Cisco IOS-XE SD-WAN安装带DN位的OSPF外部路由

目录

简介

本文描述当开放最短路径优先(OSPF)外部路由安装到路由表中时Cisco IOS®-XE SD-WAN软件的预期行为。

Cisco IOS-XE SD-WAN安装带DN位的OSPF外部路由

运行Cisco IOS-XE SD-WAN软件的路由器将OSPF外部路由（E1或E2）安装到路由表中。为了进行演示，请考虑以下简单拓扑图：

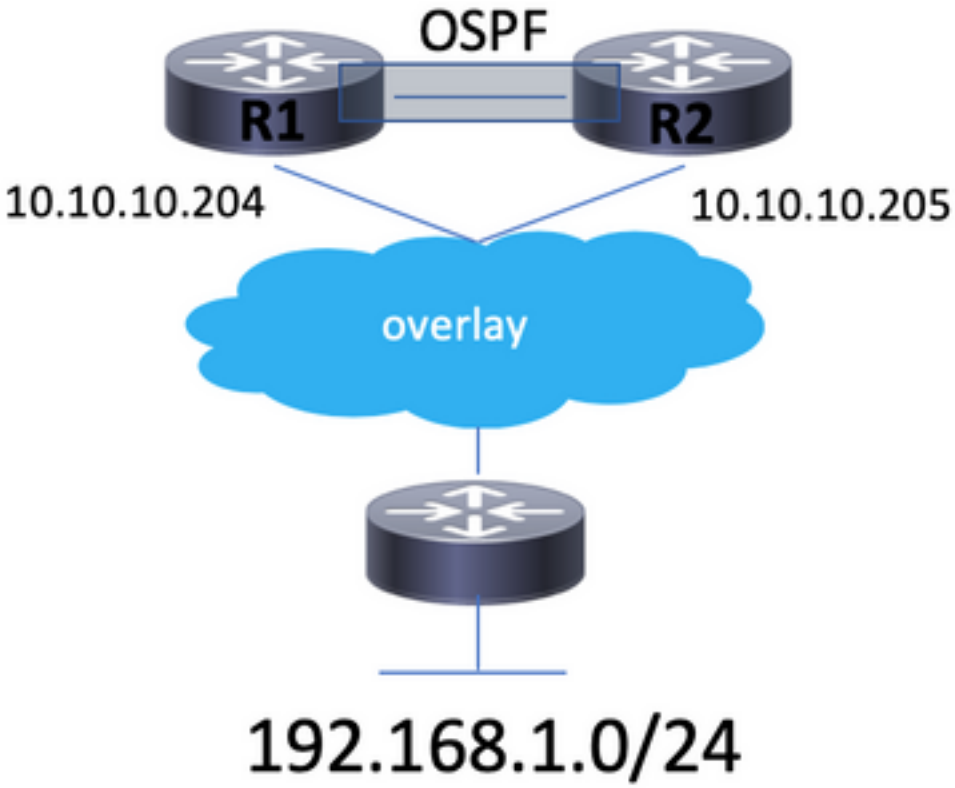

以下是运行Cisco IOS-XE SD-WAN软件的一对路由器R1和R2，它们通过服务端vpn（本例中为vrf 2）建立OSPF对等。 路由器相应地具有system-ip 10.10.10.204和10.10.10.205。System-ip等于OSPF路由器ID。其他一些路由器通过重叠管理协议(OMP)向此站点通告前缀192.168.1.0/24。

两台路由器的配置方式类似。此处提供了相关配置（要点是OSPF和OMP之间的相互重分发已完成）：

```
route-map omp2ospf permit 10
 set metric 1000
 set metric-type type-1
!
router ospf 2 vrf 2
 compatible rfc1583
 distance ospf external 110
 distance ospf inter-area 110
 distance ospf intra-area 110
 redistribute omp route-map omp2ospf
!
 omp
  no shutdown
  send-path-limit  4
  ecmp-limit       4
  graceful-restart
  no as-dot-notation
  timers
   holdtime               60
   advertisement-interval 1
   graceful-restart-timer 43200
   eor-timer              300
  exit
  address-family ipv4 vrf 2
   advertise ospf external
   advertise connected
   advertise static
  !
  address-family ipv4
   advertise connected
   advertise static
  !
  address-family ipv6
   advertise connected
   advertise static
  !
```

完成正常情况路由表条目后，192.168.1.0/24会从OMP安装到路由信息库(RIB)中，并重新分发到OSPF。此条目如下所示：

```
R1#sh ip route vrf 2 192.168.1.0 255.255.255.0

Routing Table: 2
Routing entry for 192.168.1.0/24
  Known via "omp", distance 251, metric 0, type omp
  Redistributing via ospf 2
  Advertised by ospf 2 subnets route-map omp2ospf
  Last update from 10.10.10.201 00:03:00 ago
  Routing Descriptor Blocks:
  * 10.10.10.201 (default), from 10.10.10.201, 00:03:00 ago
      Route metric is 0, traffic share count is 1

R1#show ip ospf database external 192.168.1.0

            OSPF Router with ID (172.16.1.204) (Process ID 2)

Type-5 AS External Link States
```

```
  LS age: 354
  Options: (No TOS-capability, DC, Downward)
  LS Type: AS External Link
  Link State ID: 192.168.1.0 (External Network Number )
  Advertising Router: 172.16.1.204
  LS Seq Number: 80000001
  Checksum: 0x25AE
  Length: 36
  Network Mask: /24
Metric Type: 1 (Comparable directly to link state metric)
MTID: 0
Metric: 1000
Forward Address: 0.0.0.0
External Route Tag: 0

  LS age: 355
  Options: (No TOS-capability, DC, Downward)
  LS Type: AS External Link
  Link State ID: 192.168.1.0 (External Network Number )
  Advertising Router: 172.16.1.205
  LS Seq Number: 80000001
  Checksum: 0x1FB3
  Length: 36
  Network Mask: /24
Metric Type: 1 (Comparable directly to link state metric)
MTID: 0
Metric: 1000
Forward Address: 0.0.0.0
External Route Tag: 0

R2#sh ip route vrf 2 192.168.1.0 255.255.255.0

Routing Table: 2
Routing entry for 192.168.1.0/24
  Known via "omp", distance 251, metric 0, type omp
  Redistributing via ospf 2
  Advertised by ospf 2 subnets route-map omp2ospf
  Last update from 10.10.10.201 00:04:13 ago
  Routing Descriptor Blocks:
  * 10.10.10.201 (default), from 10.10.10.201, 00:04:13 ago
      Route metric is 0, traffic share count is 1

R2#show ip ospf database external 192.168.1.0

            OSPF Router with ID (172.16.1.205) (Process ID 2)

Type-5 AS External Link States

  LS age: 317
  Options: (No TOS-capability, DC, Downward)
  LS Type: AS External Link
  Link State ID: 192.168.1.0 (External Network Number )
  Advertising Router: 172.16.1.204
  LS Seq Number: 80000001
  Checksum: 0x25AE
  Length: 36
  Network Mask: /24
Metric Type: 1 (Comparable directly to link state metric)
MTID: 0
Metric: 1000
Forward Address: 0.0.0.0
External Route Tag: 0
```

```
  LS age: 316
  Options: (No TOS-capability, DC, Downward)
  LS Type: AS External Link
  Link State ID: 192.168.1.0 (External Network Number )
  Advertising Router: 172.16.1.205
  LS Seq Number: 80000001
  Checksum: 0x1FB3
  Length: 36
  Network Mask: /24
Metric Type: 1 (Comparable directly to link state metric)
MTID: 0
Metric: 1000
Forward Address: 0.0.0.0
External Route Tag: 0
```

如您所见，两台路由器都将路由安装到RIB中，然后将其重新分发到OSPF。两台路由器都将DN位设置为外部LSA第5类，这应该会防止这些路由作为OSPF路由安装到RIB中，从而重新分发回OMP，从而从根本上防止环路。这与RFC 4576和RFC 4577中描述的机制相同。

所有路由器都与vSmart控制器建立了OMP对等：

```
R1#show sdwan omp peers
R -> routes received
I -> routes installed
S -> routes sent

                        DOMAIN    OVERLAY   SITE
PEER            TYPE    ID        ID        ID        STATE    UPTIME           R/I/S
-----------------------------------------------------------------------------------
10.10.10.229    vsmart  1         1         1         up       1:19:35:34       30/12/5
10.10.10.230    vsmart  1         1         3         up       1:19:35:33       26/1/5

R2#show sdwan omp peers
R -> routes received
I -> routes installed
S -> routes sent

                        DOMAIN    OVERLAY   SITE
PEER            TYPE    ID        ID        ID        STATE    UPTIME           R/I/S
-----------------------------------------------------------------------------------
10.10.10.229    vsmart  1         1         1         up       0:01:38:48       30/10/6
10.10.10.230    vsmart  1         1         3         up       1:19:35:36       25/1/6
```

现在，R1与两个OMP对等体失去连接：

```
Oct 11 12:53:57.777: %Cisco-SDWAN-Router-OMPD-3-ERRO-400002: R0/0: OMPD: vSmart peer
10.10.10.229 state changed to Init
Oct 11 12:53:57.777: %Cisco-SDWAN-Router-OMPD-6-INFO-400005: R0/0: OMPD: Number of vSmarts
connected : 1
Oct 11 12:53:58.777: %Cisco-SDWAN-Router-OMPD-3-ERRO-400002: R0/0: OMPD: vSmart peer
10.10.10.230 state changed to Init
Oct 11 12:53:58.777: %Cisco-SDWAN-Router-OMPD-6-INFO-400005: R0/0: OMPD: Number of vSmarts
connected : 0

R1#show sdwan omp peers
R -> routes received
I -> routes installed
S -> routes sent

                        DOMAIN    OVERLAY   SITE
PEER            TYPE    ID        ID        ID        STATE    UPTIME           R/I/S
```

```
-----------------------------------------------------------------------------------------
10.10.10.229     vsmart  1         1         1         init-in-gr                30/12/0
10.10.10.230     vsmart  1         1         3         init-in-gr                26/1/0
```

R1会将OMP路由标记为过时（参见OMP路由状态S），但会继续将该路由保留在OMP协议安装的RIB中，直到graceful-restart-timer过期：

```
R1#show sdwan omp routes 192.168.1.0/24 | exclude not set
-----------------------------------------------
omp route entries for vpn 2 route 192.168.1.0/24
-----------------------------------------------
            RECEIVED FROM:
peer            10.10.10.229
path-id         1076
label           1002
status          C,I,R,S
    Attributes:
     originator       10.10.10.201
     type             installed
     tloc             10.10.10.201, biz-internet, ipsec
     overlay-id        1
     site-id          201207
     origin-proto     connected
     origin-metric    0
            RECEIVED FROM:
peer            10.10.10.230
path-id         775
label           1002
status          C,R,S
    Attributes:
     originator       10.10.10.201
     type             installed
     tloc             10.10.10.201, biz-internet, ipsec
     overlay-id        1
     site-id          201207
     origin-proto     connected
     origin-metric    0

R1#sh ip route vrf 2 192.168.1.0 255.255.255.0

Routing Table: 2
Routing entry for 192.168.1.0/24
  Known via "omp", distance 251, metric 0, type omp
  Redistributing via ospf 2
  Advertised by ospf 2 subnets route-map omp2ospf
  Last update from 10.10.10.201 00:23:35 ago
  Routing Descriptor Blocks:
  * 10.10.10.201 (default), from 10.10.10.201, 00:23:35 ago
      Route metric is 0, traffic share count is 1
```

默认graceful-restart-timer计时器为43,200秒（12小时）。 一旦到期，到192.168.1.0/24的路由将仍然存在。

```
R1#sh ip route vrf 2 192.168.1.0 255.255.255.0

Routing Table: 2
Routing entry for 192.168.1.0/24
  Known via "ospf 2", distance 252, metric 1100, type extern 1
  Redistributing via omp
  Last update from 10.28.7.205 on Vlan2807, 00:04:11 ago
  Routing Descriptor Blocks:
```

```
  * 10.28.7.205, from 172.16.1.205, 00:04:11 ago, via Vlan2807
      SDWAN Down
      Route metric is 1100, traffic share count is 1

R1#show ip ospf database external 192.168.1.0

            OSPF Router with ID (172.16.1.204) (Process ID 2)

Type-5 AS External Link States

  LS age: 339
  Options: (No TOS-capability, DC, Downward)
  LS Type: AS External Link
  Link State ID: 192.168.1.0 (External Network Number )
  Advertising Router: 172.16.1.205
  LS Seq Number: 80000004
  Checksum: 0x19B6
  Length: 36
  Network Mask: /24
Metric Type: 1 (Comparable directly to link state metric)
MTID: 0
Metric: 1000
Forward Address: 0.0.0.0
External Route Tag: 0
```

它现在作为OSPF外部第1类路由安装，尽管对应的OSPF LSA设置了DN位。

另请注意，管理距离(AD)始终比OMP的AD多1个单位（251是OMP的默认值，因此本例中为252）。

必须解释路由器为何使用大于OMP路由AD的AD来安装此路由。这是因为，当OMP对等重新建立且到交换矩阵的可达性恢复时，您会尝试防止环路方案。

如果启用了debug ip routing和debug ip ospf rib redistribution命令，则还可以清楚地看**到AD=252的路由安装过程**:

```
Oct 11 14:13:28.302: RT(2): del 192.168.1.0 via 10.10.10.201, omp metric [251/0]
Oct 11 14:13:28.303: RT(2): delete network route to 192.168.1.0/24
Oct 11 14:13:28.307: OSPF-2 REDIS: Notification to redistribute 192.168.1.0/24
Oct 11 14:13:28.307: RT(2): updating ospf 192.168.1.0/24 (0x2) [local lbl/ctx:1048577/0x0] omp-
tag:0  :
    via 10.28.7.205 Vl2807  0 1048578 0x100001
Oct 11 14:13:28.307: RT(2): add 192.168.1.0/24 via 10.28.7.205, ospf metric [252/1100]
```

这是Cisco IOS-XE SD-WAN软件中特别引入的预期行为，以避免当其中一台路由器从SD-WAN重叠分区时出现流量黑洞场景。出现黑洞可能是因为服务端流量仍通过两台路由器进行负载均衡。这是因为两条静态路由指向两台路由器，或某些路由仅指向一台已分区的路由器。

在ECMP（当R1从交换矩阵分区时）流量遵循两条路径：

LAN -> R1 -> R2 ->远程路由器 — > 192.168.1.0/24

LAN -> R2 ->远程路由器 — > 192.168.1.0/24

在此，您还可以看到当R1从交换矩阵分区时R1的输出示例。如您所见，通过R2（10.27.7.205下一跳）仍可保留到LAN子网192.168.1.0/24的连接：

```
R1#ping vrf 2 192.168.1.1
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/33/44 ms
R1# traceroute vrf 2 192.168.1.1 numeric
Type escape sequence to abort.
Tracing the route to 192.168.1.1
VRF info: (vrf in name/id, vrf out name/id)
  1 10.28.7.205 4 msec 0 msec 0 msec
  2 192.168.1.1 4 msec *  0 msec
```