瞭解基本BGP疑難排解

目錄

簡介

本檔案介紹進行邊界閘道通訊協定(BGP)疑難排解和瞭解Nexus基本輸出的程式。

必要條件

需求

思科建議您瞭解以下主題：

- Nexus交換機
- BGP

採用元件

本文件所述內容不限於特定軟體和硬體版本。

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

BGP是一種外部網關協定，用於大規模網路，以在自治系統(AS)之間交換路由和可達性資訊。它是支援Internet全域性路由系統的核心路由協定。

確定為BGP對等體實施的配置

驗證Nexus 9300上啟用了BGP功能。

識別並理解BGP進程的配置。

```
switch# show running-config bgp

!Command: show running-config bgp
!Running configuration last done at: Tue Jul 18 19:45:05 2023
!Time: Tue Jul 18 19:45:44 2023

version 10.2(4) Bios:version 05.47
feature bgp

router bgp 64512                                -> The AS number of the local BGP speaker.
  router-id 172.17.255.255                      -> BGP speaker identifier.
  address-family ipv4 unicast                   -> Global address family configuration.
    network 10.100.1.0/24                       -> Specifies a network as local to this autonomous system
 redistribute direct route-map permit_all       -> Routes redistribution from other routing protocols (OS
 neighbor 10.1.1.1                              -> IP address of the remote BGP peer.
    remote-as 64512                             -> The AS number of the remote BGP peer.*
    address-family ipv4 unicast                 -> Local address family configuration.
      prefix-list allow_in in                   -> Prefix-list applied at the inbound of the BGP peer. **
      prefix-list allow_out out                 -> Prefix-list applied at the outbound of the BGP peer. *
      soft-reconfiguration inbound always       -> Store the inbound BGP route updates.
  neighbor 172.18.255.255
    remote-as 65535
    update-source loopback10                    -> Interface used to source BGP updates.
    ebgp-multihop 3                             -> Maximum hops to reach peer IP address, it modifies the
    address-family ipv4 unicast
      route-map block_route in                  -> Route-map applied at the inbound of the BGP peer. **
      route-map no_local out                    -> Route-map applied at the outbound of the BGP peer. **
      soft-reconfiguration inbound always
```

注意：在本地和遠端配置的相同ASN用於標識iBGP會話，而配置的不同ASN用於標識eBGP會話。

路由對映的層次結構值高於應用於同一對等體的字首清單。

如何理解基本BGP輸出

BGP對等體狀態

```
switch# show ip bgp summary
BGP summary information for VRF default, address family IPv4 Unicast
BGP router identifier 172.17.255.255, local AS number 64512 -> Local BGP ID and ASN
BGP table version is 67, IPv4 Unicast config peers 2, capable peers 2
20 network entries and 19 paths using 5424 bytes of memory
BGP attribute entries [6/2112], BGP AS path entries [2/20]
BGP community entries [0/0], BGP clusterlist entries [0/0]
13 received paths for inbound soft reconfiguration
12 identical, 0 modified, 1 filtered received paths using 96 bytes
```

```
Neighbor        V AS     MsgRcvd   MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
10.1.1.1        4 64512   346      334     67      0     0 05:25:12 8    -> BGP peer IP address, remote A
172.18.255.255 4 65535   334      327     67      0     0 05:18:00 8
```

註：「開啟/關閉」部分顯示BGP對等會話已開啟或關閉的時間。
State/PfxRcd部分顯示會話當前所在的BGP狀態。它顯示建立狀態時從對等體學習的字首數。

BGP狀態	
空閒	這是BGP等待「開始事件」的第一個狀態。當某人配置新的BGP鄰居或重置已建立的BGP對等時，會發生start事件。
連線	BGP正在等待TCP三次握手完成。成功後，它會繼續進入OpenSent狀態。如果它失敗，我們繼續進入Active狀態。
Active（作用中）	BGP會嘗試另一個TCP三次握手，以與遠端BGP鄰居建立連線。如果成功，則移至OpenSent狀態。如果ConnectRetry計時器過期，則我們返回到連線狀態。
OpenSent	在此狀態下，BGP等待來自遠端BGP鄰居的Open消息。
OpenConfirm	BGP等待來自遠端BGP鄰居的keepalive消息。
已建立	BGP鄰居鄰接關係完成，BGP路由器將傳送更新資料包以交換路由資訊。

瞭解BGP對等體資訊。

```
switch# show ip bgp neighbors 10.1.1.1
BGP neighbor is 10.1.1.1, remote AS 64512, ibgp link, Peer index 3            -> Peer IP address, remo
  BGP version 4, remote router ID 172.16.255.255                              -> BGP version, Peer Rou
  Neighbor previous state = OpenConfirm                                       -> Previous BGP state
  BGP state = Established, up for 00:05:29                                    -> Current BGP state and
  Neighbor vrf: default                                                       -> VRF used for the peer
  Peer is directly attached, interface Ethernet1/49                           -> Interface used to sou
  Last read 00:00:28, hold time = 180, keepalive interval is 60 seconds       -> Amount of time from l
  Last written 00:00:28, keepalive timer expiry due 00:00:31                  -> Amount of time from l
  Received 363 messages, 0 notifications, 0 bytes in queue
  Sent 354 messages, 1 notifications, 0(0) bytes in queue
  Enhanced error processing: On
    0 discarded attributes
  Connections established 2, dropped 1                                        -> Counters informing th
  Last update recd 00:05:28, Last update sent  = 00:05:28                     -> Timer of the last maj
   Last reset by us 00:06:21, due to holdtimer expired error                  -> Last reset timer and
  Last error length sent: 0
  Reset error value sent: 0
  Reset error sent major: 4 minor: 0
  Notification data sent:
  Last reset by peer never, due to No error
  Last error length received: 0
  Reset error value received 0
  Reset error received major: 0 minor: 0
  Notification data received:
```

```
  Neighbor capabilities:
  Dynamic capability: advertised (mp, refresh, gr) received (mp, refresh, gr)
  Dynamic capability (old): advertised received
  Route refresh capability (new): advertised received
  Route refresh capability (old): advertised received
  4-Byte AS capability: advertised received
  Address family IPv4 Unicast: advertised received
  Graceful Restart capability: advertised received

  Graceful Restart Parameters:
  Address families advertised to peer:
    IPv4 Unicast
  Address families received from peer:
    IPv4 Unicast
  Forwarding state preserved by peer for:
  Restart time advertised to peer: 120 seconds
  Stale time for routes advertised by peer: 300 seconds
  Restart time advertised by peer: 120 seconds
  Extended Next Hop Encoding Capability: advertised received
  Receive IPv6 next hop encoding Capability for AF:
    IPv4 Unicast  VPNv4 Unicast

  Message statistics:
                              Sent               Rcvd
  Opens:                         2                  2
  Notifications:                 1                  0
  Updates:                      22                 20
  Keepalives:                  340                339
  Route Refresh:                 1                  0
  Capability:                    2                  2
  Total:                       354                363
  Total bytes:                7949               7524
  Bytes in queue:                0                  0

  For address family: IPv4 Unicast
  BGP table version 88, neighbor version 88
  8 accepted prefixes (8 paths), consuming 2176 bytes of memory
  0 received prefixes treated as withdrawn
  11 sent prefixes (11 paths)                                                  -> Amount of prefixes ad
  Inbound soft reconfiguration allowed(always)
  Third-party Nexthop will not be computed.
  Inbound ip prefix-list configured is allow_in, handle obtained
  Outbound ip prefix-list configured is allow_out, handle obtained
  Last End-of-RIB received 00:00:01 after session start
  Last End-of-RIB sent 00:00:01 after session start
  First convergence 00:00:01 after session start with 11 routes sent

  Local host: 10.1.1.2, Local port: 28262
  Foreign host: 10.1.1.1, Foreign port: 179
  fd = 73
```

瞭解BGP表

此輸出顯示BGP表中獲取的所有字首的狀態、下一跳、度量、本地優先順序、權重和AS路徑。

```
switch# show ip bgp
```

```
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 88, Local Router ID is 172.17.255.255
Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best
Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-injected
Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - best2

   Network            Next Hop            Metric     LocPrf     Weight Path
*>r10.1.1.0/30        0.0.0.0                  0        100      32768 ?
*>i10.100.1.0/24      10.1.1.1                           100          0 i
*>i10.100.2.0/24      10.1.1.1                           100          0 i
*>i10.100.3.0/24      10.1.1.1                           150          0 i
*>i10.100.4.0/24      10.1.1.1                 0         100          0 ?
*>i10.100.5.0/24      10.1.1.1                 0         100          0 ?
*>i10.100.6.0/24      10.1.1.1                 0         100          0 ?
*>i10.100.7.0/24      10.1.1.1                 0         100          0 ?
*>i10.100.8.0/24      10.1.1.1                 0         100          0 ?
*>r172.17.255.255/32  0.0.0.0                  0         100      32768 ?
*>e172.30.1.0/24      172.18.255.255           0                      0 65535 ?
*>e172.30.2.0/24      172.18.255.255           0                      0 65535 ?
*>e172.30.3.0/24      172.18.255.255           0                      0 65535 ?
*>e172.30.4.0/24      172.18.255.255           0                      0 65535 ?
*>e172.30.5.0/24      172.18.255.255           0                      0 65535 65534 65533 ?
*>e172.30.6.0/24      172.18.255.255           0                      0 65535 65534 65533 ?
*>e172.30.7.0/24      172.18.255.255           0                      0 65535 65534 65533 ?
*>e172.30.8.0/24      172.18.255.255           0                      0 65535 65534 65533 ?
*>r192.168.1.0/30     0.0.0.0                  0         100      32768 ?
```

字首會通告給特定BGP鄰居。

```
switch# show ip bgp neighbors 172.18.255.255 advertised-routes

Peer 172.18.255.255 routes for address family IPv4 Unicast:
BGP table version is 88, Local Router ID is 172.17.255.255
Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best
Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-injected
Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - best2

Network            Next Hop        Metric    LocPrf    Weight Path
*>i10.100.1.0/24   10.1.1.1                  100       0 i
*>i10.100.2.0/24   10.1.1.1                  100       0 i
*>i10.100.3.0/24   10.1.1.1                  150       0 i
*>i10.100.4.0/24   10.1.1.1        0         100       0 ?
*>i10.100.5.0/24   10.1.1.1        0         100       0 ?
*>i10.100.6.0/24   10.1.1.1        0         100       0 ?
*>i10.100.7.0/24   10.1.1.1        0         100       0 ?
*>i10.100.8.0/24   10.1.1.1        0         100       0 ?
```

在任何過濾器（字首清單和/或路由對映）之前從BGP對等體接收的字首*

```
switch# show ip bgp neighbors 172.18.255.255 received-routes

Peer 172.18.255.255 routes for address family IPv4 Unicast:
```

```
BGP table version is 88, Local Router ID is 172.17.255.255
Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best
Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-injected
Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - best2

Network                Next Hop         Metric      LocPrf         Weight Path
* e172.18.255.255/32   172.18.255.255    0                          0 65535 ?
*>e172.30.1.0/24       172.18.255.255    0                          0 65535 ?
*>e172.30.2.0/24       172.18.255.255    0                          0 65535 ?
*>e172.30.3.0/24       172.18.255.255    0                          0 65535 ?
*>e172.30.4.0/24       172.18.255.255    0                          0 65535 ?
*>e172.30.5.0/24       172.18.255.255    0                          0 65535 65534 65533 ?
*>e172.30.6.0/24       172.18.255.255    0                          0 65535 65534 65533 ?
*>e172.30.7.0/24       172.18.255.255    0                          0 65535 65534 65533 ?
*>e172.30.8.0/24       172.18.255.255    0                          0 65535 65534 65533 ?
```

附註 必須在鄰居上配置軟重新配置入站

在篩選條件之後從BGP對等點接收的字首（prefixes-list和/或路由對映）

```
switch# show ip bgp neighbors 172.18.255.255 routes

Peer 172.18.255.255 routes for address family IPv4 Unicast:
BGP table version is 88, Local Router ID is 172.17.255.255
Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best
Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-injected
Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - best2

Network            Next Hop          Metric     LocPrf      Weight Path
*>e172.30.1.0/24   172.18.255.255    0                      0 65535 ?
*>e172.30.2.0/24   172.18.255.255    0                      0 65535 ?
*>e172.30.3.0/24   172.18.255.255    0                      0 65535 ?
*>e172.30.4.0/24   172.18.255.255    0                      0 65535 ?
*>e172.30.5.0/24   172.18.255.255    0                      0 65535 65534 65533 ?
*>e172.30.6.0/24   172.18.255.255    0                      0 65535 65534 65533 ?
*>e172.30.7.0/24   172.18.255.255    0                      0 65535 65534 65533 ?
*>e172.30.8.0/24   172.18.255.255    0                      0 65535 65534 65533 ?
```

特定首碼的詳細路徑資訊。

```
switch# show ip bgp 172.30.6.0
BGP routing table information for VRF default, address family IPv4 Unicast
BGP routing table entry for 172.30.6.0/24, version 28                                  -> Prefi
Paths: (3 available, best #3)                                                          -> Numbe
Flags: (0x8000001a) (high32 00000000) on xmit-list, is in urib, is best urib route, is in HW

Path type: external, path is valid, not best reason: Router Id, no labeled nexthop
AS-Path: 65535 65534 65533 , path sourced external to AS                               -> As Pa
```

```
172.20.255.255 (metric 0) from 172.20.255.255 (172.20.255.255)                                  -> Next
Origin incomplete, MED 0, localpref 100, weight 0

Path type: external, path is valid, not best reason: newer EBGP path, no labeled nexthop
AS-Path: 65535 65534 65533 , path sourced external to AS
172.19.255.255 (metric 0) from 172.19.255.255 (172.19.255.255)
Origin incomplete, MED 0, localpref 100, weight 0

Advertised path-id 1
Path type: external, path is valid, is best path, no labeled nexthop, in rib                    -> Path s
AS-Path: 65535 65534 65533 , path sourced external to AS
172.18.255.255 (metric 0) from 172.18.255.255 (172.18.255.255)
Origin incomplete, MED 0, localpref 100, weight 0

Path-id 1 advertised to peers:                                                                  -> BGP pee
10.1.1.2
```

註: *AS-Path表示到達字首源位置所遍歷的ASN。
**AS-Path從右到左讀取。

若要檢視BGP中的最佳路徑選取流程，請參閱BGP最佳路徑選取。

相關資訊

- 思科技術支援與下載

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）。