了解BGP动态分段路由流量工程

目录

简介

本文档介绍如何理解、配置和验证Cisco IOS® XR中的BGP动态分段路由流量工程(SR-TE)功能。

先决条件

读者无需满足任何前提条件即可理解本文档内容。

要求

本文档没有任何特定的要求。

使用的组件

本文档中的信息基于Cisco IOS XR和Cisco IOS XE。

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

SR-TE能够引导流量通过启用SR的核心，而无需状态创建和维护（无状态）。SR-TE 策略表示为一个指定路径的分段列表，称为分段 ID (SID) 列表。无需信令，因为状态信息将包含在数据包中，并且传输路由器会将 SID 列表作为一组指令进行处理。

使用动态边界网关协议(BGP)SR-TE，您可以根据任意条件（例如参与分段路由网络的路由器所发信号的社区）生成自动SR-TE策略。为了能够根据特定要求满足站点应用和计算路径的服务级别保证(SLA)，您可以通过设置社区并触发这些策略，为给定IP子网或服务生成自动SR-TE策略。

注意：还支持匹配社区以外的条件来创建动态SR-TE策略。

此功能的一个常见应用是在MPLS L3VPN环境中，在此环境中，网络管理员可以触发自动SR-TE隧道策略，以根据特定限制（延迟、带宽等）路由流量。对于本文档中的演示，我们将创建连接XR1和XR5的L3VPN服务，并根据MP-BGP上的XR4（尾端）上的特定团体集在XR2（头端）上触发自动隧道。

配置

网络图

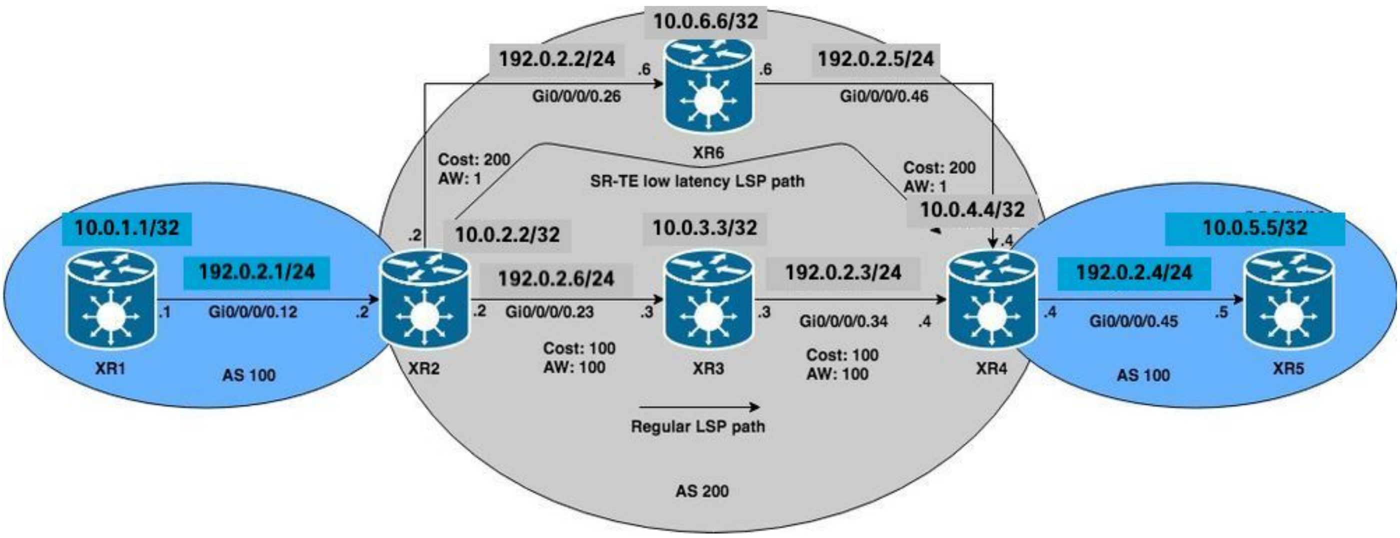

初始配置

L3VPN、网段路由和SR-TE基本配置已启用。

```
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)使用分段路由构建了LSP，这可以通过为相应分段路由FEC使用MPLS ping进行验证。在本场景中，有两条可能的路径用于传输从XR1到XR5的L3VPN流量：

常规LSP路径：XR1 > XR2 > **XR3** > XR4 > XR5

低延迟LSP路径：XR1 > XR2 >**XR6** > XR4 > XR5

最初，由于IGP成本较低，因此XR1和XR5之间的所有流量都通过XR3通过常规LSP路径路由，我们可以按照这些规范确认LSP和连接。通过 XR3 从 XR2 到达 XR4 的 IGP 开销为 201，而通过 XR6 从 XR2 到达 XR4 的 IGP 开销为 401。即使通过 XR3 路由的路径具有更佳的路径度量值，VRF BLUE 上的低延迟服务也必须采用通过 XR6 的路径进行路由。

```
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应用时，我们必须使用Nil-FEC或通用FEC。

如果您在XR1上验证L3VPN服务，则可以通过常规LSP路径确认分别到XR5环回10.0.5.5/32和10.0.5.55/32的可达性。基本L3VPN服务在SR MPLS核心中启用。

```
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
```

如观察，VRF BLUE上的所有流量通过常规LSP路径XR1 > XR2 > XR3 > XR4 > XR5。

配置BGP动态SR-TE

在本示例中，将XR4（尾端）配置为插入社区1:1，并将其发送到XR2，以发出在VRF BLUE上为前缀10.0.5.55/32创建SR-TE策略的信号。SR-TE策略路径选择将被设置为采用低延迟路径而不是常规LSP，通过通过XR6选择最低TE度量（管理权重）执行此操作。通过XR6的总的TE度量（管理权重）为2，因为通过XR6指向XR4（尾端）的传出接口上的管理权重已设置为1，如参考拓扑图和初始配置所示。

要创建动态SR-TE策略，我们需要配置将用作源的环回以及头端将用来生成隧道的动态隧道范围，此配置在SR-TE策略XR2的前端是必需的。将隧道范围设置为最小值500和最大值500，从而有效地为隧道在头端创建单个SR-TE隧道以及到环回0的源环回。

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

在XR4上，设置社区1:1并将其应用于VRF BLUE前缀10.0.5.55/32，这将允许其在BGP更新中插入社区。

```
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
```

检验XR2（头端）我们可以看到它在从XR4收到的VPNv4更新上设置了社区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（前端）上，创建与社区1:1匹配的RPL路由策略，并为MPLS流量工程设置相应的属性集。设置策略后，我们可以转到MPLS-TE配置阶段，为SR-TE策略设置相应的属性集，并指明路径选择标准，在本例中是分段路由和TE度量，因为我们希望通过XR6通过最低管理权重选择路径。

```
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
```

验证

完成后，您可以观察到tunnel-te 500接口已针对指定范围动态创建。

```
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 表明“DYN_SR-TE_POLICIES”策略已附加到前缀，这意味着必须根据策略路由流量。

```
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
```

如果我们详细验证前缀10.0.5.55/32的BGP RIB，我们可以看到生成SR-TE隧道时将会引用的控制平面信息。

```
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
```

可以看到隧道策略处于打开状态并已注册。分配的绑定 SID 为 24000，此绑定 SID 可用于验证哪个隧道被用于此特定前缀。如前所述，tunnel-te500 已创建并安装在 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**：绑定SID有许多使用案例，对于此特定文档，限制其在本地验证中的使用，但其应用范围更广。

或者，您可以使用BGP RIB输出中的给定**if-handle 0x00000130**检查为前缀10.0.5.55/32分配的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（前端）上的SR-TE策略将具有从控制平面和数据平面角度转发流量的这些属性。此外，SR-TE隧道的状态信息也可以按输出查看，该输出必须与之前的验证相匹配。

```
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
```

直接在VRF BLUE RIB上检查前缀，我们可以确认绑定SID 24000已分配给前缀。

```
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.
```

FIB for VRF BLUE表示根据我们的BGP动态SR-TE策略，此前缀的转发通过tunnel-te 500完成。

```
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，我们可以验证连接并确认流量通过tunnel-te 500通过XR6的低延迟路径。

```
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
```

隧道 te500 的 XR2 计数器增加，这与我们的 SR-TE 策略相对应。

```
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
```

前缀10.0.5.5/32的路径仍在通过XR3通过常规LSP路径，如下所示。

```
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
```

故障排除

目前没有针对此配置的故障排除信息。

摘要

BGP 动态 SR-TE 为已启用 SR 的核心中的流量工程提供路由策略的粒度和自动实施。可以根据任意条件触发自动隧道创建，使网络管理员能够轻松创建符合站点应用要求的流量模式。

关于此翻译
思科采用人工翻译与机器翻译相结合的方式将此文档翻译成不同语言，希望全球的用户都能通过各自的语言得到支持性的内容。

请注意：即使是最好的机器翻译，其准确度也不及专业翻译人员的水平。

Cisco Systems, Inc. 对于翻译的准确性不承担任何责任，并建议您总是参考英文原始文档（已提供链接）。