瞭解BGP動態片段路由流量工程

目錄

簡介

本檔案介紹如何瞭解、設定和驗證Cisco IOS® XR中的BGP動態片段路由流量工程(SR-TE)功能。

必要條件

本檔案沒有先決條件。

需求

本文件沒有特定需求。

採用元件

本檔案中的資訊是根據Cisco IOS XR和Cisco IOS XE。

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

SR-TE提供無需狀態建立和維護（無狀態）即可引導流量通過啟用了SR的核心的功能。SR-TE策略表示為指定路徑的段清單，稱為段ID(SID)清單。不需要信令，因為狀態在資料包中，而SID清單由傳輸路由器作為指令集處理。

使用動態邊界閘道通訊協定(BGP)SR-TE，您可以根據任意標準（例如參與區段路由網路的路由器所傳送的社群）來產生自動SR-TE原則。為了能夠滿足站點應用程式和基於特定需求的計算路徑的服務級別保證(SLA)，您可以通過設定社群並觸發這些策略來生成給定IP子網或服務的自動SR-TE策略。

注意：支援使用社群以外的匹配條件來建立動態SR-TE策略。

此功能的一個常見應用是在MPLS L3VPN環境中，在此環境中，網路管理員可以觸發自動SR-TE隧道策略，以根據特定約束（延遲、頻寬等）路由流量。對於本文檔中的演示，我們將建立連線XR1和XR5的L3VPN服務，並在XR2（頭端）上觸發基於在MP-BGP上的XR4（尾端）上設定的特定社群的自動隧道。

設定

網路圖表

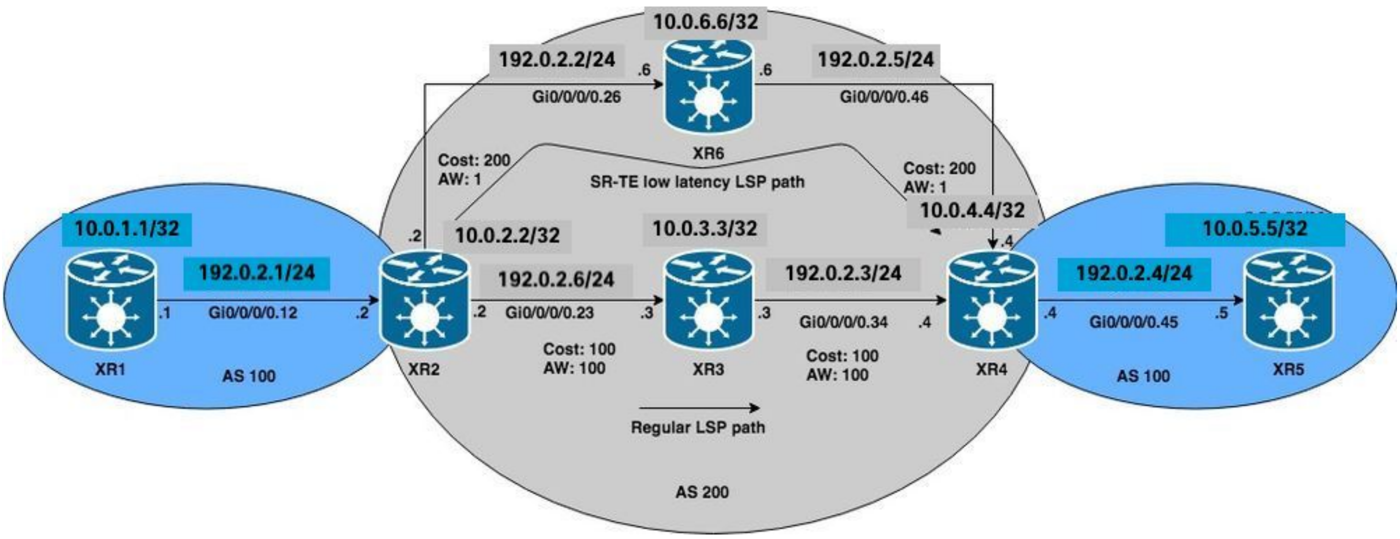

初始配置

已啟用L3VPN、段路由和SR-TE基本配置。

XR1

```
hostname XR1
logging console debugging
interface Loopback0
 ipv4 address 10.0.1.1 255.255.255.255
!
interface GigabitEthernet0/0/0/0.12
 ipv4 address 192.0.2.1 255.255.255.0
 encapsulation dot1q 12
!
route-policy PASS
  pass
end-policy
!
router bgp 100
 bgp router-id 10.0.1.1
 address-family ipv4 unicast
  network 10.0.1.1/32
 !
 neighbor 192.0.2.7
  remote-as 200
  address-family ipv4 unicast
   route-policy PASS in
```

```
   route-policy PASS out
  !
 !
!
end
```

XR2

```
hostname XR2 logging console debugging vrf BLUE address-family ipv4 unicast import route-target
1:1 ! export route-target 1:1 ! ! ! interface Loopback0 ipv4 address 10.0.2.2 255.255.255.255 !
interface GigabitEthernet0/0/0/0.12 vrf BLUE ipv4 address 192.0.2.7 255.255.255.0 encapsulation
dot1q 12 ! interface GigabitEthernet0/0/0/0.23 ipv4 address 192.0.2.8 255.255.255.0
encapsulation dot1q 23 ! interface GigabitEthernet0/0/0/0.26 ipv4 address 192.0.2.9
255.255.255.0 encapsulation dot1q 26 ! route-policy PASS pass end-policy ! ! router ospf 1
segment-routing mpls segment-routing forwarding mpls segment-routing sr-prefer address-family
ipv4 area 0 mpls traffic-eng interface Loopback0 prefix-sid index 2 ! interface
GigabitEthernet0/0/0/0.23 cost 100 network point-to-point ! interface GigabitEthernet0/0/0/0.26
cost 200 network point-to-point ! ! mpls traffic-eng router-id Loopback0 ! router bgp 100 bgp
router-id 10.0.2.2 address-family vpnv4 unicast ! neighbor 10.0.4.4 remote-as 200 update-source
Loopback0 address-family vpnv4 unicast ! ! vrf BLUE rd 1:1 address-family ipv4 unicast !
neighbor 192.0.2.10 remote-as 200 address-family ipv4 unicast route-policy PASS in route-policy
PASS out as-override ! ! ! ! mpls oam ! mpls traffic-eng interface GigabitEthernet0/0/0/0.23
admin-weight 100 ! interface GigabitEthernet0/0/0/0.26 admin-weight 1 ! ! end
```

XR3

```
hostname XR3 logging console debugging interface Loopback0 ipv4 address 10.0.3.3 255.255.255.255
! ! interface GigabitEthernet0/0/0/0.23 ipv4 address 192.0.2.11 255.255.255.0 encapsulation
dot1q 23 ! interface GigabitEthernet0/0/0/0.34 ipv4 address 192.0.2.12 255.255.255.0
encapsulation dot1q 34 ! router ospf 1 segment-routing mpls segment-routing forwarding mpls
segment-routing sr-prefer address-family ipv4 area 0 mpls traffic-eng interface Loopback0
prefix-sid index 3 ! interface GigabitEthernet0/0/0/0.23 cost 100 network point-to-point !
interface GigabitEthernet0/0/0/0.34 cost 100 network point-to-point ! ! mpls traffic-eng router-
id Loopback0 ! mpls oam ! mpls traffic-eng interface GigabitEthernet0/0/0/0.23 admin-weight 100
! interface GigabitEthernet0/0/0/0.34 admin-weight 100 ! ! end
```

XR4

```
hostname XR4 logging console debugging vrf BLUE address-family ipv4 unicast import route-target
1:1 ! export route-target 1:1 ! ! ! interface Loopback0 ipv4 address 10.0.4.4 255.255.255.255 !
interface GigabitEthernet0/0/0/0.34 ipv4 address 192.0.2.13 255.255.255.0 encapsulation dot1q 34
! interface GigabitEthernet0/0/0/0.45 vrf BLUE ipv4 address 192.0.2.14 255.255.255.0
encapsulation dot1q 45 ! interface GigabitEthernet0/0/0/0.46 ipv4 address 192.0.2.15
255.255.255.0 encapsulation dot1q 46 ! route-policy PASS pass end-policy ! ! router ospf 1
segment-routing mpls segment-routing forwarding mpls segment-routing sr-prefer address-family
ipv4 area 0 mpls traffic-eng interface Loopback0 prefix-sid index 4 ! interface
GigabitEthernet0/0/0/0.34 cost 100 network point-to-point ! interface GigabitEthernet0/0/0/0.46
cost 200 network point-to-point ! ! mpls traffic-eng router-id Loopback0 ! router bgp 100 bgp
router-id 10.0.4.4 address-family vpnv4 unicast ! neighbor 10.0.2.2 remote-as 200 update-source
Loopback0 address-family vpnv4 unicast ! ! vrf BLUE rd 1:1 bgp unsafe-ebgp-policy address-family
ipv4 unicast ! neighbor 192.0.2.16 remote-as 200 address-family ipv4 unicast route-policy PASS
in route-policy PASS out as-override ! ! ! ! mpls oam ! mpls traffic-eng interface
GigabitEthernet0/0/0/0.34 admin-weight 100 ! interface GigabitEthernet0/0/0/0.46 admin-weight 1
! ! end
```

XR5

```
hostname XR5
logging console debugging
interface Loopback0
description REGULAR LSP PATH ipv4 address 10.0.5.5 255.255.255.255 ! interface Loopback1
description DELAY SENSITIVE - LOW LATENCY PATH (1:1) ipv4 address 10.0.5.55 255.255.255.255 !
interface GigabitEthernet0/0/0/0.45 ipv4 address 192.0.2.16 255.255.255.0 encapsulation dot1q 45
! route-policy PASS pass end-policy ! router bgp 100 bgp router-id 10.0.5.5 bgp unsafe-ebgp-
policy address-family ipv4 unicast network 10.0.5.5/32 network 10.0.5.55/32 ! neighbor
192.0.2.14 remote-as 200 address-family ipv4 unicast route-policy PASS in route-policy PASS out
! ! ! mpls oam ! end
```

XR6

```
hostname XR6 logging console debugging interface Loopback0 ipv4 address 10.0.6.6 255.255.255.255 ! interface GigabitEthernet0/0/0/0.26 ipv4 address 192.0.2.17 255.255.255.0 encapsulation dot1q 26 ! interface GigabitEthernet0/0/0/0.46 ipv4 address 192.0.2.18 255.255.255.0 encapsulation dot1q 46 ! router ospf 1 segment-routing mpls segment-routing forwarding mpls segment-routing sr-prefer address-family ipv4 area 0 mpls traffic-eng interface Loopback0 prefix-sid index 6 ! interface GigabitEthernet0/0/0/0.26 cost 200 network point-to-point ! interface GigabitEthernet0/0/0/0.46 cost 200 network point-to-point ! ! mpls traffic-eng router-id Loopback0 ! mpls oam ! mpls traffic-eng interface GigabitEthernet0/0/0/0.26 admin-weight 1 ! interface GigabitEthernet0/0/0/0.46 admin-weight 1 ! ! end
```

XR2和XR4(PE)使用分段路由構建了LSP，這可以通過為相應的分段路由FEC使用MPLS ping進行驗證。在此案例中，有兩個可能路徑可用來將L3VPN流量從XR1傳輸到XR5:

常規LSP路徑：XR1 > XR2 > **XR3** > XR4 > XR5

低延遲LSP路徑：XR1 > XR2 >**XR6** > XR4 > XR5

最初，由於IGP成本較低，XR1和XR5之間的所有流量都通過常規LSP路徑通過XR3路由，我們可以按照這些說明確認LSP和連線。通過XR3從XR2到達XR4的IGP成本為201，而通過XR6到達XR4的成本為401。儘管通過XR3的路徑具有更好的路徑度量，但VRF BLUE上的低延遲服務必須通過通過XR6的路徑進行路由。

```
RP/0/0/CPU0:XR2#ping mpls ipv4 10.0.4.4/32 fec-type generic verbose

Sending 5, 100-byte MPLS Echos to 10.0.4.4/32,
      timeout is 2 seconds, send interval is 0 msec:

Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
  'L' - labeled output interface, 'B' - unlabeled output interface,
  'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
  'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
  'P' - no rx intf label prot, 'p' - premature termination of LSP,
  'R' - transit router, 'I' - unknown upstream index,
  'X' - unknown return code, 'x' - return code 0

Type escape sequence to abort.

!      size 100, reply addr 192.0.2.13, return code 3
!      size 100, reply addr 192.0.2.13, return code 3
!      size 100, reply addr 192.0.2.13, return code 3
!      size 100, reply addr 192.0.2.13, return code 3
!      size 100, reply addr 192.0.2.13, return code 3

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/10 ms
```

注意：在分段路由中使用ping MPLS應用時，我們必須使用Nil-FEC或通用FEC。

如果您在XR1上驗證L3VPN服務，則可以通過常規LSP路徑確認分別到XR5環回10.0.5.5/32和10.0.5.55/32的可達性。基本L3VPN服務在SR MPLS核心中啟用。

```
RP/0/0/CPU0:XR1#ping 10.0.5.5 source 10.0.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.0.5.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/7/9 ms
```

```
RP/0/0/CPU0:XR1#ping 10.0.5.55 source 10.0.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.0.5.55, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/7/9 ms


RP/0/0/CPU0:XR1#traceroute 10.0.5.5 source 10.0.1.1

Type escape sequence to abort.
Tracing the route to 10.0.5.5

 1  192.0.2.7 9 msec  0 msec  0 msec
 2  192.0.2.11 [MPLS: Labels 16004/24002 Exp 0] 0 msec  0 msec  0 msec
 3  192.0.2.13 [MPLS: Label 24002 Exp 0] 0 msec  0 msec  0 msec
 4  192.0.2.16 0 msec  *  0 msec

RP/0/0/CPU0:XR1#traceroute 10.0.5.55 source 10.0.1.1

Type escape sequence to abort.
Tracing the route to 10.0.5.55

 1  192.0.2.7 9 msec  0 msec  0 msec
 2  192.0.2.11 [MPLS: Labels 16004/24005 Exp 0] 0 msec  0 msec  0 msec
 3  192.0.2.13 [MPLS: Label 24005 Exp 0] 0 msec  0 msec  0 msec
 4  192.0.2.16 0 msec  *  0 msec
```

如前所述，VRF BLUE上的所有流量都通過常規LSP路徑XR1 > XR2 > XR3 > XR4 > XR5。

設定BGP動態SR-TE

在本例中，將XR4（尾端）配置為插入社群1:1，並將其傳送到XR2，以發出在VRF BLUE上為字首10.0.5.55/32建立SR-TE策略的訊號。SR-TE策略路徑選擇將被設定為採用低延遲路徑而不是常規LSP，為此，請通過XR6選擇最低TE度量（管理權重）。通過XR6的總的TE度量（管理權重）為2，因為通過XR6指向XR4（尾端）的傳出介面上的管理權重已設定為1，如參考拓撲圖和初始配置所示。

要建立動態SR-TE策略，我們需要配置將哪個環回用作源，以及頭端將用來生成隧道的動態隧道範圍，SR-TE策略XR2的前端需要此配置。將隧道範圍設定為最小值500和最大值500，有效地為隧道在前端建立單個SR-TE隧道和到環回0的源環回。

```
XR2
ipv4 unnumbered mpls traffic-eng Loopback0
mpls traffic-eng
 auto-tunnel p2p
  tunnel-id min 500 max 500
 !
!
end
```

在XR4上，設定社群1:1，並將其套用在VRF BLUE首碼10.0.5.55/32上，這樣可允許它在BGP更新中插入社群。

```
XR4
route-policy COMMUNITY_1:1
  # 1:1 Community
  if destination in (10.0.5.55/32) then
```

```
    set community (1:1)
  endif
  pass
end-policy
!
router bgp 100
 vrf BLUE
 !
  neighbor 192.0.2.16
  address-family ipv4 unicast
    route-policy COMMUNITY_1:1 in
 !
!
end
```

驗證XR2（頭端）我們可以看到它在從XR4接收的VPNv4更新上設定了社群1:1。

```
RP/0/0/CPU0:XR2#show bgp vrf BLUE 10.0.5.55/32 detail
BGP routing table entry for 10.0.5.55/32, Route Distinguisher: 1:1 Versions: Process bRIB/RIB
SendTblVer Speaker 36 36 Flags: 0x00043001+0x00000200; Last Modified: Nov 23 17:50:59.798 for
00:02:53 Paths: (1 available, best #1) Advertised to CE peers (in unique update groups):
192.0.2.10 Path #1: Received by speaker 0 Flags: 0x4000000085060005, import: 0x9f Advertised to
CE peers (in unique update groups): 192.0.2.10 200 10.0.4.4 (metric 201) from 10.0.4.4
(10.0.4.4) Received Label 24005 Origin IGP, metric 0, localpref 100, valid, internal, best,
group-best, import-candidate, imported Received Path ID 0, Local Path ID 0, version 36
Community: 1:1
      Extended community: RT:1:1
      Source AFI: VPNv4 Unicast, Source VRF: BLUE, Source Route Distinguisher: 1:1
```

在XR2（頭端）上，建立與社群1:1匹配的RPL路由策略，並為MPLS流量工程設定相應的屬性集。設定策略後，我們可以轉到MPLS-TE配置模式並為SR-TE策略設定相應的屬性集，並指明路徑選擇標準，在本例中是分段路由和TE度量，因為我們希望通過XR6通過最低管理權重選擇路徑。

```
XR2
route-policy DYN_BGP_SR-TE
  # Matches community 1:1
  if community matches-every (1:1) then
    set mpls traffic-eng attributeset DYN_SR-TE_POLICIES
  endif
  pass
end-policy
!
router bgp 100
!
 neighbor 10.0.4.4
 address-family vpnv4 unicast
   route-policy DYN_BGP_SR-TE in
  !
mpls traffic-eng
 attribute-set p2p-te DYN_SR-TE_POLICIES
  path-selection
   metric te
   segment-routing adjacency unprotected
  !
end
```

驗證

完成後，您可以觀察到tunnel-te 500介面已針對指定範圍動態建立。

```
RP/0/0/CPU0:XR2#show mpls traffic-eng tunnels segment-routing tabular

          Tunnel   LSP     Destination          Source    Tun    FRR  LSP  Path
            Name    ID         Address         Address  State  State Role  Prot
----------------- ----- --------------- -------------- ------ ------ ---- -----
    ^tunnel-te500     2         10.0.4.4        10.0.2.2     up  Inact Head Inact
^ = automatically created P2P/P2MP tunnel
```

BGP RIB表示「DYN_SR-TE_POLICIES」策略附加到字首，這意味著流量必須根據該策略進行路由。

```
RP/0/0/CPU0:XR2#show bgp vrf BLUE

Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network            Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf BLUE)
*> 10.0.1.1/32         192.0.2.10                0              0 200 i
*>i10.0.5.5/32         10.0.4.4                 0    100      0 200 i
*>i10.0.5.55/32        10.0.4.4 T:DYN_SR-TE_POLICIES
                                               0    100      0 200 i
```

如果詳細驗證字首10.0.5.55/32的BGP RIB，我們可以看到將用來產生SR-TE通道的控制平面資訊。

```
RP/0/0/CPU0:XR2#show bgp vrf BLUE 10.0.5.55/32 detail

BGP routing table entry for 10.0.5.55/32, Route Distinguisher: 1:1
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                 39          39
    Flags: 0x00041001+0x00000200;
Last Modified: Nov 23 17:55:22.798 for 00:04:43
Paths: (1 available, best #1)
  Advertised to CE peers (in unique update groups):
    192.0.2.10
  Path #1: Received by speaker 0
  Flags: 0x4000000085060005, import: 0x9f
  Advertised to CE peers (in unique update groups):
    192.0.2.10
  200
    10.0.4.4 T:DYN_SR-TE_POLICIES (metric 201) from 10.0.4.4 (10.0.4.4)
      Received Label 24005
      Origin IGP, metric 0, localpref 100, valid, internal, best, group-best, import-candidate,
imported
      Received Path ID 0, Local Path ID 0, version 39
      Community: 1:1
      Extended community: RT:1:1
      TE tunnel attribute-set DYN_SR-TE_POLICIES, up, registered, binding-label 24000, if-handle
0x00000130

      Source AFI: VPNv4 Unicast, Source VRF: BLUE, Source Route Distinguisher: 1:1
```

我們可以看到隧道策略處於**up**狀態並已**註冊**。分配的繫結SID為24000，此繫結SID可用於驗證哪個隧道用於此特定字首。如前所述，tunnel-te500被建立並安裝在LFIB中。

```
RP/0/0/CPU0:XR2#show mpls forwarding labels 24000 detail
Local Outgoing Prefix Outgoing Next Hop Bytes Label Label or ID Interface Switched ------ ------
----- ------------------ ------------ -------------- ------------ 24000  Pop         No ID
tt500        point2point     0
     Updated: Nov 23 17:55:23.267
     Label Stack (Top -> Bottom): { }
     MAC/Encaps: 0/0, MTU: 0
     Packets Switched: 0
```

注意：繫結SID有許多使用案例，對於此特定文檔，限制其在本地驗證中的使用，但其應用範圍要廣泛得多。

或者，可以使用BGP RIB輸出中的給定**if-handle 0x00000130** 來檢查指派給字首10.0.5.55/32的SR-TE策略。

```
RP/0/0/CPU0:XR2#show mpls forwarding tunnels ifh 0x00000130 detail
Tunnel Outgoing Outgoing Next Hop Bytes Name Label Interface Switched ------------- -----------
------------ -------------- ------------ tt500        (SR) 24003       Gi0/0/0/0.26 192.0.2.17
0
     Updated: Nov 23 17:55:23.267
     Version: 138, Priority: 2
     Label Stack (Top -> Bottom): { 24003 }
     NHID: 0x0, Encap-ID: N/A, Path idx: 0, Backup path idx: 0, Weight: 0
     MAC/Encaps: 18/22, MTU: 1500
     Packets Switched: 0

  Interface Name: tunnel-te500, Interface Handle: 0x00000130, Local Label: 24001
  Forwarding Class: 0, Weight: 0
  Packets/Bytes Switched: 0/0
```

XR2（頭端）上的SR-TE策略將具有從控制平面和資料平面角度轉發流量的這些屬性。此外，SR-TE通道的狀態資訊也可以視為依照輸出，且必須與先前的驗證相符。

```
RP/0/0/CPU0:XR2#show mpls traffic-eng tunnels segment-routing p2p 500

Name: tunnel-te500  Destination: 10.0.4.4  Ifhandle:0x130 (auto-tunnel for BGP default)
  Signalled-Name: auto_XR2_t500
  Status:
    Admin:    up Oper:    up   Path:  valid   Signalling: connected

    path option 10, (Segment-Routing) type dynamic  (Basis for Setup, path weight 2)
    G-PID: 0x0800 (derived from egress interface properties)
    Bandwidth Requested: 0 kbps  CT0
    Creation Time: Fri Nov 23 17:55:23 2018 (00:09:01 ago)
  Config Parameters:
    Bandwidth:        0 kbps (CT0) Priority:  7  7 Affinity: 0x0/0x0
    Metric Type: TE (interface)
    Path Selection:
      Tiebreaker: Min-fill (default)
      Protection: Unprotected Adjacency
    Hop-limit: disabled
    Cost-limit: disabled
    Path-invalidation timeout: 10000 msec (default), Action: Tear (default)
    AutoRoute: disabled  LockDown: disabled   Policy class: not set
    Forward class: 0 (default)
    Forwarding-Adjacency: disabled
```

```
    Autoroute Destinations: 0
    Loadshare:          0 equal loadshares
    Auto-bw: disabled
    Path Protection: Not Enabled
    Attribute-set: DYN_SR-TE_POLICIES (type p2p-te)
    BFD Fast Detection: Disabled
    Reoptimization after affinity failure: Enabled
    SRLG discovery: Disabled
  History:
    Tunnel has been up for: 00:09:01 (since Fri Nov 23 17:55:23 UTC 2018)
    Current LSP:
      Uptime: 00:09:01 (since Fri Nov 23 17:55:23 UTC 2018)
    Reopt. LSP:
      Last Failure:
        LSP not signalled, identical to the [CURRENT] LSP
        Date/Time: Fri Nov 23 17:56:53 UTC 2018 [00:07:31 ago]

  Segment-Routing Path Info (OSPF 1 area 0)
    Segment0[Link]: 192.0.2.9 - 192.0.2.17, Label: 24005
    Segment1[Link]: 192.0.2.18 - 192.0.2.15, Label: 24003
Displayed 1 (of 1) heads, 0 (of 0) midpoints, 0 (of 0) tails
Displayed 1 up, 0 down, 0 recovering, 0 recovered heads
```

直接檢查VRF BLUE RIB上的字首，我們可以確認繫結SID 24000已分配給字首。

```
RP/0/0/CPU0:XR2#show route vrf BLUE 10.0.5.55/32 detail

Routing entry for 10.0.5.55/32
  Known via "bgp 100", distance 200, metric 0
  Tag 200, type internal
  Installed Nov 23 17:55:23.267 for 00:10:38
  Routing Descriptor Blocks
    10.0.4.4, from 10.0.4.4
      Nexthop in Vrf: "default", Table: "default", IPv4 Unicast, Table Id: 0xe0000000
      Route metric is 0
      Label: 0x5dc5 (24005)
      Tunnel ID: None
      Binding Label: 0x5dc0 (24000)
      Extended communities count: 0
      Source RD attributes: 0x0000:1:1
      NHID:0x0(Ref:0)
  Route version is 0x5 (5)
  No local label
  IP Precedence: Not Set
  QoS Group ID: Not Set
  Flow-tag: Not Set
  Fwd-class: Not Set
  Route Priority: RIB_PRIORITY_RECURSIVE (12) SVD Type RIB_SVD_TYPE_REMOTE
  Download Priority 3, Download Version 27
  No advertising protos.
```

適用於VRF BLUE的FIB表示根據我們的BGP動態SR-TE策略，此首碼的轉送是透過tunnel-te 500完成的。

```
RP/0/0/CPU0:XR2#show cef vrf BLUE 10.0.5.55/32 detail
10.0.5.55/32, version 27, internal 0x1000001 0x0 (ptr 0xa142a574) [1], 0x0 (0x0), 0x208
(0xa159d208) Updated Nov 23 17:55:23.287 Prefix Len 32, traffic index 0, precedence n/a,
priority 3 gateway array (0xa129f23c) reference count 1, flags 0x4038, source rib (7), 0 backups
[1 type 1 flags 0x48441 (0xa15b780c) ext 0x0 (0x0)] LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]
```

```
gateway array update type-time 1 Nov 23 17:55:23.287 LDI Update time Nov 23 17:55:23.287 via
local-label 24000, 3 dependencies, recursive [flags 0x6000]     path-idx 0 NHID 0x0 [0xa1605bf4
0x0]
    recursion-via-label
    next hop VRF - 'default', table - 0xe0000000
    next hop via 24000/0/21
     next hop tt500        labels imposed {ImplNull 24005}


    Load distribution: 0 (refcount 1)

    Hash  OK  Interface                 Address
    0     Y   Unknown                   24000/0
```

在XR1上，我們可以驗證連線並確認流量通過tunnel-te 500通過XR6的低延遲路徑。

```
RP/0/0/CPU0:XR1#traceroute 10.0.5.55 source 10.0.1.1

Type escape sequence to abort.
Tracing the route to 10.0.5.55

 1  192.0.2.7 0 msec  0 msec  0 msec
 2  192.0.2.17 [MPLS: Labels 24003/24005 Exp 0] 0 msec  0 msec  0 msec
 3  192.0.2.15 [MPLS: Label 24005 Exp 0] 0 msec  0 msec  0 msec
 4  192.0.2.16 0 msec  *  9 msec
```

對應我們的SR-TE原則的tunnel-te500的XR2計數增加。

```
RP/0/0/CPU0:XR2#show mpls forwarding tunnels

Tunnel        Outgoing    Outgoing     Next Hop        Bytes
Name          Label       Interface                    Switched
------------- ----------- ------------ --------------- ------------
tt500        (SR) 24003       Gi0/0/0/0.26 192.0.2.17        2250
```

字首10.0.5.5/32的路徑仍在通過XR3通過常規LSP路徑，如下圖所示。

```
RP/0/0/CPU0:XR1#traceroute 10.0.5.5 source 10.0.1.1

Type escape sequence to abort.
Tracing the route to 10.0.5.5

 1  192.0.2.7 0 msec  0 msec  0 msec
 2  192.0.2.11 [MPLS: Labels 16004/24002 Exp 0] 0 msec  0 msec  0 msec
 3  192.0.2.13 [MPLS: Label 24002 Exp 0] 0 msec  0 msec  0 msec
 4  192.0.2.16 0 msec  *  0 msec
```

疑難排解

目前尚無適用於此組態的具體疑難排解資訊。

摘要

BGP動態SR-TE為啟用SR的核心中的流量工程提供精細度和自動實施路由策略。可以根據任意標準

觸發自動隧道建立，使網路管理員能夠輕鬆建立符合站點應用要求的流量模式。

關於此翻譯

思科已使用電腦和人工技術翻譯本文件，讓全世界的使用者能夠以自己的語言理解支援內容。請注意，即使是最佳機器翻譯，也不如專業譯者翻譯的內容準確。Cisco Systems, Inc. 對這些翻譯的準確度概不負責，並建議一律查看原始英文文件（提供連結）。