flags: 0x2000000001020005, import: 0x080

Not advertised to any peer

65002

10.0.0.5 C:101 (bsid:24007) (admin 20) (metric 23) from 10.0.0.3 (10.0.0.5), if-handle 0x00000000

Received Label 24018

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

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

Extended community: Color:101 RT:65001:101

Originator: 10.0.0.5, Cluster list: 10.0.0.3

SR policy color 101, up, not-registered, bsid 24007

best of AS 65002

An iBGP path, whereas best path (path #3) is an eBGP path

Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 65001:2

Path #2: Received by speaker 0

Flags: 0x2000000000020005, import: 0x0a0

Not advertised to any peer

65002

10.0.0.6 from 10.0.0.4 (10.0.0.6), if-handle 0x00000000

Received Label 24004

Origin IGP, metric 0, localpref 100, valid, internal, imported

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

Extended community: RT:65001:101

Originator: 10.0.0.6, Cluster list: 10.0.0.4

Non SR-policy path is ignored due to config knob

Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 65001:3

Path #3: Received by speaker 0

Flags: 0x300000000d040003, import: 0x31f

Advertised to PE peers (in unique update groups):

```
10.0.0.4      10.0.0.3
```

65003

10.1.8.8 from 10.1.8.8 (10.0.0.8), if-handle 0x00000000

Origin IGP, metric 0, localpref 100, valid, external, **best**, group-best, import-candidate

Received Path ID 0, Local Path ID 1, version 579

Extended community: RT:65001:101

Origin-AS validity: (disabled)

best of AS 65003, Overall best

eBGP路径是最佳路径，即使它没有颜色。如果不希望eBGP路由不带颜色作为最佳路径，则使用force选项配置bgp bestpath sr-policy命令。

**注意：**本地路径和重分发的路径始终符合最佳路径计算的条件。

使用此命令检查平台是否支持通过SR策略转发和本地下一跳的混合。

```
RP/0/RP0/CPU0:R1#show bgp process detail | include native
Platform support mix of sr-policy and native nexthop: No
```

**注意：**路由器NCS55xx和NCS560/NCS540显示no，而ASR9000显示yes。

## BGP最佳路径选择强制SR策略路径

该命令指示BGP在执行最佳路径计算时首选具有SR策略下一跳的路由，但不包括无颜色的eBGP路径。

```
RP/0/RP0/CPU0:PE1(config-bgp)#bgp bestpath sr-policy ?
 force    Consider only paths over SR Policy for bestpath selection, eBGP no-color ineligible
RP/0/RP0/CPU0:PE1(config-bgp)#bgp bestpath sr-policy force ?
```

## 禁用颜色扩展通信路径的下一跳可达性验证

这通常用于路由反射器(RR)。

在RR上：

```
RP/0/RP0/CPU0:RR1(config-bgp)#nexthop validation color-extcomm disable
RP/0/RP0/CPU0:RR1(config-bgp)#commit
```

禁用颜色扩展通信路径的下一跳可达性验证。这与SR策略的状态或存在无关。

## 概述

头端和RR上的行为由下一跳验证命令和bgp best path igp-metric sr-policy命令的配置驱动。有4种方案。每个场景都有两个配置命令的组合。

## 默认行为

适用于头端路由器和路由反射器。

配置:

```
no nexthop validation color-extcomm sr-policy
no bgp bestpath igp-metric sr-policy
```

功能:

```
Perform RIB validation (hard next-hop).
Do not use admin/metric from the sr-policy.
```

## 使用SR策略度量进行RIB相关验证

适用于头端路由器和路由反射器。

配置:

```
no nexthop validation color-extcomm sr-policy
bgp bestpath igp-metric sr-policy
```

功能:

```
Perform RIB validation (hard next-hop).
If NH is reachable in RIB:
  If policy is up:
    Use policy metric
  If policy is down:
    Use RIB metric
```

## SR策略相关验证与RIB度量

这是默认行为。

适用于头端路由器。

配置:

```
nexthop validation color-extcomm sr-policy
no bgp bestpath igp-metric sr-policy
```

功能:

```
Do not perform RIB validation (soft next-hop).
Do not use admin/metric from the SR policy.
The RIB metric might not be available.
```

## SR策略相关验证与SR策略度量

适用于头端路由器。

配置:

```
nexthop validation color-extcomm sr-policy
bgp bestpath igp-metric sr-policy
```

功能:

```
Do not perform RIB validation (soft next-hop). RIB reachability is not needed.
If policy is up:
  Use policy metric and validation, even if RIB reachability is present
If policy is down:
  Use RIB validation and metric if available. If not available, the route is not installed.
```

## SR策略相关验证，RIB度量和SR策略不用于最佳路径计算

适用于RR路由器。

配置:

```
nexthop validation color-extcomm disable
no bgp bestpath igp-metric sr-policy
```

功能:

Use RIB metric if the next-hop is in the RIB. Else, use the gateway metric (the next-hop IGP metric) 0.

Do not use SR policy for bestpath calculation. Do not use admin/metric from the SR policy.

## SR策略相关验证，使用RIB度量和SR策略进行最佳路径计算

适用于RR路由器。

配置:

```
nexthop validation color-extcomm disable
bgp bestpath igp-metric sr-policy
```

功能:

Use RIB metric if the next-hop is in the RIB. Else, use the gateway metric 0.  
Use sr-policy for bestpath calculation.

If policy is up:

Use policy metric and validation, even if RIB reachability is present

If policy is down

Use RIB validation and metric if available

If RIB validation and metric is not available:

use the gateway metric 0

## 确认

这是验证哪种下一跳验证处于活动状态以及SR策略的管理距离/度量是否在最佳路径计算期间使用的方式。

```
RP/0/RP0/CPU0:PE1#show bgp process detail | i Nexthop
```

```
Use SR-Policy admin/metric of color-extcomm Nexthop during path comparison: enabled
```

```
ExtComm Color Nexthop validation: SR-Policy then RIB
```

这是默认设置。

这是SR策略相关验证的示例，其中RIB度量和SR策略不用于最佳路径计算。

```
RP/0/RP0/CPU0:PE1#show bgp vpnv4 unicast process detail | in Nexthop
```

```
Use SR-Policy admin/metric of color-extcomm Nexthop during path comparison: disabled
```

```
ExtComm Color Nexthop validation: RIBExtComm Color Nexthop validation: RIB
```

这是附加到BGP路由的管理距离/度量的示例。

```

RP/0/RP0/CPU0:PE1#show bgp vrf VRF1002 ipv4 unicast 10.77.2.0
BGP routing table entry for 10.77.2.0/24, Route Distinguisher: 18522:1002
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          5232243   5232243
Paths: (1 available, best #1)
  Advertised to CE peers (in unique update groups):
    10.11.2.11     10.15.2.2
  Path #1: Received by speaker 0
  Advertised to CE peers (in unique update groups):
    10.11.2.11     10.15.2.2
16611 770
  10.1.1.33 C:1129 (bsid:27163) (admin 20) (metric 25) from 10.1.1.100 (10.1.1.33)
    Received Label 24007
    Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate, imported
    Received Path ID 1, Local Path ID 1, version 5232243
    Extended community: Color:1129 RT:17933:1002 RT:18522:1002
    Originator: 10.1.1.33, Cluster list: 10.1.1.100
    SR policy color 1129, up, registered, bsid 27163, if-handle 0x200053dc
    Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 18522:3002

```

## 检查策略是打开还是关闭

这是验证SR策略是打开还是关闭的方式。

```

RP/0/RP0/CPU0:PE1#show segment-routing traffic-eng pcc lsp
PCC's SR policy database:
-----
Symbolic Name: cfg_ODN-policy-1_discr_100
LSP[0]:
Source 10.0.0.1, Destination 10.0.0.5, Tunnel ID 3, LSP ID 8
State: Admin up, Operation up
Setup type: SR
Binding SID: 24005

```

## 验证是否使用策略

使用BGP show命令查看路由。

如果有绑定网段Identifier(BSID), 则此路由使用SR策略。

```

RP/0/RP0/CPU0:PE1#show bgp vrf one 10.0.0.9/32
BGP routing table entry for 10.0.0.9/32, Route Distinguisher: 65000:1
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          89        89
Last Modified: Oct 28 13:21:57.714 for 00:00:30
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
65002
  10.0.0.5 C:101 (bsid:24004) from 10.0.0.3 (10.0.0.5)
    Received Label 24002
    Origin IGP, metric 0, localpref 100, valid, internal, best, group-best, import-candidate,
imported
    Received Path ID 0, Local Path ID 1, version 87
    Extended community: Color:101 RT:65001:101
    Originator: 10.0.0.5, Cluster list: 10.0.0.3

```

SR policy color 101, up, not-registered, bsid 24004

Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 65001:2

绑定SID是此处的MPLS标签。此标签链接到一个SR策略。

RP/0/RP0/CPU0:PE1#show mpls forwarding labels 24004

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched
24004	Pop	No ID	srte_c_101_e	point2point	0

## Show BGP Nexthops命令

您可以使用show bgp nexthops命令查看终端的颜色、管理和度量。

RP/0/RP0/CPU0:RR#show bgp nexthops wide

Total Nexthop Processing  
Time Spent: 0.000 secs

Maximum Nexthop Processing

Received: 00:21:57  
Bestpaths Deleted: 0  
Bestpaths Changed: 31  
Time Spent: 0.000 secs

Last Notification Processing

Received: 00:01:22  
Time Spent: 0.000 secs

Gateway Address Family: IPv4 Unicast

Table ID: 0xe0000000

Gateway Reference Count: 8

Gateway AF Bits : 0x8011

Nexthop Count: 6

Critical Trigger Delay: 3000msec

Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1

EPE Table Version: 1, EPE Label version: 1

EPE Downloaded Version: 1, EPE Standby Version: 0

Status codes: R/UR Reachable/Unreachable

C/NC Connected/Not-connected

L/NL Local/Non-local

PR Pending Registration

I Invalid (Policy drop)

Next Hop	Status	Metric	Tbl-ID
10.0.0.1	[R][NC][NL]	30	
e0000000	6/0 00:01:22 (Cri)		0/5
10.0.0.3	[R][NC][NL]	20	e0000000
	6/0 00:01:22 (Cri)	0/34	
10.0.0.4	[R][NC][NL]	30	
e0000000	6/0 00:01:22 (Cri)		0/34
10.0.0.5	[UR]	4294967295	
e0000000	2/0 00:01:22 (Cri)		0/4
10.0.0.5 T:101	[UR]	4294967295	
e0000000	2/0 00:01:22 (Cri)		0/3
10.0.0.6	[UR]	4294967295	
e0000000	2/0 00:01:22 (Cri)		0/3

RP/0/RP0/CPU0:PE1#show bgp vpnv4 unicast nexthops 10.0.0.5 color 101

**Nexthop: 10.0.0.5 C:101**

VRF: default

Nexthop ID: 0x6000006, Version: 0x0

Nexthop Flags: 0x00480002

Nexthop Handle: 0x7efc84043624

RIB Related Information:

Firsthop interface handle 0x0000000c

Gateway TBL Id: 0xe0000000 Gateway Flags: 0x00000080

Gateway Handle: 0x7efcadede6e98

Gateway: reachable, non-Connected route, prefix length 8

Resolving Route: 10.0.0.0/8 (static)

Paths: 0

RIB Nexthop ID: 0x0

**Status: Reachable via SR-TE**

Status: [Reachable][Connected][Not Local]

**Metric: 0 (SR-TE metric 333)**

ORR afi bits: 0x0

Registration: Asynchronous, Completed: 2d05h

Events: Critical (14)/Non-critical (0)

Last Received: 02:15:15 (Critical)

Last gw update: (Crit-notif) 02:15:15(rib)

Reference Count: 2

Prefix Related Information

Active Tables: [IPv4 Unicast][VPNv4 Unicast]

Metrics: [0x0][0x0]

Reference Counts: [0][2]

Interface Handle: 0x0

Attr ref-count: 5

**SR policy color 101, State: [Up]**

Not registered, bsid 24004

Skip Reg on restart [No]

First notif received [Yes]

SR Policy Flags [0x2]

BGP TE registered [No]

ODN registered [No]

**End-point admin/metric: 30/333**

IPv6 capability required/enabled: Yes/Yes

Last SR policy update: 00:55:07

## BGP跟踪

show bgp trace输出中的某些条目引用SR策略。注意管理员/度量的存在。

```
default-bgp/spkr-tr2-sr 0/RP0/CPU0 t8885 [SR]:1323: SR-policy hdlr for reg nh with XTC af 0,
reg/unreg flag 1
default-bgp/spkr-tr2-sr 0/RP0/CPU0 t8885 [SR]:3394: SR-policy XTC nexthop 10.0.0.5/32 T:, color
101, register 1 with XTC done, v6-cap 1, rc 'Success', flags 0x480000
default-bgp/spkr-tr2-sr 0/RP0/CPU0 t8885 [SR]:3394: SR-policy XTC nexthop 10.0.0.6/32 T:, color
101, register 1 with XTC done, v6-cap 0, rc 'Success', flags 0x480000
default-bgp/spkr-tr2-sr 0/RP0/CPU0 t8885 [SR]:2424: SR-policy XTC notif NH end-point
color,gw_afi 0, [C:101][10.0.0.5] admin/metric 100/2147483647
default-bgp/spkr-tr2-sr 0/RP0/CPU0 t8885 [SR]:2424: SR-policy XTC notif NH end-point
color,gw_afi 0, [C:101][10.0.0.5] admin/metric 100/2147483647
default-bgp/spkr-tr2-sr 0/RP0/CPU0 t8885 [SR]:2424: SR-policy XTC notif NH end-point
color,gw_afi 0, [C:101][10.0.0.5] admin/metric 20/30
default-bgp/spkr-tr2-sr 0/RP0/CPU0 t8881 [SR]:1379: SR-policy trigger XTC for nh reg af 0,
reg/unreg flag 1
default-bgp/spkr-tr2-nh 0/RP0/CPU0 t8885 [NH]:7370: nexthop walk for AFI:'VPNv4 Unicast' start
default-bgp/spkr-tr2-nh 0/RP0/CPU0 t8885 [NH]:7425: nexthop walk for AFI:'VPNv4 Unicast', paths
```

deleted: 0, recalculated bestpaths: 2, color nh trigger for 2 nets, 0 msec

**注意：**思科IOS XR流量控制器(XTC)是指SR控制器。

BGP跟踪中的某些条目指与下一跳处理相关的配置更改。

```
default-bgp/spkr-tr2-prog 0/RP0/CPU0 t9036 [PROG]:724: 'Done VRF cfg notif init', name default iid 0
default-bgp/spkr-tr2-prog 0/RP0/CPU0 t9036 [PROG]:792: 'Done cfg init', name default iid 0
default-bgp/spkr-tr2-gen 0/RP0/CPU0 t9048 [GEN]:17871: nh cfg change 2 sense 1
default-bgp/spkr-tr2-gen 0/RP0/CPU0 t9048 [GEN]:17920: nh cfg change 1 sense 1
```

## 管理距离和度量

管理距离(admin)由SR策略中的度量类型确定。可以在头端路由器上设置度量类型。

```
RP/0/RP0/CPU0:PE1#conf t
RP/0/RP0/CPU0:PE1(config)#segment-routing
RP/0/RP0/CPU0:PE1(config-sr)#traffic-eng
RP/0/RP0/CPU0:PE1(config-sr-te)#policy ODN-policy-1
RP/0/RP0/CPU0:PE1(config-sr-te-policy)#color 101 end-point ipv4 10.0.0.5
RP/0/RP0/CPU0:PE1(config-sr-te-policy)#candidate-paths
RP/0/RP0/CPU0:PE1(config-sr-te-policy-path)#preference 100
RP/0/RP0/CPU0:PE1(config-sr-te-policy-path-pref)#dynamic
RP/0/RP0/CPU0:PE1(config-sr-te-pp-info)#metric ?
  margin      Metric margin
  sid-limit   SID limit
  type        Metric type configuration
  <cr>

RP/0/RP0/CPU0:PE1(config-sr-te-pp-info)#metric type ?
  hopcount   Hopcount metric type
  igp        IGP metric type
  latency    Latency metric type
  te         TE metric type
```

## 管理员值

这些是默认SR策略管理员值。

- 延迟 10
- TE 20
- IGP (默认) 30
- hopcount 40
- NONE/UNKNOWN度量类型 (对于显式段列表策略) 100

如果度量类型为none，则度量值为1。

管理员值越低，通向BGP的路径越优先。

度量越低，如果管理员具有相同值，则路径越首选BGP。

## 验证ODN中的管理和度量类型



```
RP/0/RP0/CPU0:PE1#show segment-routing traffic-eng policy color 101 endpoint ipv4 10.0.0.5
```

```
SR-TE policy database
```

```
-----
```

```
Color: 101, End-point: 10.0.0.5
```

```
Name: srte_c_101_ep_10.0.0.5
```

```
Status:
```

```
Admin: up Operational: up for 01:01:00 (since Oct 28 15:22:36.012)
```

```
Candidate-paths:
```

```
Preference: 100 (configuration) (active)
```

```
Name: ODN-policy-1
```

```
Requested BSID: dynamic
```

```
PCC info:
```

```
Symbolic name: cfg_ODN-policy-1_discr_100
```

```
PLSP-ID: 4
```

```
Protection Type: protected-preferred
```

```
Maximum SID Depth: 10
```

```
Dynamic (pce 10.0.0.7) (valid)
```

```
Metric Type: IGP, Path Accumulated Metric: 30
```

```
16002 [Prefix-SID, 10.0.0.2]
```

```
24009 [Adjacency-SID, 10.2.3.2 - 10.2.3.3]
```

```
16005 [Prefix-SID, 10.0.0.5]
```

```
Attributes:
```

```
Binding SID: 24004
```

```
Forward Class: Not Configured
```

```
Steering labeled-services disabled: no
```

```
Steering BGP disabled: no
```

```
IPv6 caps enable: yes
```

```
Invalidation drop enabled: no
```

## 有效指标

有效度量强制配置此命令的策略的类型和度量。

```
RP/0/RP0/CPU0:PE1#conf t
```

```
RP/0/RP0/CPU0:PE1(config)#segment-routing
```

```
RP/0/RP0/CPU0:PE1(config-sr)#traffic-eng
```

```
RP/0/RP0/CPU0:PE1(config-sr-te)#policy ODN-policy-1
```

```
RP/0/RP0/CPU0:PE1(config-sr-te-policy)#candidate-paths
```

```
RP/0/RP0/CPU0:PE1(config-sr-te-policy-path)#preference 100
```

```
RP/0/RP0/CPU0:PE1(config-sr-te-policy-path-pref)#effective-metric ?
```

```
value Metric value, advertised to other protocols
```

```
<cr>
```

```
RP/0/RP0/CPU0:PE1(config-sr-te-policy-path-pref)#effective-metric value 333 ?
```

```
type Metric type, advertised to other protocols
```

```
<cr>
```

```
RP/0/RP0/CPU0:PE1(config-sr-te-policy-path-pref)#effective-metric value 333 type ?
```

```
hopcount HOPCOUNT metric type
```

```
igp IGP metric type
```

```
latency LATENCY metric type
```

```
te TE metric type
```

```
RP/0/RP0/CPU0:PE1(config-sr-te-policy-path-pref)#effective-metric value 333 type igp ?
```

```
<cr>
```

```
RP/0/RP0/CPU0:PE1(config-sr-te-policy-path-pref)#effective-metric value 333 type igp
```

```
RP/0/RP0/CPU0:PE1(config-sr-te-policy-path-pref)#commit
```

```
RP/0/RP0/CPU0:PE1#show run segment-routing traffic-eng policy ODN-policy-1
```

```

segment-routing
traffic-eng
policy ODN-policy-1
color 101 end-point ipv4 10.0.0.5
candidate-paths
preference 100
dynamic
pcep
!
metric
type igp
!
!
effective-metric
value 333 type igp

```

您可以以这种方式验证应用的有效度量类型（管理距离）和度量值。

```

RP/0/RP0/CPU0:PE1#show bgp vrf one 10.0.0.9/32
BGP routing table entry for 10.0.0.9/32, Route Distinguisher: 65000:1
Versions:
Process          bRIB/RIB  SendTblVer
Speaker          131      131
Last Modified: Oct 28 15:22:35.714 for 00:03:42
Paths: (1 available, best #1)
Not advertised to any peer
Path #1: Received by speaker 0
Not advertised to any peer
65002
  10.0.0.5 C:101 (bsid:24004) (admin 30) (metric 333) from 10.0.0.7 (10.0.0.5)
    Received Label 24002
    Origin IGP, metric 0, localpref 100, valid, internal, best, group-best, import-candidate,
imported
    Received Path ID 0, Local Path ID 1, version 130
    Extended community: Color:101 RT:65001:101
    Originator: 10.0.0.5, Cluster list: 10.0.0.7, 10.0.0.3
    SR policy color 101, up, not-registered, bsid 24004
    Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 65001:2

```

## BGP路径比较

默认情况下，BGP路径的比较不会更改。

如果配置了**bgp bestpath igp-metric sr-policy**命令，则SR策略的管理距离和度量在BGP最佳路径选择算法中使用。

SR策略的管理距离和度量与SR策略关联。这是本地配置或通过PCEP（路径计算元素协议）从SR-PCE接收的。这意味着，如果路由反射器比较路径，它看不到管理距离和度量，因为它没有ODN的头端功能。因此，它没有到SR PCE的PCEP会话。

此示例显示一个远程PE路由器通告的前缀。这是配置。

```

segment-routing
global-block 16000 23999
traffic-eng
logging
policy status
!
policy ODN-policy-1

```

```

color 101 end-point ipv4 10.0.0.5
candidate-paths
  preference 100
  dynamic
    pcep
    !
    metric
      type te
    !
  !
  preference 200
  dynamic
    pcep
    !
    metric
      type te
    !

```

度量类型为TE。

此头端路由器看到带有颜色的前缀两次，且具有相同的TE度量，因为它对于两条路径是相同的BGP下一跳。

```

RP/0/RP0/CPU0:PE1#show bgp vpnv4 unicast rd 65001:2 10.0.0.9/32 bestpath-compare
BGP routing table entry for 10.0.0.9/32, Route Distinguisher: 65001:2
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          8         8
  Flags: 0x00040001+0x00010000;
Last Modified: Nov  2 09:21:55.948 for 00:00:32
Paths: (2 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Flags: 0xa000000025060005, import: 0x31f
  Not advertised to any peer
  65002
    10.0.0.5 C:101 (bsid:24018) (admin 20) (metric 23) from 10.0.0.3 (10.0.0.5), if-handle
0x00000000
  Received Label 24002
  Origin IGP, metric 0, localpref 100, valid, internal, best, group-best, import-candidate,
not-in-vrf
  Received Path ID 0, Local Path ID 1, version 8
  Extended community: Color:101 RT:65001:101
  Originator: 10.0.0.5, Cluster list: 10.0.0.3
  SR policy color 101, up, not-registered, bsid 24018
  best of AS 65002, Overall best
  Path #2: Received by speaker 0
  Flags: 0x2000000024020005, import: 0x000
  Not advertised to any peer
  65002
    10.0.0.5 C:101 (bsid:24018) (admin 20) (metric 23) from 10.0.0.4 (10.0.0.5), if-handle
0x00000000
  Received Label 24002
  Origin IGP, metric 0, localpref 100, valid, internal, import-candidate, not-in-vrf
  Received Path ID 0, Local Path ID 0, version 0
  Extended community: Color:101 RT:65001:101
  Originator: 10.0.0.5, Cluster list: 10.0.0.4, 10.0.0.7, 10.0.0.3
  SR policy color 101, up, not-registered, bsid 24018
  Longer cluster length than best path (path #1)

```

由于管理距离和度量对于两条路径都相同，因此在BGP最佳路径选择算法中，将进一步决定哪条路径是最佳路径。

此示例显示两个远程PE路由器通告的前缀。一条路径的下一跳为10.0.0.5，另一条路径的下一跳为10.0.0.6。前缀的颜色为101，来自两个远程PE路由器。头端路由器PE1对此颜色有两个ODN策略。

```
segment-routing
global-block 16000 23999
traffic-eng
logging
policy status
!
policy ODN-policy-1
color 101 end-point ipv4 10.0.0.5
candidate-paths
preference 100
dynamic
pcep
!
metric
type igp
!
!
!
preference 200
dynamic
pcep
!
metric
type te
!
!
!
!
policy ODN-policy-2
color 101 end-point ipv4 10.0.0.6
candidate-paths
preference 100
dynamic
pcep
!
metric
type igp
!
```

终端10.0.0.5的策略使用度量类型TE，而终端10.0.0.6的策略使用度量类型IGP。

```
RP/0/RP0/CPU0:PE1#show bgp vpnv4 unicast vrf one 10.0.0.9/32 bestpath-compare
BGP routing table entry for 10.0.0.9/32, Route Distinguisher: 65000:1
Versions:
Process          bRIB/RIB  SendTblVer
Speaker          25        25
  Flags: 0x00043001+0x00000000;
Last Modified: Nov  1 11:42:28.948 for 00:43:41
Paths: (2 available, best #1)
Not advertised to any peer
Path #1: Received by speaker 0
Flags: 0xa000000005060005, import: 0x080
Not advertised to any peer
65002
  10.0.0.5 C:101 (bsid:24007) (admin 20) (metric 30) from 10.0.0.4 (10.0.0.5), if-handle
0x00000000
```

```
Received Label 24002
Origin IGP, metric 0, localpref 100, weight 65000, valid, internal, best, group-best,
import-candidate, imported
Received Path ID 0, Local Path ID 1, version 25
Extended community: Color:101 RT:65001:101
Originator: 10.0.0.5, Cluster list: 10.0.0.4, 10.0.0.7, 10.0.0.3
SR policy color 101, up, not-registered, bsid 24007
best of AS 65002, Overall best
Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 65001:2
Path #2: Received by speaker 0
Flags: 0x2000000000020005, import: 0x0a0
Not advertised to any peer
65002
```

```
10.0.0.6 C:101 (bsid:24012) (admin 30) (metric 30) from 10.0.0.4 (10.0.0.6), if-handle
0x00000000
```

```
Received Label 24002
Origin IGP, metric 0, localpref 100, weight 65000, valid, internal, imported
Received Path ID 0, Local Path ID 0, version 0
Extended community: Color:101 RT:65001:101
Originator: 10.0.0.6, Cluster list: 10.0.0.4
SR policy color 101, up, not-registered, bsid 24012
Higher nexthop admin distance than best path (path #1)
Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 65001:3
```

最佳路径是第一条，因为它的管理距离比第二条路径的管理距离低。度量类型TE的管理距离低于度量类型IGP的管理距离。

ODN-policy-1的SR策略优先级为200,ODN-policy-2的SR策略优先级为100。

```
RP/0/RP0/CPU0:PE1#show segment-routing traffic-eng pcc lsp detail
```

```
PCC's SR policy database:
```

```
-----
Symbolic Name: cfg_ODN-policy-1_discr_100
LSP[0]:
Source 10.0.0.1, Destination 10.0.0.5, Tunnel ID 1, LSP ID 0
State: Admin up, Operation down
Setup type: SR
Bandwidth: requested 0, used 0
LSP object:
  PLSP-ID 0x1, flags: D:0 S:0 R:0 A:1 O:0 C:0
Metric type: IGP, Accumulated Metric 30
ERO:
  SID[0]: Node, Label 16004, NAI: 10.0.0.4
  SID[1]: Node, Label 16005, NAI: 10.0.0.5
```

```
Symbolic Name: cfg_ODN-policy-1_discr_200
LSP[0]:
Source 10.0.0.1, Destination 10.0.0.5, Tunnel ID 1, LSP ID 4
State: Admin up, Operation up
Binding SID: 24007
Setup type: SR
Bandwidth: requested 0, used 0
LSP object:
  PLSP-ID 0x2, flags: D:0 S:0 R:0 A:1 O:1 C:0
Metric type: TE, Accumulated Metric 30
ERO:
  SID[0]: Adj, Label 24001, NAI: local 10.1.2.1 remote 10.1.2.2
  SID[1]: Adj, Label 24003, NAI: local 10.2.3.2 remote 10.2.3.3
  SID[2]: Node, Label 16005, NAI: 10.0.0.5
```

```
Symbolic Name: cfg_ODN-policy-2_discr_100
LSP[0]:
Source 10.0.0.1, Destination 10.0.0.6, Tunnel ID 2, LSP ID 2
```

```
State: Admin up, Operation up
Binding SID: 24012
Setup type: SR
Bandwidth: requested 0, used 0
LSP object:
  PLSP-ID 0x3, flags: D:0 S:0 R:0 A:1 O:1 C:0
Metric type: IGP, Accumulated Metric 30
ERO:
  SID[0]: Node, Label 16004, NAI: 10.0.0.4
  SID[1]: Node, Label 16006, NAI: 10.0.0.6
```

以下是管理距离相同但度量不同的示例。

```
RP/0/RP0/CPU0:PE1#show bgp vpnv4 unicast vrf one 10.0.0.9/32 bestpath-compare
BGP routing table entry for 10.0.0.9/32, Route Distinguisher: 65000:1
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          57        57
  Flags: 0x00043001+0x00010000;
Last Modified: Nov  2 07:54:20.948 for 00:00:04
Paths: (2 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Flags: 0xa000000005060005, import: 0x080
  Not advertised to any peer
  65002
    10.0.0.5 C:101 (bsid:24007) (admin 30) (metric 23) from 10.0.0.4 (10.0.0.5), if-handle
    0x00000000
      Received Label 24002
      Origin IGP, metric 0, localpref 100, weight 65000, valid, internal, best, group-best,
import-candidate, imported
      Received Path ID 0, Local Path ID 1, version 39
      Extended community: Color:101 RT:65001:101
      Originator: 10.0.0.5, Cluster list: 10.0.0.4, 10.0.0.7, 10.0.0.3
      SR policy color 101, up, not-registered, bsid 24007
      best of AS 65002, Overall best
      Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 65001:2
  Path #2: Received by speaker 0
  Flags: 0x2000000004020005, import: 0x080
  Not advertised to any peer
  65002
    10.0.0.6 C:101 (bsid:24012) (admin 30) (metric 30) from 10.0.0.4 (10.0.0.6), if-handle
    0x00000000
      Received Label 24002
      Origin IGP, metric 0, localpref 100, weight 65000, valid, internal, import-candidate,
imported
      Received Path ID 0, Local Path ID 0, version 0
      Extended community: Color:101 RT:65001:101
      Originator: 10.0.0.6, Cluster list: 10.0.0.4
      SR policy color 101, up, not-registered, bsid 24012
Higher IGP metric than best path (path #1)
      Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 65001:3
```

这是带有metric-type hopcount的示例。

```
RP/0/RP0/CPU0:PE1#show bgp vpnv4 unicast vrf one 10.0.0.9/32 bestpath-compare
BGP routing table entry for 10.0.0.9/32, Route Distinguisher: 65000:1
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          99        99
  Flags: 0x00043001+0x00010000;
Last Modified: Nov  2 08:21:19.948 for 00:00:41
```

```

Paths: (2 available, best #2)
Not advertised to any peer
Path #1: Received by speaker 0
Flags: 0x2000000004020005, import: 0x080
Not advertised to any peer
65002
  10.0.0.5 C:101 (bsid:24007) (admin 40) (metric 4) from 10.0.0.4 (10.0.0.5), if-handle
0x00000000
  Received Label 24002
  Origin IGP, metric 0, localpref 100, weight 65000, valid, internal, import-candidate,
imported
  Received Path ID 0, Local Path ID 0, version 0
  Extended community: Color:101 RT:65001:101
  Originator: 10.0.0.5, Cluster list: 10.0.0.4, 10.0.0.7, 10.0.0.3
  SR policy color 101, up, not-registered, bsid 24007
  Higher IGP metric than best path (path #2)
  Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 65001:2
Path #2: Received by speaker 0
Flags: 0xa000000005060005, import: 0x080
Not advertised to any peer
65002
  10.0.0.6 C:101 (bsid:24010) (admin 40) (metric 3) from 10.0.0.4 (10.0.0.6), if-handle
0x00000000
  Received Label 24002
  Origin IGP, metric 0, localpref 100, weight 65000, valid, internal, best, group-best,
import-candidate, imported
  Received Path ID 0, Local Path ID 1, version 95
  Extended community: Color:101 RT:65001:101
  Originator: 10.0.0.6, Cluster list: 10.0.0.4
  SR policy color 101, up, not-registered, bsid 24010
  best of AS 65002, Overall best
  Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 65001:3

```

两个不同终端有两条相互竞争的BGP路径。BGP决定哪条路径成功并安装到路由表中。这进而根据颜色和终端确定要将流量转发到BGP VPNv4前缀而安装的SR策略。

## 比较BGP路径是否带颜色

在场景4中，在头端路由器上启用软下一跳验证，它接收两个BGP路径，一条是带前缀，一条是不带颜色。如果没有下一跳的路由，则无颜色的路径不可访问下一跳，并且不考虑进行安装。

```

RP/0/RP0/CPU0:PE1#show bgp vpnv4 unicast | include 10.0.0.9/32
*>i10.0.0.9/32      10.0.0.5 C:101      0    100    0 65002 i
*>i10.0.0.9/32      10.0.0.5 C:101      0    100    0 65002 i
* i10.0.0.9/32      10.0.0.6            0    100    0 65002 i

```

最后一个BGP路径没有>，因此下一跳不可访问。

```

RP/0/RP0/CPU0:PE1#show bgp vpnv4 unicast rd 65001:3 10.0.0.9/32
BGP routing table entry for 10.0.0.9/32, Route Distinguisher: 65001:3
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          31        31
Last Modified: Nov  2 10:08:44.948 for 00:08:11
Paths: (2 available, no best path)
Not advertised to any peer
Path #1: Received by speaker 0
Not advertised to any peer
65002
  10.0.0.6 (inaccessible) from 10.0.0.3 (10.0.0.6)

```

```

Received Label 24002
Origin IGP, metric 0, localpref 100, valid, internal, not-in-vrf
Received Path ID 0, Local Path ID 0, version 0
Extended community: RT:65001:101
Originator: 10.0.0.6, Cluster list: 10.0.0.3, 10.0.0.7, 10.0.0.4
Path #2: Received by speaker 0
Not advertised to any peer
65002
  10.0.0.6 (inaccessible) from 10.0.0.4 (10.0.0.6)
    Received Label 24002
    Origin IGP, metric 0, localpref 100, valid, internal, not-in-vrf
    Received Path ID 0, Local Path ID 0, version 0
    Extended community: RT:65001:101
    Originator: 10.0.0.6, Cluster list: 10.0.0.4

```

使用带SR策略的BGP路径。

但是，如果由于RIB中的路由而解析下一跳10.0.0.6，则此路径可被选为最佳路径。如果它没有颜色，则不能用于ODN，SR策略将关闭。但是，此路由的管理距离是100，因此它比带有颜色的路径高得多。

```

RP/0/RP0/CPU0:PE1#show bgp vrf one 10.0.0.9/32 bestpath-compare
BGP routing table entry for 10.0.0.9/32, Route Distinguisher: 65000:1
Versions:
Process          bRIB/RIB  SendTblVer
Speaker          47        47
  Flags: 0x00043001+0x00000000;
Last Modified: Nov  2 10:30:55.948 for 00:00:21
Paths: (2 available, best #1)
Advertised to CE peers (in unique update groups):
  10.1.8.8
Path #1: Received by speaker 0
Flags: 0xa000000005060005, import: 0x080
Advertised to CE peers (in unique update groups):
  10.1.8.8
65002
  10.0.0.5 C:101 (bsid:24021) (admin 20) (metric 23) from 10.0.0.3 (10.0.0.5), if-handle
0x00000000
  Received Label 24002
  Origin IGP, metric 0, localpref 100, valid, internal, best, group-best, import-candidate,
imported
  Received Path ID 0, Local Path ID 1, version 40
  Extended community: Color:101 RT:65001:101
  Originator: 10.0.0.5, Cluster list: 10.0.0.3
  SR policy color 101, up, not-registered, bsid 24021
  best of AS 65002, Overall best
  Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 65001:2
Path #2: Received by speaker 0
Flags: 0x2000000000020005, import: 0x0a0
Not advertised to any peer
65002
  10.0.0.6 from 10.0.0.4 (10.0.0.6), if-handle 0x00000000
  Received Label 24002
  Origin IGP, metric 0, localpref 100, valid, internal, imported
  Received Path ID 0, Local Path ID 0, version 0
  Extended community: RT:65001:101
  Originator: 10.0.0.6, Cluster list: 10.0.0.4
  Higher nexthop admin distance than best path (path #1)
  Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 65001:3

```

使用show命令更详细地了解场景



## 默认行为

适用于头端路由器和路由反射器。

配置:

```
no nexthop validation color-extcomm sr-policy
no bgp bestpath igp-metric sr-policy
```

功能:

执行RIB验证 (硬下一跳)。

BGP不使用SR策略中的管理员/度量。

对服务路由的下一跳执行RIB验证。

如果下一跳没有比默认路由更具体的路由,则服务路由具有无法访问的下一跳。

```
If the RIB metric is available:
RIB metric is used. Route is installed.
If policy is up:
Policy is used.
If policy is not up:
Policy is not used.
If the RIB metric is not available:
Route is not installed.
```

```
RP/0/RP0/CPU0:PE1#show bgp vpnv4 unicast rd 65001:2 10.0.0.9/32
BGP routing table entry for 10.0.0.9/32, Route Distinguisher: 65001:2
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          31        31
Last Modified: Oct 26 14:21:56.714 for 00:01:32
Paths: (1 available, no best path)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
  65002
    10.0.0.5 C:101 (bsid:24005) (inaccessible) from 10.0.0.3 (10.0.0.5)
      Received Label 24002
      Origin IGP, metric 0, localpref 100, valid, internal, not-in-vrf
      Received Path ID 0, Local Path ID 0, version 0
      Extended community: Color:101 RT:65001:101
      Originator: 10.0.0.5, Cluster list: 10.0.0.3
      SR policy color 101, up, not-registered, bsid 24005
```

这也导致服务路由未导入到VRF。

```
RP/0/RP0/CPU0:PE1#show bgp vrf one 10.0.0.9/32
BGP routing table entry for 10.0.0.9/32, Route Distinguisher: 65000:1
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          37        37
Last Modified: Oct 26 14:24:36.714 for 00:00:03
Paths: (0 available, no best path)
  Not advertised to any peer
```

如果在头端路由器上添加非默认静态路由，覆盖服务路由的下一跳，可缓解此问题。ODN网络中通常使用此功能。

此静态路由包含下一跳10.0.0.5，不是默认路由。

```
router static
  address-family ipv4 unicast
    10.0.0.0/24 Null10
  !
!
```

它解决了ODN无法访问的下一跳。

```
RP/0/RP0/CPU0:PE1#show bgp vrf one 10.0.0.9/32
BGP routing table entry for 10.0.0.9/32, Route Distinguisher: 65000:1
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          27        27
Last Modified: Oct 26 14:19:06.714 for 00:00:26
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
  65002
    10.0.0.5 C:101 (bsid:24005) from 10.0.0.3 (10.0.0.5)
      Received Label 24002
      Origin IGP, metric 0, localpref 100, valid, internal, best, group-best, import-candidate,
imported
      Received Path ID 0, Local Path ID 1, version 22
      Extended community: Color:101 RT:65001:101
      Originator: 10.0.0.5, Cluster list: 10.0.0.3
      SR policy color 101, up, not-registered, bsid 24005
      Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 65001:2
```

RR也是如此：如果服务路由的下一跳不可访问，则该路由不会反映在其他iBGP扬声器中。非默认静态路由的解决方法同样适用于RR。

## 使用SR策略度量进行RIB相关验证

适用于头端路由器和路由反射器。

配置:

```
no nexthop validation color-extcomm sr-policy
bgp bestpath igp-metric sr-policy
```

功能:

PCE/路径管理和度量值将传递到BGP，并用于最佳路径计算。

```
Perform RIB validation (hard next-hop).
If NH is reachable in RIB:
If policy is up:
  Use policy metric.
If policy is down:
  Use RIB metric.
```

头端路由器

如果下一跳在RIB中不可达，则服务路由的下一跳将不可访问，且未安装。

如果下一跳可到达（可能通过使用静态路由），则安装服务路由，现在带有管理和度量值。

```
RP/0/RP0/CPU0:PE1#show bgp vrf one 10.0.0.9/32
BGP routing table entry for 10.0.0.9/32, Route Distinguisher: 65000:1
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          43        43
Last Modified: Oct 26 14:42:54.714 for 00:00:03
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
  65002
    10.0.0.5 C:101 (bsid:24005) (admin 20) (metric 30) from 10.0.0.3 (10.0.0.5)
      Received Label 24002
      Origin IGP, metric 0, localpref 100, valid, internal, best, group-best, import-candidate,
imported
      Received Path ID 0, Local Path ID 1, version 43
      Extended community: Color:101 RT:65001:101
      Originator: 10.0.0.5, Cluster list: 10.0.0.3
      SR policy color 101, up, not-registered, bsid 24005
      Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 65001:2
```

政策已经启动。

如果策略关闭，而RIB有下一跳路由，则安装服务路由。但是，服务路由未在CEF表中解析。SR策略不再提供到达终端的连接（MPLS标签堆栈）。

```
RP/0/RP0/CPU0:PE1#show bgp vrf one 10.0.0.9/32
BGP routing table entry for 10.0.0.9/32, Route Distinguisher: 65000:1
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          57        57
Last Modified: Oct 26 15:13:46.714 for 00:01:39
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
  65002
    10.0.0.5 from 10.0.0.3 (10.0.0.5)
      Received Label 24002
      Origin IGP, metric 0, localpref 100, valid, internal, best, group-best, import-candidate,
imported
      Received Path ID 0, Local Path ID 1, version 48
      Extended community: Color:101 RT:65001:101
      Originator: 10.0.0.5, Cluster list: 10.0.0.3
      Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 65001:2
```

该路由已安装，但没有SR策略，CEF未解析为此服务路由。

```
RP/0/RP0/CPU0:PE1#show cef vrf one 10.0.0.9/32
10.0.0.9/32, version 36, drop adjacency, internal 0x5000001 0x30 (ptr 0xe3abf78) [1], 0x600
(0xe54a068), 0xa08 (0xec42558)
Updated Oct 26 15:13:47.003
Prefix Len 32, traffic index 0, precedence n/a, priority 3
gateway array (0xe3b26b8) reference count 2, flags 0x3a, source rib (7), 0 backups
[3 type 1 flags 0x88401 (0xec85888) ext 0x0 (0x0)]
LW-LDI[type=1, refc=1, ptr=0xe54a068, sh-ldi=0xec85888]
```

```
gateway array update type-time 3 Oct 26 15:16:24.524
LDI Update time Oct 26 14:42:54.404
LW-LDI-TS Oct 26 15:13:47.003
via 10.0.0.5/32, 0 dependencies, recursive [flags 0x6000]
  path-idx 0 NHID 0x0 [0xd649400 0x0]
  recursion-via-/32
  next hop VRF - 'default', table - 0xe0000000
  unresolved
  labels imposed {24002}
```

Load distribution: 0 (refcount 3)

Hash	OK	Interface	Address
0	Y	recursive	<b>drop</b>

RR路由器：

如果SR策略为up或not，并且RIB可达性存在，则RR通告服务路由。

## SR策略相关验证与RIB度量

适用于头端路由器。

配置：

```
nexthop validation color-extcomm sr-policy
no bgp bestpath igp-metric sr-policy
```

功能：

PCE/路径管理和度量值不会传递给BGP。

```
If the RIB metric is available:
RIB metric is used. Route is installed.
If policy is up:
Policy is used.
If policy is not up:
Policy is not used.
```

```
If the RIB metric is not available:
Route is not installed.
```

## SR策略相关验证与SR策略度量

适用于头端路由器。

配置：

```
nexthop validation color-extcomm sr-policy
bgp bestpath igp-metric sr-policy
```

功能：

```
Do not perform RIB validation (soft next-hop). RIB reachability is not needed.
```

If policy is up:

Use policy metric and validation, even if RIB reachability is present.

If policy is down:

Use RIB validation and metric if available. If not available, the route is not installed.

如果SR策略可用：

```
RP/0/RP0/CPU0:PE1#show bgp vrf one 10.0.0.9/32
```

```
BGP routing table entry for 10.0.0.9/32, Route Distinguisher: 65000:1
```

```
Versions:
```

```
Process          bRIB/RIB  SendTblVer
Speaker          101      101
```

```
Last Modified: Oct 28 13:32:24.714 for 00:25:39
```

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

```
Not advertised to any peer
```

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

```
Not advertised to any peer
```

```
65002
```

```
10.0.0.5 C:101 (bsid:24008) (admin 30) (metric 30) from 10.0.0.3 (10.0.0.5)
```

```
Received Label 24002
```

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

```
imported
```

```
Received Path ID 0, Local Path ID 1, version 99
```

```
Extended community: Color:101 RT:65001:101
```

```
Originator: 10.0.0.5, Cluster list: 10.0.0.3
```

```
SR policy color 101, up, not-registered, bsid 24008
```

```
Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 65001:2
```

## SR策略相关验证，RIB度量和SR策略不用于最佳路径计算

适用于RR路由器。

配置:

```
nexthop validation color-extcomm disable
```

```
no bgp bestpath igp-metric sr-policy
```

功能:

第一个命令表示禁用颜色扩展通信路径的下一跳可达性验证。下一跳可达性有硬性检查。可禁用软下一跳可达性的验证检查，因为此路由器是RR，仅反映BGP服务路由。RR不为其安装SR策略。如果没有此命令，将执行软检查。如果路由表中没有用于下一跳的路由比默认路由更多，则下一跳将无法访问。然后，路由不会反映。

第二个命令表示SR策略不用于BGP最佳路径计算。因此，不使用SR策略中的管理员/度量。如果下一跳在RIB中，则使用RIB度量。否则，使用网关度量0（下一跳IGP度量）。

## SR策略相关验证，使用RIB度量和SR策略进行最佳路径计算

适用于RR路由器。

配置:

```
nexthop validation color-extcomm disable
```

```
bgp bestpath igp-metric sr-policy
```

功能:

第一个命令表示禁用颜色扩展通信路径的下一跳可达性验证。下一跳可达性有硬性检查。可禁用软下一跳可达性的验证检查，因为这是RR，并且仅反映BGP服务路由。RR不为其安装SR策略。如果没有此命令，将执行软检查。如果路由表中没有用于下一跳的路由比默认路由更多，则下一跳将无法访问。然后，路由不会反映。

第二个命令表示SR策略用于BGP最佳路径计算。

```
If policy is up:
  Use policy metric and validation, even if RIB reachability is present
If policy is down
  Use RIB validation and metric if available
  If RIB validation and metric is not available:
    use the gateway metric 0
```