nexus上透過vPC進行的eBGP對等

目錄

簡介

本檔案將說明如何設定和驗證vPC Nxus配對與其他裝置之間的e-Bordere Gateway Protocol(eBGP)對等。為清楚起見，外部裝置上的配置顯示為Cisco CLI NX-OS。

必要條件

需求

思科建議您瞭解以下主題：

- vPC操作和配置基本概念。供參考：https://www.cisco.com/c/dam/en/us/td/docs/switches/datacenter/sw/design/vpc_design/vpc_best_
- BGP操作和配置

採用元件

本文中的資訊係根據以下軟體和硬體版本：

leaf1#	C93108TC-FX	NXOS 9.3(3)
leaf2#	C93108TC-FX	NXOS 9.3(3)
外部裝置	N9K-C9396PX	NXOS：版本9.2(3)

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

vPC對和外部裝置之間的路由協定(OSPF、ISIS、RIP、EIGRP、BGP)對等。支援 https://www.cisco.com/c/en/us/support/docs/ip/ip-routing/118997-technote-nexus-00.html。本文介紹eBGP作為路由協定的其他說明和配置示例。

嘗試在vPC上啟用路由協定對等時會出現一些問題，使用時不存在這些問題

standard port-channel:

1. 無法確定外部裝置將用於為每個mac地址轉發流量的埠通道中的哪一個成員。外部裝置可能會通過鏈路將枝葉1的BGP資料包傳送到枝葉1。當Leaf-1收到它時，它會將它傳送到CPU並丟棄它（不是自己的IP地址），因此BGP（單播）或其他協定（組播）將不斷擺動。以下幫助命令peer-gateway。
2. 即使使用peer-gateway指令，此類封包上的TTL也會下降。NX-OS中的新命令 — layer3 peer-router 禁用該命令。
3. 兩個vPC成員之間的iBGP需要遵守所有鄰居之間的iBGP的BGP規則。我們在vPC端的vrf中運行，因此只有這兩個成員需要運行iBGP。在鏈路故障至NX-OS網路的其餘部分（VXlan或其他）並提供冗餘的情況下，也需要使用TH。

設定

網路圖表

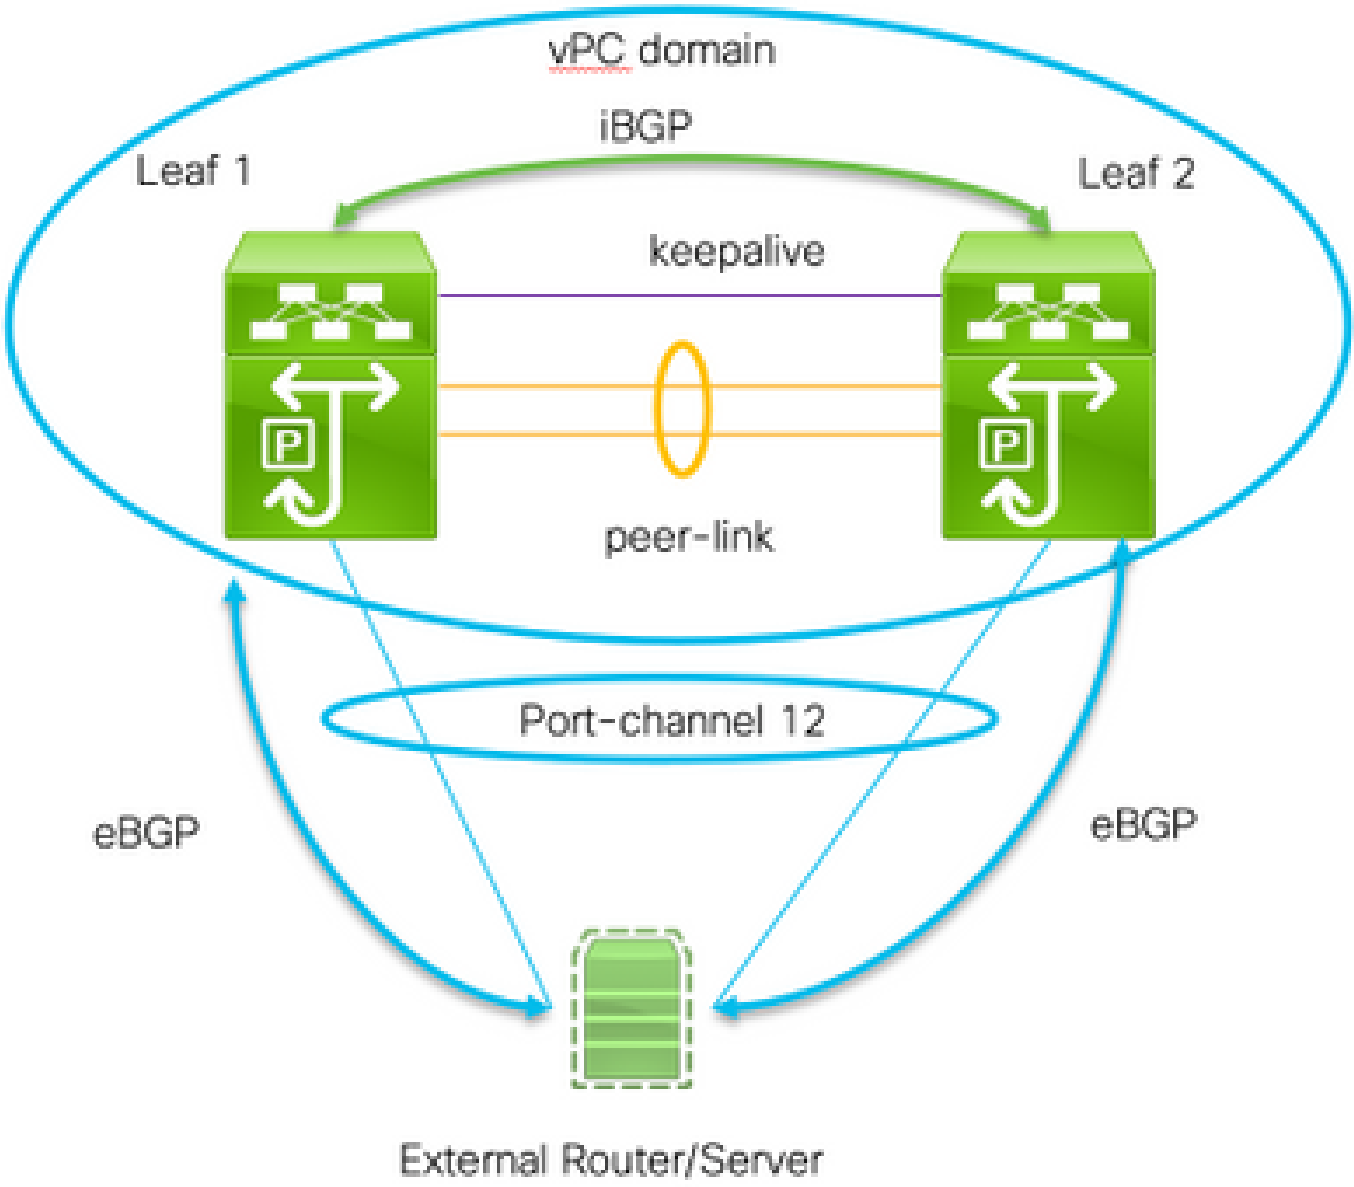

啟用此對等需要兩個重要命令：

- 對等網關 - vPC對等網關功能允許vPC交換機充當將資料包傳送到vPC對等裝置的路由器MAC地址的活動網關
- 第3層對等路由器 — 目的地為對等體的資料包的TTL不變，從層路由協定對等角度來看，外部裝置將vPC域視為單個物理實體。

組態

Leaf 1:

```
! Form the vPC domain:
vpc domain 1
  peer-switch
  role priority 10 peer-keepalive destination 192.0.2.2 source 192.0.2.1 peer-gateway layer3 peer-router
```

```
!
!vPC peer-link interface members
interface Ethernet1/53 - 54
  description vPC-Peerlink member
  switchport
  switchport mode trunk
  channel-group 11 mode active
  no shutdown
!
! vPC peer-link port-channel
interface port-channel11
  description vPC-peerlink
  switchport
  switchport mode trunk
  spanning-tree port type network
  no shutdown
  vpc peer-link
!
! vPC port-channel member to External Device
interface Ethernet1/52
  description ExternalDevice Eth2/13
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 203,205
  mtu 9216
  channel-group 12 mode active
  no shutdown
!
! vPC port-channel to External Device
interface port-channel12
  description vPC port-channel to External Device
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 203,205
  mtu 9216
  vpc 12
!
! Layer 3 interface to the Eternal device:
interface Vlan205
  no shutdown
  vrf member Customer
! BFD for eBGP
  bfd interval 500 min_rx 500 multiplier 3
! Disable bfd echo, as it is not supported over vPC
  no bfd echo
  no ip redirects
! We use/29 as we need 3 ip address, one per each member of the ! vPC domain and 3rd for the External Device
  ip address 198.51.100.1/29 tag 800204
! Disable redirects - this is needed to enable BFD
  no ipv6 redirects
!
router bgp 65535
router bgp 65535
  router-id 203.0.113.1
  log-neighbor-changes
  address-family ipv4 unicast
   Customer router-id 198.51.100.1 address-family ipv4 unicast neighbor 198.51.100.2 description Leaf-2 remote-as 65535 address-family ipv4 unicast soft
! Form the vPC domain:
vpc domain 1
```

```
  peer-switch
  role priority 10
  peer-keepalive destination 192.0.2.1 source 192.0.2.2
  peer-gateway
  layer3 peer-router
  ipv6 nd synchronize
  ip arp synchronize
!
!vPC peer-link interface members
interface Ethernet1/53 - 54
  description vPC-Peerlink member
  switchport
  switchport mode trunk
  channel-group 11 mode active
  no shutdown
!
! vPC peer-link port-channel
interface port-channel11
  description vPC-peerlink
  switchport
  switchport mode trunk
  spanning-tree port type network
  no shutdown
  vpc peer-link
!
! vPC port-channel member to External Device
interface Ethernet1/52
  description ExternalDevice Eth2/13
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 203,205
  mtu 9216
  channel-group 12 mode active
  no shutdown
!
! vPC port-channel to External Device
interface port-channel12
  description vPC port-channel to External Device
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 203,205
  mtu 9216
  vpc 12
!
! Layer 3 interface to the Eternal device:
interface Vlan205
  no shutdown
  vrf member Customer
! BFD for eBGP
  bfd interval 500 min_rx 500 multiplier 3
! Disable bfd echo, as it is not supported over vPC
  no bfd echo
  no ip redirects
! We use/29 as we need 3 ip address, one per each member of the ! vPC domain and 3rd for the External Device
  ip address 198.51.100.2/29 tag 800204
! Disable redirects - this is needed to enable BFD
  no ipv6 redirects
!
router bgp 65535
```

```
router bgp 65535
  router-id 203.0.113.2
  log-neighbor-changes
  address-family ipv4 unicast
 vrf Customer
    router-id 198.51.100.2
    address-family ipv4 unicast
    neighbor 198.51.100.1
      description Leaf-2
      remote-as 65535
      address-family ipv4 unicast
        soft-reconfiguration inbound always
    neighbor 198.51.100.3
      description to External Device
      bfd
      remote-as 65000
      update-source Vlan205
      address-family ipv4 unicast
        soft-reconfiguration inbound always
!
External Device (NX-OS style CLI):
interface Ethernet2/13 - 14
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 203,205
  mtu 9216
  channel-group 12 mode active
  no shutdown
!
interface port-channel12
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 203,205
  mtu 9216
  no shutdown
!
interface Vlan205
  no shutdown
  mtu 9216
! See notes in Leaf-1 and Leaf 2 for BFD
  bfd interval 500 min_rx 500 multiplier 3
  no bfd echo
  no ip redirects
  ip address 198.51.100.3/29
  no ipv6 redirects
!
router bgp 65000
  log-neighbor-changes
  address-family ipv4 unicast
   neighbor 198.51.100.1 remote-as 65535
      description to Leaf-1
      update-source Vlan205
      bfd
   neighbor 198.51.100.2 remote-as 65535
      description to Leaf-2
      update-source Vlan205
      bfd
end
!
```

驗證

以下是show bgp ipv4單播鄰居的輸出。它驗證了：

1. BGP鄰居關係已建立且穩定
2. 已在外部鄰居之間啟用BFD

Leaf 1/2: show bgp ipv4 unicast neighbors vrf Customer BGP neighbor is 203.0.113.2, remote AS 65535, ibgp link, Peer index 4 BGP version 4, remote router ID 203.0.113.2 Neighbor previous state = OpenConfirm BGP state = Established, up for 6d22h Neighbor vrf: Customer Peer is directly attached, interface Vlan205 Last read 00:00:14, hold time = 180, keepalive interval is 60 seconds Last written 00:00:03, keepalive timer expiry due 00:00:56 Received 10012 messages, 0 notifications, 0 bytes in queue ... BGP neighbor is 203.0.113.2.3, remote AS 65000, ebgp link, Peer index 3 BGP version 4, remote router ID 203.0.113.2 Neighbor previous state = OpenConfirm BGP state = Established, up for 1d00h Neighbor vrf: Customer Using Vlan205 as update source for this peer Peer is directly attached, interface Vlan205 BFD live-detection is configured and enabled, state is Up Last read 00:00:22, hold time = 180, keepalive interval is 60 seconds Last written 00:00:56, keepalive timer expiry due 00:00:03 ! External Device: show bgp ipv4 unicast neighbors BGP neighbor is 203.0.113.1, remote AS 65535, ebgp link, Peer index 3 Inherits peer configuration from peer-template Cust_BGP_Peer BGP version 4, remote router ID 203.0.113.1 BGP state = Established, up for 1d00h Peer is directly attached, interface Vlan205 Enable logging neighbor events BFD live-detection is configured and enabled, state is Up Last read 0.660288, hold time = 180, keepalive interval is 60 seconds Last written 00:00:26, keepalive timer expiry due 00:00:33 Received 10122 messages, 1 notifications, 0 bytes in queue Sent 10086 messages, 1 notifications, 0(0) bytes in queue Connections established 14, dropped 13 Last reset by us 1d00h, due to bfd session down Last reset by peer 6d22h, due to other configuration change ....

疑難排解

以下命令將有助於驗證操作：

show vpc show vpc consistency-parameters global show vpc consistency-parameters interface

show bgp ipv4 unicast neighbors show bgp ipv4 unicast summary

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）。