

在Cisco IOS XR中部署BGP軟式下一躍點

目錄

[簡介](#)

[必要條件](#)

[採用元件](#)

[背景資訊](#)

[解決方案](#)

[BGP路由無法訪問下一跳](#)

[組態](#)

[BGP ODN AS NextHop軟驗證按鈕](#)

[BGP ODN AS Next-Hop硬式驗證旋鈕](#)

[考慮SR原則路徑度量的BGP最佳路徑選取](#)

[BGP最佳路徑選取優先SR原則路徑](#)

[BGP最佳路徑選取強制SR原則路徑](#)

[概觀](#)

[預設行為](#)

[使用SR策略度量的RIB相關驗證](#)

[使用RIB度量的SR策略相關驗證](#)

[使用SR策略度量的SR策略相關驗證](#)

[使用RIB度量和SR策略進行的SR策略相關驗證不用於最佳路徑計算](#)

[使用RIB度量和用於最佳路徑計算的SR策略相關驗證](#)

[驗證](#)

[檢查策略是否為Up或Down](#)

[驗證是否使用了策略](#)

[Show BGP Nexthops命令](#)

[BGP跟蹤](#)

[管理距離和度量](#)

[管理員值](#)

[驗證ODN中的管理員和度量型別](#)

[有效度量](#)

[BGP路徑比較](#)

[比較帶顏色和不帶顏色的BGP路徑](#)

[使用Show命令檢視更詳細的情況](#)

[預設行為](#)

[使用SR策略度量的RIB相關驗證](#)

[使用RIB度量的SR策略相關驗證](#)

[使用SR策略度量的SR策略相關驗證](#)

[使用RIB度量和SR策略的SR策略相關驗證不用於最佳路徑計算](#)

[使用RIB度量和用於最佳路徑計算的SR策略相關驗證](#)

簡介

本檔案將說明Cisco IOS® XR中邊界閘道通訊協定(BGP)進行下一躍點^{處理}的行為。BGP在路由資訊庫(RIB)中安裝路徑之前，需要存取路徑的下一個躍點(NH)。此規則適用於所有BGP揚聲器。這是下一跳驗證檢查。BGP軟式下一躍點功能可確保無需在RIB中連線到BGP下一躍點。

必要條件

需求

本文件沒有特定需求。

採用元件

本檔案是Cisco IOS XR特有的。

本文中的資訊是根據特定實驗室環境內的裝置所建立。文中使用到的所有裝置皆從已清除 (預設) 的組態來啟動。如果您的網路運作中，請確保您瞭解任何指令可能造成的影響。

背景資訊

在單個自治系統(AS)網路、多域網路或Inter-AS場景中，如果NH沒有在域或自治系統之間重分發，則它可能無法再被訪問。

問題不僅限於頭端提供者邊緣(PE)路由器，而且還包括輸出和輸入PE之間的中繼BGP揚聲器(例如路由反射器(RR)和自主系統邊界路由器(ASBR))。中間BGP揚聲器必須能夠到達NH，然後才能安裝和傳播路由。

按需下一跳(ODN)是分段路由(SR)應用程式，它將SR策略安裝到路由器上。連線到這些SR策略的服務路由可以是BGP路由。如果下一躍點有效，則這些BGP路由只能安裝在RIB和Cisco Express Forwarding(CEF)表中。有如無縫MPLS或AS間MPLS虛擬專用網路(VPN)這樣的設計，在RIB中的路由無法保證可到達網路另一部分 (例如另一區域或另一域) 中的BGP下一躍點。如果可由控制器或SR路徑計算元件(SR-PCE)確保可達性，而該控制器或SR路徑計算元件可在整個網路中提供網路元件的可達性，則這不是問題。

目前，如果BGP路由的下一跳在RIB中作為非預設路由，BGP服務路由只能使用SR策略。

如果搭載SR原則的BGP發言者在BGP下一躍點的RIB中沒有路由 (預設路由除外)，則可以採取變通辦法。解決方法是配置一個特定 (非預設) 靜態路由到null0，覆蓋那些無法到達的NH，通過BGP-LU注入路由，或在IGP域之間重新分發路由。

這會造成麻煩和/或影響可擴充性。

解決方案

PE (頭端) 接收彩色BGP L3VPN字首。它可以在本地學習SR策略，也可以請求ODN SR策略來獲取顏色和下一跳。

如果配置了NH驗證，則BGP會對NH執行軟驗證，並在啟用命令時應用NH AD/度量。對於有色NH，AD/度量來自SR控制器。下一跳的軟驗證意味著沒有檢查RIB可達性，但會對SR策略資訊執行檢查。其中包括SR策略路由型別和管理距離以及該度量型別的度量值。

引入新命令在頭端路由器或RR上執行此軟式下一跳驗證。

為RR引入了新命令，以跳過color-extcomm路徑的下一跳可達性驗證。

為RR引入了新命令，以便不將SR策略用於BGP最佳路徑計算。

此功能是在Cisco IOS XR 7.3.2和7.4.1版中匯入。

BGP路由無法訪問下一跳

具有無法訪問的下一跳點的BGP路由不會通告。

此路由是RR上的VPNv4路由。其下一跳 (PE環回) 無法訪問，因為路由表中沒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
```

不會因此通告BGP VPNv4路由。

```
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.
```

當前的解決方法是配置一條覆蓋頭端路由器上PE環回地址的靜態路由。以下是通向null0的此類靜態路由的示例。

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

到Null0的此靜態路由可在RIB中為所有遠端PE環回地址 (BGP下一跳地址) 建立可達性。此靜態路由涵蓋10.0.0.0 - 10.0.0.255範圍內的所有地址。

下一跳通過靜態路由解析。您可以透過此指令看到此情況。

```
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
```

如果使用SR原則驗證下一個躍點，則會看到以下輸出：

```
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

組態

這些配置命令是新的：

```
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軟驗證按鈕

PE(HE):

```
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
```

以下是主要指令：會開啟BGP軟式下一躍點行為。如果存在針對下一跳和顏色的SR策略，則不執行RIB驗證。

BGP ODN AS Next-Hop硬式驗證旋鈕

BGP硬式下一躍點是預設行為。

此命令是恢復此行為的命令：`no nexthop validation color-extcomm`。

考慮SR原則路徑度量的BGP最佳路徑選取

當內部網道通訊協定(IGP)可連線至NH時，且如果演演算法達到BGP最佳路徑選取流程中的步驟8，則首選BGP路徑是到下一躍點距離最低(IGP)的路徑。這是預設行為。請參閱[BGP最佳路徑選取演演算法](#)。

除非配置了**`bgp bestpath igp-metric ignore`**命令，否則情況如此。在這種情況下，根本不會考慮IGP成本。

目前，只考慮對BGP NH的IGP指標；不是SR策略路徑提供的度量。這仍然是預設行為，但有命令指示BGP使用SR原則路徑指標，而不是BGP最佳路徑選取演演算法的IGP指標。

```
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 comparason stage
```

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

此命令將啟用PCE/路徑管理值和度量值的考慮因素。只有當SR策略啟動時，這些管理/指標值才能傳遞到BGP。此命令可讓BGP演演算法根據SR原則中下一躍點的管理員和指標選取最佳路徑。如果沒有此命令，預設行為是僅考慮下一跳的IGP度量。這稱為「下一跳的RIB驗證」。

BGP最佳路徑選取首選SR原則路徑

有些平台不支援具有本機下一躍點或SR策略下一躍點的路徑組合。平台可能不支援這兩種路徑型別的轉發混合。考慮到使用等價多重路徑(ECMP)或非等價多重路徑(UCMP)或備份路徑，這一點非常重要。任何型別的路徑都可以是BGP的最佳路徑。預設行為是隻考慮與BGP最佳路徑具有相同下一躍點型別的路徑。

此命令指示BGP優先使用在路由器執行最佳路徑計算時具有顏色/下一躍點的SR策略的路由。這表示在最佳路徑計算期間，不會考慮SR原則關閉或沒有SR原則的路徑。


`bgp bestpath sr-policy {force | 首選}`

必須配置這兩個關鍵字之一。

```
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
```

如果設定首選選項，則無顏色的eBGP路徑會標示為合格路徑（因此可以是最佳路徑的一部分）。如果不需要此行為，您可以向eBGP路徑新增一個虛擬SR原則。否則，您可以為此命令配置force選項，以使沒有顏色的eBGP路由不符合條件。

如圖所示參考網路。



網路10.99.99.99/32有三種可能的路徑來自路由器PE1。字首10.99.99.99/32由R8和CE路由器通告。

BGP具有路由10.99.99.99/32的3個路徑：2個iBGP (PE2和PE3是BGP下一跳路由器) 和1個eBGP路徑 (來自R8)。

iBGP路徑具有下一躍點10.0.0.5和10.0.0.6。eBGP路徑具有下一躍點10.1.8.8。

組態沒有以下命令**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路徑沒有顏色或SR策略。這是最好的路徑。

如果eBGP路由有顏色，但沒有SR策略，則仍會將其選擇為最佳路徑。

如果eBGP路由具有顏色和SR策略，則會選擇它作為最佳路徑。

下面是另一個示例。eBGP路由沒有顏色，且未配置SR策略和**bgp bestpath sr-policy prefer**命令。

附註：eBGP鄰居在VRF中。這表示您必須在VRF下設定命令**bgp bestpath sr-policy prefer**。

```

router bgp 65001
  nexthop 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路徑是最佳路徑，即使它沒有顏色。如果不希望使用不帶顏色的eBGP路由作為最佳路徑，請使用**force**選項配置命令**bgp bestpath sr-policy**。

附註：本地路徑和重新分發的路徑總是符合最佳路徑計算的條件。

使用此命令檢查平台是否支援通過SR轉發策略和本地下一跳的混合轉發。

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

附註：路由器NCS55xx和NCS560/NCS540顯示否，而ASR9000顯示是。

BGP最佳路徑選取強制SR原則路徑

此命令指示BGP在執行最佳路徑計算時優先使用帶有SR策略下一躍點的路由，但不包括無顏色的eBGP路徑。

```
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 ?
```

已禁用color-extcomm路徑的下一跳可達性驗證

這通常用於路由反射器(RR)。

在RR上：

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

已禁用color-extcomm路徑的下一跳可達性驗證。與SR策略的狀態或存在無關。

概觀

頭端和RR上的行為由下一跳驗證命令和**bgp best path igp-metric sr-policy**命令的配置驅動。有4種情況。每個場景包含兩個配置命令的組合。

預設行為

適用於頭端路由器和RR。

組態：

```
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策略度量的RIB相關驗證

適用於頭端路由器和RR。

組態:

```
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
```

使用RIB度量的SR策略相關驗證

這是預設行為。

適用於頭端路由器。

組態:

```
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.
```

使用SR策略度量的SR策略相關驗證

適用於頭端路由器。

組態:

```
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.
```

使用RIB度量和SR策略進行的SR策略相關驗證不用於最佳路徑計算

適用於RR路由器。

組態:

```
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.

使用RIB度量和用於最佳路徑計算的SR策略相關驗證

適用於RR路由器。

組態:

```
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

驗證

以下是驗證哪種型別的下一跳驗證處於活動狀態以及在最佳路徑計算期間是否使用SR策略的管理距離/指標的方式。

```
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
```

這是預設設定。

以下是使用RIB度量和SR策略不用於最佳路徑計算的SR策略相關驗證的示例。

```
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
```

以下是連線到BGP路由的管理距離/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

```

檢查策略是否為Up或Down

以下是驗證SR原則是處於開啟還是關閉的方式。

```

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

```

驗證是否使用了策略

使用BGP show命令檢視路由。

如果存在繫結段Identifier(BSID)，則此路由使用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

```

此處繫結SID是MPLS標籤。此標籤連結到一個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命令檢視終端的顏色、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: 0x7efcadee6e98
  Gateway: reachable, non-Connected route, prefix length 8
  Resolving Route: 10.0.0.0/8 (static)
  Paths: 0
  RIB NexHop 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]
Metric: [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跟蹤

show bgp trace輸出中的某些專案請參閱SR策略。注意admin/metric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

```

附註：Cisco IOS XR流量控制器(XTC)是指SR控制器。

BGP追蹤中的某些專案指與下一個躍點處理相關的組態變更。

```
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
```

管理距離和度量

管理距離(admin)由SR策略中的度量型別決定。可在頭端路由器上設定度量型別。

```
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
```

管理員值

以下是預設SR策略管理值。

- 延遲10
- TE 20
- IGP (預設) 30
- 跳數40
- NONE/UNKNOWN度量型別 (用於顯式段清單策略) 100

如果度量型別為none，則度量值為1。

admin值越小，BGP的路徑就越受歡迎。

如果管理員具有相同的值，則度量越低，BGP的路徑就越優先。

驗證ODN中的管理員和度量型別

```
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

有效度量

有效度量強制配置此命令所依據的策略的型別和度量。

```
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

```

您可以這樣驗證應用的有效度量型別（管理距離）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路徑比較

預設情況下，BGP路徑的比較不會改變。

如果設定了**bgp bestpath igp-metric sr-policy**指令，則會在BGP最佳路徑選取演演算法中使用SR原則的管理距離和指標。

SR策略的管理距離和度量與SR策略相關。這是本地配置的，或是通過PCEP（路徑計算元素協定）從SR-PCE接收的。這表示如果RR比較路徑，它看不到管理距離和度量，因為它沒有ODN的頭端功能。因此，它沒有與SR PCE的PCEP會話。

此示例顯示一個遠端PE路由器通告的字首。這是組態。

```

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
  !

```

度量型別為TE。

此頭端路由器兩次看到帶有顏色的字首和相同的TE度量，因為對於兩條路徑，它是相同的BGP下一躍點。

```

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)

```

由於兩條路徑的管理距離和度量相同，因此在BGP最佳路徑選擇演算法中，將進一步決定哪一條路徑是最佳路徑。


```
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
```

最佳路徑是第一個路徑，因為它的管理距離比第二個路徑短。度量型別TE的管理距離小於度量型別IGP的管理距離。

ODN-policy-1的SR策略優先順序為200,ODN-policy-2的SR策略優先順序為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
```

以下範例顯示管理距離相同，但度量不同。

```
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
```

以下是具有度量型別跳數的示例。

```
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

```

兩個不同的端點有兩個競爭的BGP路徑。BGP決定哪條路徑會獲選並安裝到路由表中。接下來會根據顏色和端點決定安裝哪個SR原則以將流量轉送到BGP VPNv4字首。

比較帶顏色和不帶顏色的BGP路徑

在案例4中，在前端路由器上啟用軟式下一躍點驗證，且它收到兩個對應一個首碼的BGP路徑，一個包含顏色，另一個沒有顏色。如果沒有下一跳的路由，沒有顏色的路徑將無法訪問下一跳，因此不考慮分批傳送。

```

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

```

最後一個BGP路徑沒有>，因此無法存取下一個躍點。

```

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

```

會使用具有SR原則的BGP路徑。

但是，如果由於RIB中的路由而解析了下一跳10.0.0.6，則可以選擇此路徑作為最佳路徑。但是如果沒有顏色，則無法將其用於ODN，並且SR策略將關閉。但是，此路由的管理距離是100，因此它比具有顏色的路徑高得多。

```

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命令檢視更詳細的情況

預設行為

適用於頭端路由器和RR。

組態:

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

功能：

執行RIB驗證（硬式下一跳）。

BGP不使用SR策略中的管理/度量。

對服務路由的下一跳執行RIB驗證。

如果下一跳沒有比預設路由更具體的路由，則服務路由具有不可訪問的下一跳。

```
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
```

這也導致服務路由沒有匯入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
```

如果在頭端路由器上新增覆蓋服務路由下一跳的非預設靜態路由，可以緩解此問題。通常用於ODN網路。

此靜態路由涵蓋下一跳10.0.0.5，不是預設路由。

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

它解決了ODN無法訪問的下一跳。

```
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
```

在RR上也是如此：如果服務路由的下一跳無法訪問，則該路由不會反映在其他iBGP揚聲器中。在RR上可以使用非預設靜態路由的相同解決方法。

使用SR策略度量的RIB相關驗證

適用於頭端路由器和RR。

組態:

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

功能：

PCE/path admin和度量值傳遞給BGP，用於最佳路徑計算。

```
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.
```

頭端路由器

如果RIB中無法訪問下一跳，則服務路由將無法訪問下一跳，且未安裝該下一跳。

如果下一跳可訪問（可能通過使用靜態路由），則安裝服務路由，現在安裝管理值和度量值。

```
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
策略已啟動。
```

如果策略關閉，當RIB具有下一跳的路由時，將安裝服務路由。但是，在CEF表中未解析服務路由。SR策略不再提供到達端點的連線（MPLS標籤堆疊）。

```
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
路由已安裝，但沒有SR策略時無法為此服務路由解析CEF。
```

```
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	drop

RR路由器：

如果SR策略是否啟動，且存在RIB可達性，則RR會通告服務路由。

使用RIB度量的SR策略相關驗證

適用於頭端路由器。

組態：

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

功能：

PCE/path admin和度量值不會傳遞給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.
```

使用SR策略度量的SR策略相關驗證

適用於頭端路由器。

組態：

```
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
```

使用RIB度量和SR策略的SR策略相關驗證不用於最佳路徑計算

適用於RR路由器。

組態:

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

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

功能：

第一個命令表示已禁用color-extcomm路徑的下一跳可達性驗證。對下一跳可達性進行嚴格檢查。可停用軟式下一躍點連線能力的驗證檢查，因為此路由器是RR且僅反映BGP服務路由。RR沒有為其安裝SR策略。如果沒有此命令，將執行軟檢查。如果路由表中除了預設路由外沒有其它下一跳路由，則無法訪問下一跳。然後不反映路由。

第二個命令表示SR原則不用於BGP最佳路徑計算。因此，不使用SR策略中的管理/度量。如果下一跳位於RIB中，則使用RIB度量。否則，使用網關度量0（下一跳IGP度量）。

使用RIB度量和用於最佳路徑計算的SR策略相關驗證

適用於RR路由器。

組態:

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

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

功能：

第一個命令表示已禁用color-extcomm路徑的下一跳可達性驗證。對下一跳可達性進行嚴格檢查。可停用軟式下一躍點連線能力的驗證檢查，因為這是RR，且僅反映BGP服務路由。RR沒有為其安裝SR策略。如果沒有此命令，將執行軟檢查。如果路由表中除了預設路由外沒有其它下一跳路由，則無法訪問下一跳。然後不反映路由。

第二個命令表示SR原則用於BGP最佳路徑計算。

```
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
```