

排除為什麼首選EIGRP外部路由而非BGP的故障

目錄

[簡介](#)

[必要條件](#)

[設定和拓撲](#)

[問題陳述](#)

[問題摘要](#)

[疑難排解和驗證](#)

[為什麼EIGRP路由優先於eBGP路由？](#)

[解決方案](#)

簡介

本文檔介紹當ACI枝葉交換機透過EIGRP和eBGP接收同一路由時，其路由行為。

必要條件

讀者必須很好地瞭解ACI元件、術語和操作以及路由協定（EIGRP和BGP）。

設定和拓撲




Figure : Two different ACI fabric connected over a WAN network

1. 此設定是使用兩個不同的ACI網狀架構來完成的，連線方式如下：

- 兩個DC邊界枝葉交換機(BGP)之間的直接鏈路。
- 透過WAN網路(EIGRP)擴展。SW1和SW2是WAN交換機。

2. 192.168.10.0/24是連線到Fabric-1的內部ACI子網，透過eBGP和EIGRP通告給Fabric-2。

問題陳述

Fabric-2邊界枝葉交換機透過EIGRP和eBGP接收相同的路由，其中eBGP路由按預期安裝在交換機的路由表中。當eBGP會話斷開時，EIGRP路由將安裝在交換機的路由表中。即使eBGP啟動，交換機仍保留EIGRP路由。這裡的期望是，eBGP路由必須在eBGP會話啟動後立即安裝到路由表中，因為eBGP的AD值[20]小於EIGRP [90]。

問題摘要

- Fabric-1和Fabric-2資料中心透過WAN網路(EIGRP)以及兩個站點之間的直接鏈路連線，運行eBGP的BL交換機。
- Fabric-1邊界枝葉交換機透過eBGP和EIGRP向Fabric-2通告子網192.168.10.0/24。
- 兩個L3Out都位於同一個VRF中。
- 根據AD值，BGP路由被安裝到Fabric-2邊界枝葉交換機的路由表中。

- 當Fabric-1和Fabric-2之間的eBGP會話斷開時，EIGRP路由將安裝在預期的Fabric-2_BL交換機的路由表中。
- 當eBGP啟動時，期望的是eBGP路由必須重新安裝，而EIGRP路由將從路由表中刪除，但路由表沒有出現。
- Fabric-2邊界枝葉交換機在其路由表中保留EIGRP路由。

疑難排解和驗證

- 驗證Fabric-1和Fabric-2邊界枝葉交換機之間的eBGP鄰居關係。

```
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

- 檢驗Fabric-2上的EIGRP鄰居關係。

```
Fabric-2_BL# show ip eigrp neighbors vrf snTn:snTn_VRF
EIGRP neighbors for process 500 VRF snTn:snTn_VRF
```

H	Address	Interface	Hold (sec)	Uptime	SRTT (ms)	RT0	Q Cnt	Seq 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 (sec)	Uptime	SRTT (ms)	RT0	Q Cnt	Seq Num
0	10.10.20.2	Vlan776	14	2d00h	6	50	0	9

- 最初，BGP路由會安裝在路由表中，而相同的路由會出現在交換矩陣2邊界枝葉交換機的EIGRP拓撲表中。

```
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
```

- 當eBGP會話在Fabric-1和Fabric-2邊界枝葉交換機之間斷開時，EIGRP路由將安裝在交換矩陣-2邊界枝葉交換機的路由表中，並且即使在eBGP啟動時仍保留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
```

- 這裡的期望是，一旦eBGP會話啟動，eBGP路由必須重新安裝到路由表中。但Fabric-2_BL交換機只保留EIGRP路由。

為什麼EIGRP路由優先於eBGP路由？

- 當eBGP會話關閉時，Fabric-2_BL交換機在路由表中安裝EIGRP路由，並將該路由重分配到MP-BGP中，以將其轉發到Fabric-2中的其他服務枝葉交換機。
- 由於Fabric-2_BL交換機正在重新分發它，因此會成為預設權重值為32768的該路由的源。但

是，來自eBGP的路由的權重為0。

- 由於權重較高是首選路由，因此Fabric-2_BL交換機將重分發的路由視為最佳路由，並且不安裝eBGP路由。
- 下面顯示的輸出是eBGP會話恢復運行的時間。

```
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
```

解決方案

解決此問題的方法有兩種：

- LPM是解決方案之一：

1. 透過eBGP在EIGRP和/24掩碼下使用/23掩碼通告相同的子網，以便兩條路由都出現在Fabric-2_BL交換機的路由表中。

<#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
```

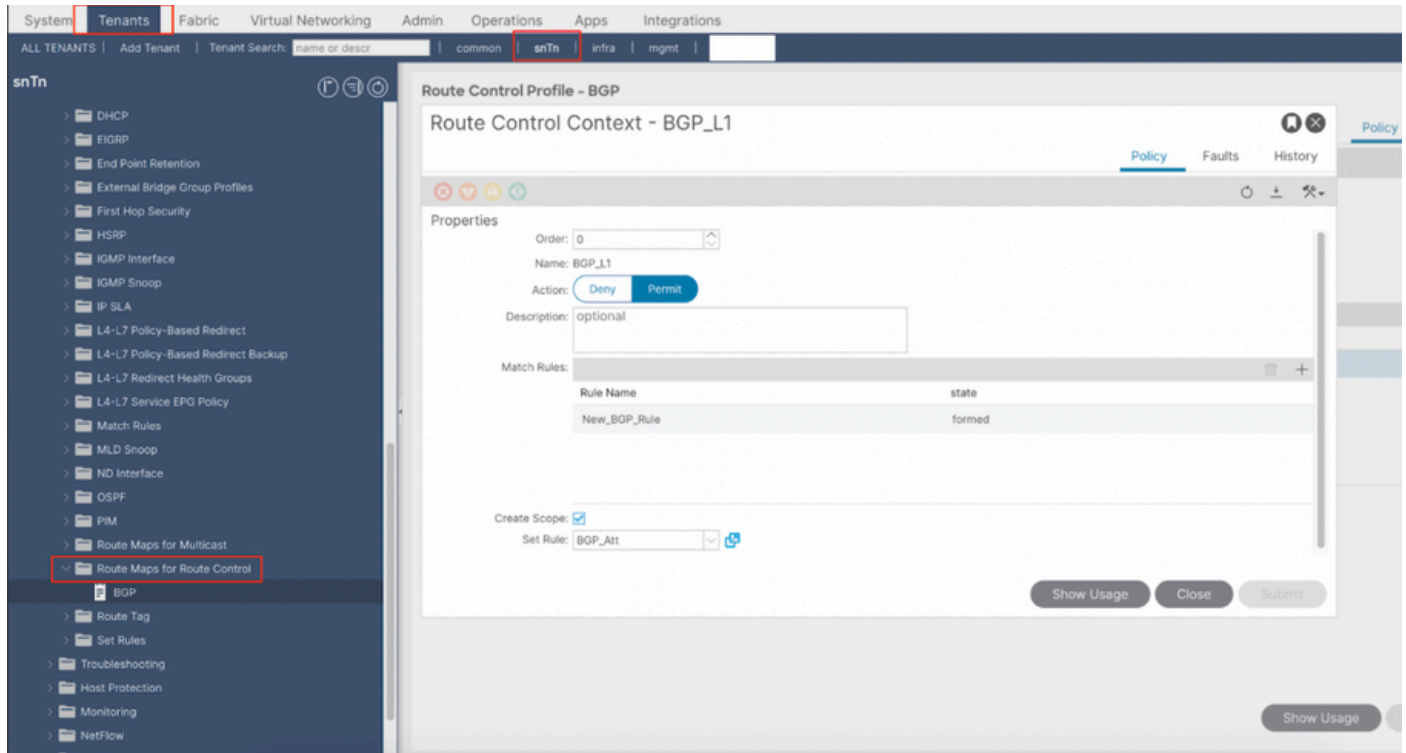
b.當eBGP會話斷開時，EIGRP路由仍存在於路由表中以實現冗餘。

c.一旦BGP會話啟動，BGP路由就會被重新安裝到路由表中並優先用於流量轉發。

- 在eBGP路由上應用權重：

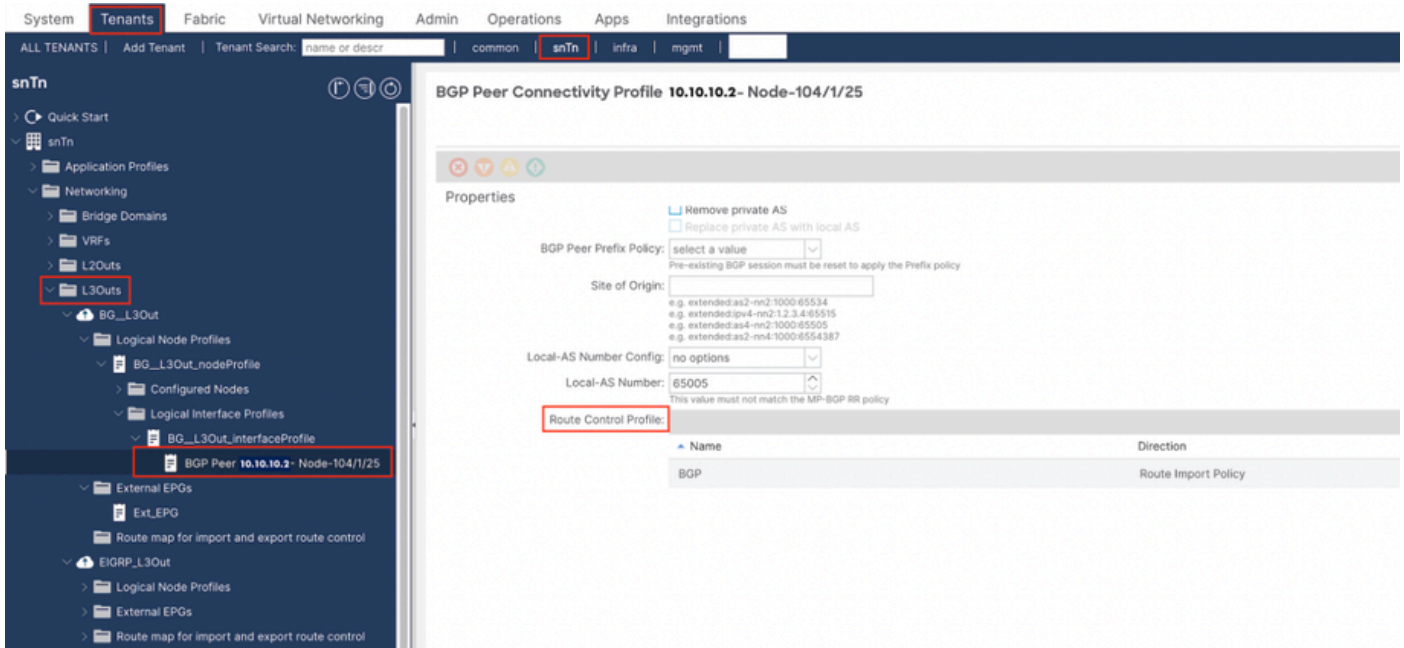
1. 如果需要透過EIGRP和BGP通告具有相同子網掩碼的子網，則可以對eBGP路由應用更高的權重(大於32768)以始終作為首選路由。
2. 如何對ACI應用權重：
 1. 建立路由對映策略。

租戶---->策略---->路由控制的路由對映 (按一下右鍵並建立新策略，填寫所有必需的詳細資訊) ---->建立「設定規則」策略—>選擇「權重」屬性策略並輸入值



ii. 將路由對映應用到L3Out：

租戶—> Networking —> L3Out ----> Logical Node Profiles —> Node Profile ----> Logical Interface Profile —> Interface Profile —> Peer Profile —>按一下「Route Control Profile」下的「+」並選擇已建立的新路由對映



<#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.此處所捕獲的是，當BGP會話啟動時，您在BGP表中看不到重分發的EIGRP路由。原因是EIGRP外部路由的FD設定為無限。

<#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), v1an20
```

- d. 「FD is Infinity」消息實際上是EIGRP中指示RIB由於存在管理距離較低的路由而拒絕該路由。
- e.當BGP會話斷開時，EIGRP路由只會重分配到MP-BGP中並安裝到fabric-2_BL交換機的路由表中。

```
<#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), v1an20
```

```
<#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

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）。