

排除EIGRP外部路由优先于BGP的原因

目录

[简介](#)

[先决条件](#)

[设置和拓扑](#)

[问题陈述](#)

[问题摘要](#)

[故障排除和验证](#)

[为什么首选EIGRP路由而不是eBGP路由？](#)

[解决方案](#)

简介

本文档描述了ACI枝叶交换机通过EIGRP和eBGP接收同一路由时的路由行为。

先决条件

读者必须很好地了解ACI组件、术语和操作以及路由协议（EIGRP和BGP）。

设置和拓扑

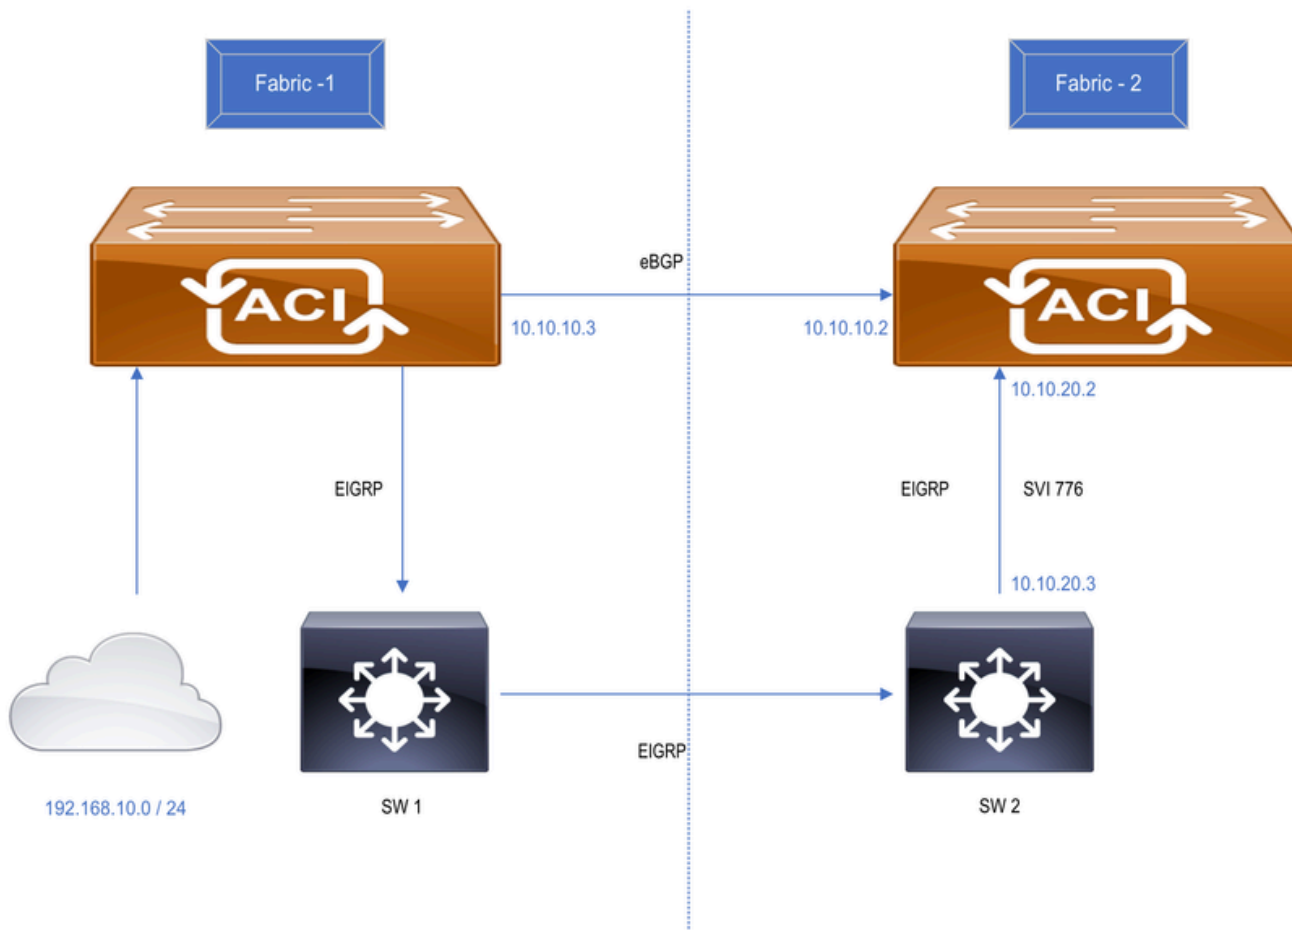


Figure : Two different ACI fabric connected over a WAN network

1. 此设置是使用两个不同的ACI交换矩阵完成的，连接如下：

- 两个DC边界枝叶交换机(BGP)之间的直接链路。
- 通过WAN网络(EIGRP)扩展。SW1和SW2是WAN交换机。

2. 192.168.10.0/24是连接到Fabric-1的内部ACI子网，并通过eBGP和EIGRP通告给Fabric-2。

问题陈述

Fabric-2边界枝叶交换机通过EIGRP和eBGP接收同一路由，其中eBGP路由按预期安装在交换机的路由表中。当eBGP会话断开时，EIGRP路由将安装到交换机的路由表中。即使eBGP启动，交换机仍保留EIGRP路由。这里的期望是，当eBGP会话启动时，必须立即将eBGP路由安装到路由表中，因为eBGP的AD值[20]比EIGRP [90]低。

问题摘要

- Fabric-1和Fabric-2数据中心通过WAN网络(EIGRP)和运行eBGP的两个站点BL交换机之间的直接链路连接。
- Fabric-1边界枝叶交换机正在通过eBGP和EIGRP向Fabric-2通告子网192.168.10.0/24。
- 两个L3Out位于同一个VRF中。
- 根据AD值，BGP路由将安装到Fabric-2边界枝叶交换机的路由表中。

- 当Fabric-1和Fabric-2之间的eBGP会话断开时，EIGRP路由会按预期安装到Fabric-2_BL交换机的路由表中。
- 当eBGP启动时，人们期望必须重新安装eBGP路由，并且将从路由表中删除EIGRP路由，但这种情况并未发生。
- Fabric-2边界枝叶交换机在其路由表中保留EIGRP路由。

故障排除和验证

- 验证Fabric-1和Fabric-2边界枝叶交换机之间的eBGP邻居关系。

```
Fabric-2_BL# show bgp sessions vrf snTn:snTn_VRF
Total peers 3, established peers 3
ASN 100
VRF snTn:snTn_VRF, local ASN 100
peers 1, established peers 1, local router-id 172.16.2.100
State: I-Idle, A-Active, O-Open, E-Established, C-Closing, S-Shutdown

Neighbor      ASN      Flaps LastUpDn|LastRead|LastWrit St Port(L/R)  Notif(S/R)
10.10.10.3    65001    2     1d23h   |never   |never   E 179/26051  45/6
```

- 检验Fabric-2上的EIGRP邻居关系。

```
Fabric-2_BL# show ip eigrp neighbors vrf snTn:snTn_VRF
EIGRP neighbors for process 500 VRF snTn:snTn_VRF
H  Address          Interface      Hold  Uptime  SRTT  RTT  Q  Seq
   (sec)            (ms)          Cnt  Num
0  10.10.20.3       vlan7         13   2d00h   1     50   0   8
```

```
SW-2# show ip eigrp neighbors VRF default
IP-EIGRP neighbors for process 500 VRF default
H  Address          Interface      Hold  Uptime  SRTT  RTT  Q  Seq
   (sec)            (ms)          Cnt  Num
0  10.10.20.2       Vlan776      14   2d00h   6     50   0   9
```

- 最初，BGP路由会安装在路由表中，并且相同路由会出现在交换矩阵2边界枝叶交换机的EIGRP拓扑表中。

```
Fabric-2_BL# show ip route 192.168.10.0/24 vrf snTn:snTn_VRF
IP Route Table for VRF "snTn:snTn_VRF"
'*' denotes best ucast next-hop
 '**' denotes best mcast next-hop
 '[x/y]' denotes [preference/metric]
```

'%<string>' in via output denotes VRF <string>

```
192.168.10.0/24, ubest/mbest: 1/0
```

```
*via 10.10.10.3%snTn:snTn_VRF, [20/0], 00:00:17, bgp-100, external, tag 65005
```

```
recursive next hop: 10.10.10.3/32%snTn:snTn_VRF
```

```
Fabric-2_BL# show ip eigrp topology 192.168.10.0/24 vrf snTn:snTn_VRF
```

```
EIGRP (AS 500): VRF: snTn:snTn_VRF , Topology entry for 192.168.10.0/24  
State is Passive, Query origin: Local origin, 0 Successor(s), FD is Infinity  
Routing Descriptor Blocks:
```

```
10.10.20.3(vlan7), from 10.10.20.3  
Urib State: in-rib,up-to-date  
Composite metric is (128576/128320), Route is Internal  
Vector metric:  
Minimum bandwidth is 8000000 Kbit  
Total delay is 5010 microseconds  
Reliability is 255/255  
Load is 1/255  
Minimum MTU is 1500  
Hop count is 1  
Internal tag is 0
```

- 当eBGP会话在Fabric-1和Fabric-2边界枝叶交换机之间断开时，EIGRP路由将安装到Fabric-2边界枝叶交换机的路由表中，并且即使在eBGP启动时仍保留EIGRP路由。

```
Fabric-2_BL# show ip route 192.168.10.0/24 vrf snTn:snTn_VRF
```

```
IP Route Table for VRF "snTn:snTn_VRF  
'*' denotes best ucast next-hop  
'**' denotes best mcast next-hop  
'[x/y]' denotes [preference/metric]  
'%<string>' in via output denotes VRF <string>
```

```
192.168.10.0/24, ubest/mbest: 1/0
```

```
*via 10.10.20.3, vlan7, [90/128576], 2d00h, eigrp-default, internal
```

- 这里的期望是，一旦eBGP会话启动，就必须将eBGP路由重新安装到路由表中。但Fabric-2_BL交换机只保留EIGRP路由。

为什么首选EIGRP路由而不是eBGP路由？

- 当eBGP会话断开时，Fabric-2_BL交换机在路由表中安装EIGRP路由，并且相同会重分发到MP-BGP以将其转发到Fabric-2中的其他服务枝叶交换机。
- 由于Fabric-2_BL交换机正在重新分发它，因此成为默认权重值为32768的路由的源。但是

- ，来自eBGP的路由权重为0。
- 由于较高的权重是首选，因此Fabric-2_BL交换机将重分发的路由视为最佳路由，并且不安装eBGP路由。
- 下面显示的输出是eBGP会话恢复的时间。

```
Fabric-2_BL# show ip bgp 192.168.10.0/24 vrf snTn:snTn_VRF
BGP routing table information for VRF snTn:snTn_VRF, address family IPv4 Unicast
BGP routing table entry for 192.168.10.0/24, version 28 dest ptr 0xa0fe0328
Paths: (2 available, best #1)
Flags: (0x80c0002 00000000) on xmit-list, is not in urib, exported
vpn: version 371, (0x100002) on xmit-list
Multipath: eBGP iBGP
```

```
Advertised path-id 1, VPN AF advertised path-id 1
Path type (0xa961d880): redistrib 0x408 0x1 ref 0 adv path ref 2, path is valid, is best path
AS-Path: NONE, path locally originated
Tx Domain path attribute Flag 0xc0,Code 36, Length 8, segment length 1
domain path: { <1:5345:128>}
```

```
0.0.0.0 (metric 0) from 0.0.0.0 (172.16.0.10)
Origin incomplete, MED 128576, localpref 100, weight 32768 tag 0, propagate 0
Extcommunity:
RT:100:2129921
VNID:2129921
COST:pre-bestpath:128:128576
COST:pre-bestpath:162:90
0x8800:32768:0 (Flags = 32768, Tag = 0)
0x8801:500:128256 (ASN = 500, Delay = 128256)
0x8802:65281:320 (Reliability = 255, Hop = 1, Bandwidth = 320)
0x8803:1:1500 (Reserve = 0, Load = 1, MTU = 1500)
0x8804:0:0 (Remote ASN = 0, Remote ID = 0)
0x8805:0:0 (Remote Prot = 0, Remote Metric = 0)
```

```
VPN AF advertised path-id 2
Path type (0xa961e0bc): external 0x28 0x0 ref 0 adv path ref 1, path is valid, not best reason: Weight
AS-Path: 65001 , path sourced external to AS
Source Domain: <1:16:128>
Tx Domain path attribute Flag 0xc0,Code 36, Length 15, segment length 2
domain path: { <1:5345:128>,<1:16:128>}
```

```
10.10.10.3 (metric 0) from 10.10.10.3 (172.16.1.100)
Origin IGP, MED not set, localpref 100, weight 0 tag 0, propagate 0
Extcommunity:
RT:100:2129921
VNID:2129921
```

```
VRF advertise information:
Path-id 1 not advertised to any peer
VPN AF advertise information:
Path-id 1 advertised to peers:
10.0.152.65 10.0.152.66
Path-id 2 not advertised to any peer
```

解决方案

解决此问题的方法有两种：

- LPM是解决方案之一：

1. 通过eBGP在EIGRP和/24掩码下使用/23掩码通告相同的子网，以使两个路由都存在于Fabric-2_BL交换机的路由表中。

<#root>

```
SW-2# show run interface vlan 776
```

```
!Command: show running-config interface Vlan776
!Time: Sun Jun 23 06:30:43 2024
```

```
version 7.0(3)I7(5) Bios:version 07.66
```

```
interface Vlan776
  no shutdown
  ip address 10.10.20.3/24
  ip router eigrp 500
  ip summary-address eigrp 500 192.168.10.0/23
```

```
>>>>> Advertised /23 via EIGRP
```

<#root>

```
Fabric-2_BL# show ip route vrf snTn:snTn_VRF
```

```
IP Route Table for VRF "snTn:snTn_VRF"
```

```
'*' denotes best ucast next-hop
```

```
'**' denotes best mcast next-hop
```

```
'[x/y]' denotes [preference/metric]
```

```
'%<string>' in via output denotes VRF <string>
```

```
192.168.10.0/23, ubest/mbest: 1/0
```

```
*via 10.10.20.3, vlan20, [90/128576], 00:24:11, eigrp-default, internal
```

```
>>>>>>>> EIGRP Route
```

```
192.168.10.0/24, ubest/mbest: 1/0
```

```
*via 10.10.10.3%snTn:snTn_VRF, [20/0], 00:04:12, bgp-100, external, tag 65005
```

```
>>>>>>>> BGP Route
```

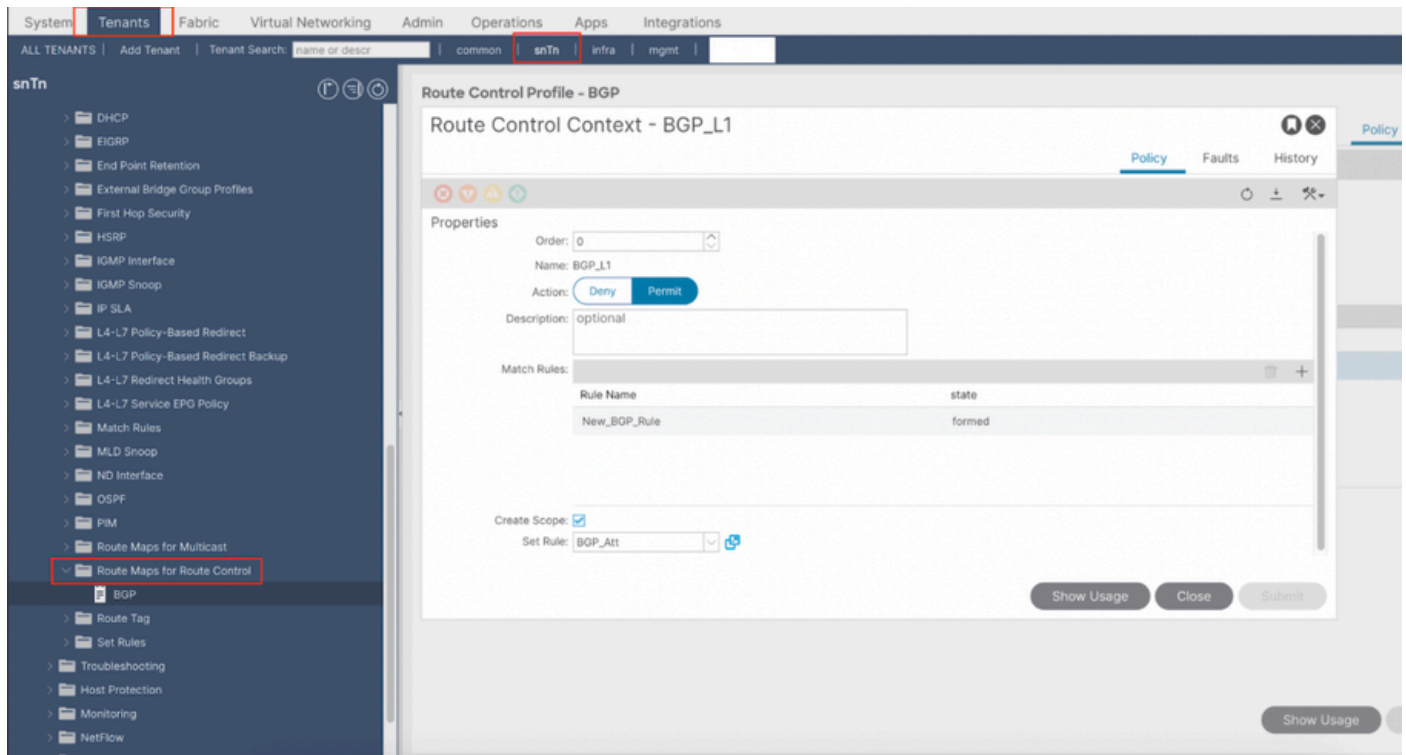
b.当eBGP会话断开时，EIGRP路由仍存在于路由表中以提供冗余。

c.一旦BGP会话启动，BGP路由将重新安装到路由表中并优先用于流量转发。

- 对eBGP路由应用权重：

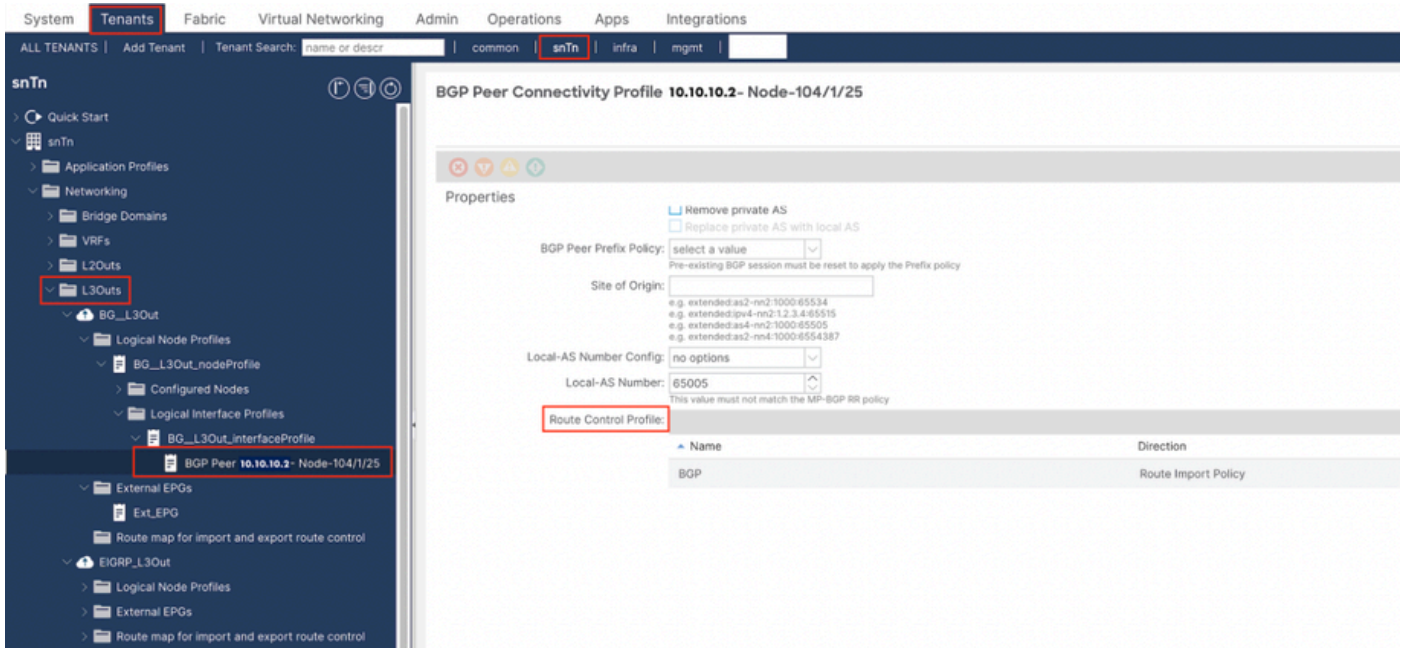
1. 如果需要通过EIGRP和BGP通告具有相同子网掩码的子网，可以对eBGP路由应用更高的权重(大于32768)以始终作为首选路由。
2. 如何在ACI上应用权重：
 1. 创建路由映射策略。

Tenant ----> Policies ----> Route Maps for Route Control (右键单击并创建新策略，填写所有必需的详细信息) ----> Create "Set Rule" policy —>选择“Weight”属性策略并输入值



ii. 将路由映射应用到L3Out：

Tenant —> Networking —> L3Out ----> Logical Node Profiles —> Node Profile ----> Logical Interface Profile —> Interface Profile —> Peer Profile —> 点击“Route Control Profile”下的“+”并选择已创建的新路由映射



<#root>

```
Fabric-2_BL# show ip bgp 192.168.10.0/24 vrf snTn:snTn_VRF
```

```
BGP routing table information for VRF snTn:snTn_VRF, address family IPv4 Unicast
BGP routing table entry for 192.168.10.0/24, version 61 dest ptr 0xa0fa3f70
Paths: (1 available, best #1)
Flags: (0x80c001a 00000000) on xmit-list, is in urib, is best urib route, is in HW, exported
vpn: version 79, (0x100002) on xmit-list
Multipath: eBGP iBGP
```

```
Advertised path-id 1, VPN AF advertised path-id 1
Path type (0xa95a2d5c): external 0x28 0x0 ref 0 adv path ref 2, path is valid, is best path
AS-Path: 65005 65001 , path sourced external to AS
Source Domain: <1:16:128>
Tx Domain path attribute Flag 0xc0, Code 36, Length 15, segment length 2
domain path: { <1:5345:128>, <1:16:128> }
10.10.10.3 (metric 0) from 10.10.10.3 (172.16.0.10)
Origin IGP, MED not set, localpref 100, weight 32769 tag 0, propagate 0
Extcommunity:
RT:100:2129921
VNID:2129921
```

```
VRF advertise information:
Path-id 1 not advertised to any peer
```

```
VPN AF advertise information:
Path-id 1 advertised to peers:
10.0.152.65 10.0.152.66
```

c.这里的关键点是，当BGP会话启动时，您不会在BGP表中看到重分发的EIGRP路由。原因是EIGRP外部路由的FD设置为无限。

<#root>


```
Fabric-2_BL# show ip eigrp topology vrf snTn:snTn_VRF
```

```
EIGRP Topology Table for AS(500)/ID(172.16.2.100) VRF snTn:snTn_VRF  
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,  
r - reply Status, s - sia Status
```

```
  P 192.168.10.0/24, 0 Successors, FD is Infinity  
    via 10.10.20.3(128576/128320), vlan20
```

d. “FD is Infinity”消息实际上是EIGRP中指示RIB由于存在管理距离更短的路由而拒绝了该路由。

e.当BGP会话断开时，EIGRP路由只会重分发到MP-BGP并安装到fabric-2_BL交换机的路由表中。

```
<#root>
```

```
Fabric-2_BL# show ip bgp summary vrf snTn:snTn_VRF
```

```
BGP summary information for VRF snTn:snTn_VRF, address family IPv4 Unicast  
BGP router identifier 172.16.2.100, local AS number 100  
BGP table version is 65, IPv4 Unicast config peers 1, capable peers 0  
6 network entries and 6 paths using 1248 bytes of memory  
BGP attribute entries [4/704], BGP AS path entries [0/0]  
BGP community entries [0/0], BGP clusterlist entries [2/8]
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.10.10.3	4	65001	18530	18554	0	0	0	00:04:25	Idle

```
<#root>
```

```
Fabric-2_BL# show ip eigrp topology vrf snTn:snTn_VRF
```

```
IP-EIGRP Topology Table for AS(500)/ID(172.16.2.100) VRF snTn:snTn_VRF  
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,  
r - reply Status, s - sia Status
```

```
  P 192.168.10.0/24, 1 successors, FD is 128576  
    via 10.10.20.3 (128576/128320), Vlan20
```

```
<#root>
```

```
Fabric-2_BL# show ip route vrf snTn:snTn_VRF
```

IP Route Table for VRF "snTn:snTn_VRF"

'*' denotes best ucast next-hop

'**' denotes best mcast next-hop

'[x/y]' denotes [preference/metric]

'%<string>' in via output denotes VRF <string>

192.168.10.0/24, ubest/mbest: 1/0

*via 10.10.20.3, Vlan20, [90/128576], 02:31:52, eigrp-default, internal >>>>>> EIGRP Route

<#root>

Fabric-2_BL# show ip bgp 192.168.10.0/24 vrf snTn:snTn_VRF

BGP routing table information for VRF snTn:snTn_VRF, address family IPv4 Unicast

BGP routing table entry for 192.168.10.0/24, version 65 dest ptr 0xa0fa3f70

Paths: (1 available, best #1)

Flags: (0x80c0002 00000000) on xmit-list, is not in urib, exported

vpn: version 83, (0x100002) on xmit-list

Multipath: eBGP iBGP

Advertised path-id 1, VPN AF advertised path-id 1

Path type (0xa95a2c64): redist 0x408 0x1 ref 0 adv path ref 2, path is valid, is best path

AS-Path: NONE, path locally originated

Tx Domain path attribute Flag 0xc0, Code 36, Length 8, segment length 1

domain path: { <1:5345:128>}

0.0.0.0 (metric 0) from 0.0.0.0 (172.16.0.10)

Origin incomplete, MED 128576, localpref 100, weight 32768 tag 0, propagate 0

Extcommunity:

RT:100:2129921

VNID:2129921

COST:pre-bestpath:128:128576

COST:pre-bestpath:162:90

0x8800:32768:0 (Flags = 32768, Tag = 0)

0x8801:500:128256 (ASN = 500, Delay = 128256)

0x8802:65281:320 (Reliability = 255, Hop = 1, Bandwidth = 320)

0x8803:1:1500 (Reserve = 0, Load = 1, MTU = 1500)

0x8804:0:0 (Remote ASN = 0, Remote ID = 0)

0x8805:0:0 (Remote Prot = 0, Remote Metric = 0)

VRF advertise information:

Path-id 1 not advertised to any peer

VPN AF advertise information:

Path-id 1 advertised to peers:

10.0.152.65 10.0.152.66

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）。