在 OSPF 进程之间再分配时的次理想路由

目录

简介

本文档演示了在开放最短路径优先(OSPF)进程之间减少路由时的次优路由问题，并提供了解决方案。

先决条件

要求

本文档没有任何特定的要求。

使用的组件

本文档不限于特定的软件和硬件版本。

本文档中的信息都是基于特定实验室环境中的设备编写的。本文档中使用的所有设备最初均采用原始（默认）配置。如果您使用的是真实网络，请确保您已经了解所有命令的潜在影响。

规则

有关文件规则的更多信息请参见“ Cisco技术提示规则”。

问题

当在网络上的多个点中在不同OSPF进程之间进行路由时，可能会陷入子短路路由甚至更糟的路由环路。

在下面的拓扑中，我们有OSPF 1和OSPF 2进程。路由器1(R1)和路由器2(R2)正在从OSPF 1重分发到OSPF 2。

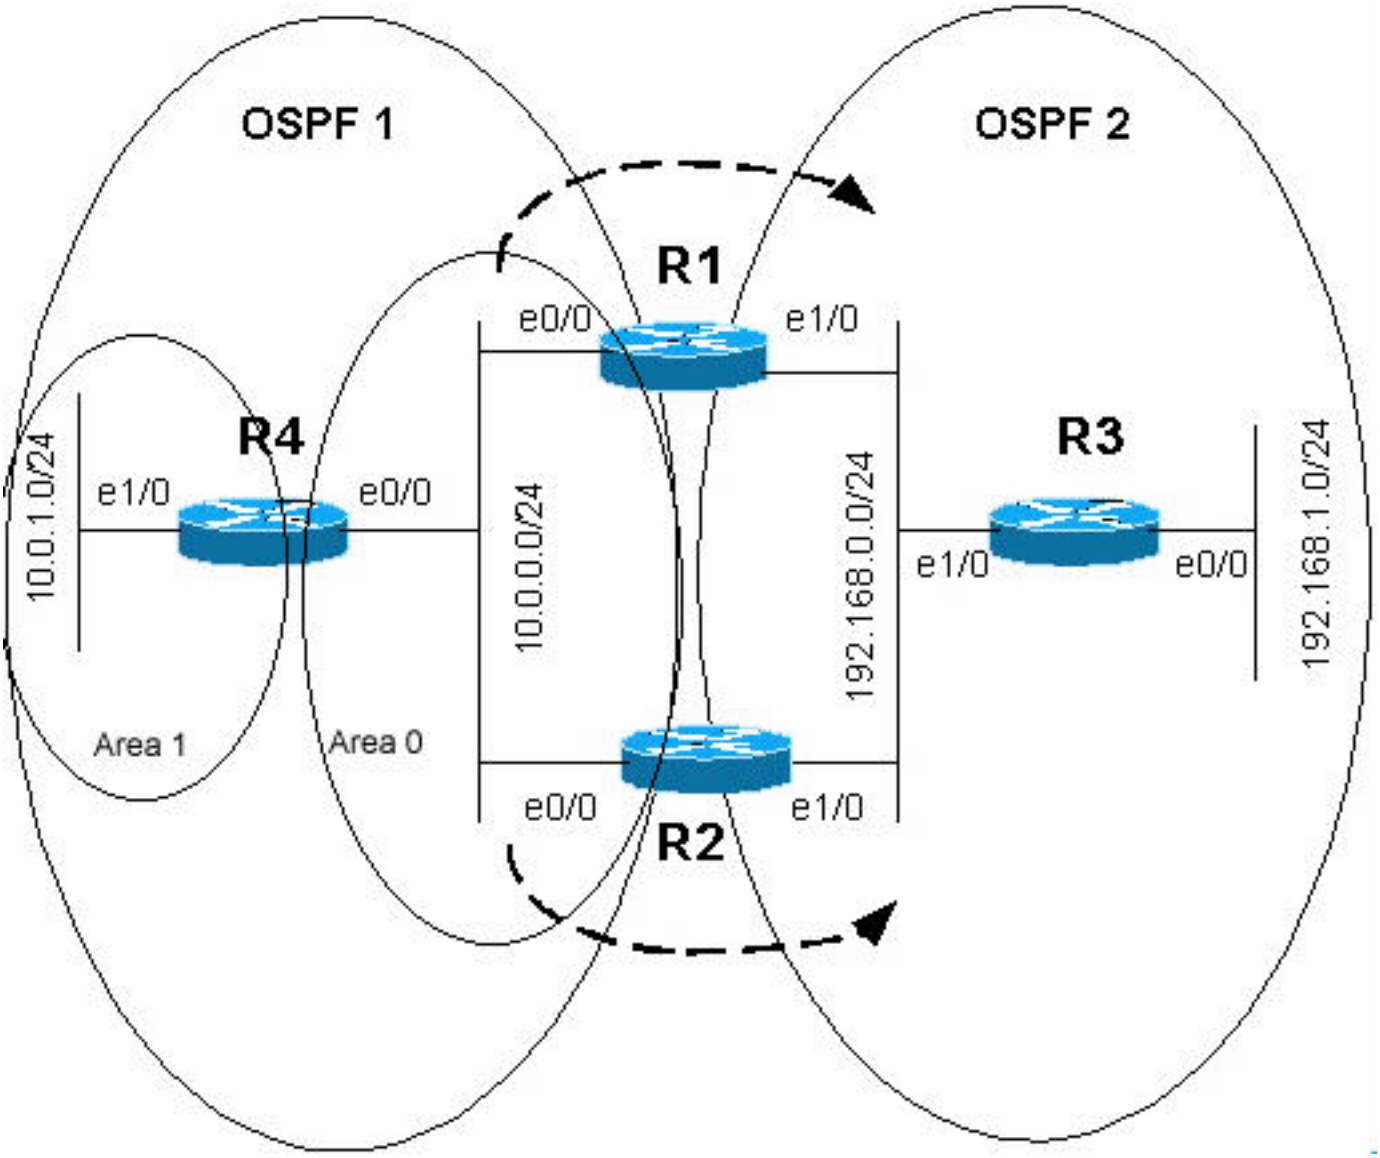

路由器R1和R2的配置如下所示。

R1

```
hostname r1
!
ip subnet-zero
!
interface Loopback0
 ip address 10.255.255.1 255.255.255.255
!
interface Loopback1
 ip address 192.168.255.1 255.255.255.255
!
interface Ethernet0/0
 ip address 10.0.0.1 255.255.255.0
!
interface Ethernet1/0
 ip address 192.168.0.1 255.255.255.0
!
router ospf 1
 router-id 10.255.255.1
```

```
 log-adjacency-changes
 passive-interface Loopback0
network 10.0.0.0 0.0.0.255 area 0
 network 10.255.255.1 0.0.0.0 area 0
!
router ospf 2
 router-id 192.168.255.1
 log-adjacency-changes
 redistribute ospf 1 subnets match internal
 !--- Redistributing OSPF 1 into OSPF 2. passive-
interface Loopback1 network 192.168.0.0 0.0.0.255 area 0
network 192.168.255.1 0.0.0.0 area 0 ! ip classless !
end
```

R2

```
hostname r2
!
ip subnet-zero
!

interface Loopback0
 ip address 10.255.255.2 255.255.255.255
!
interface Loopback1
 ip address 192.168.255.2 255.255.255.255
!
interface Ethernet0/0
 ip address 10.0.0.2 255.255.255.0
!
interface Ethernet1/0
 ip address 192.168.0.2 255.255.255.0
!
router ospf 1
 router-id 10.255.255.2
 log-adjacency-changes
 passive-interface Loopback0
network 10.0.0.0 0.0.0.255 area 0
 network 10.255.255.2 0.0.0.0 area 0
!
router ospf 2
 router-id 192.168.255.2
 log-adjacency-changes
 redistribute ospf 1 subnets match internal
!--- Redistributing OSPF 1 into OSPF 2. passive-
interface Loopback1 network 192.168.0.0 0.0.0.255 area 0
network 192.168.255.2 0.0.0.0 area 0 ! ip classless end
```

在上述拓扑中,R4的E1/0在区域1中，E0/0在区域0中。因此，R4是区域边界路由器(ABR)，通告网络10.0.1.0/24作为到R1和R2的区域间(IA)路由。R1和R2将此信息重分发到OSPF 2。**redistribute** 配置命令在R1 和**R2** 的上述配置中突出显示。因此，R1和R2将获知10.0.1.0/24作为IA（通过OSPF 1），以及作为外部类型2(E2)（通过OSPF 2），因为外部链路状态通告(LSA)在整个传播OSPF 2域。

由于IA路由始终优先于E1或E2路由，因此期望在R1和R2的路由表中看到10.0.1.0/24是具有下一跳R4的IA路由。但是，在查看其路由表时，会看到一些不同之处 — 在R1上，10.0.1.0/24是具有下一跳R4但在R4上的IA路由10.0.1.0/24是具有下一跳R1的E2路由。

这是R1的**show ip route**命令的命令输出。

```
r1#show ip route

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

!--- The gateway of the last resort is not set. 10.0.0.0/8 is variably subnetted, 5 subnets, 2
masks O E2 10.255.255.2/32 [110/1] via 192.168.0.2, 00:24:21, Ethernet1/0 C 10.0.0.0/24 is
directly connected, Ethernet0/0 O IA    10.0.1.0/24 [110/20] via 10.0.0.4, 00:23:49, Ethernet0/0
C       10.255.255.1/32 is directly connected, Loopback0
O IA    10.255.255.4/32 [110/11] via 10.0.0.4, 00:23:49, Ethernet0/0
     192.168.255.0/32 is subnetted, 3 subnets
O       192.168.255.3 [110/11] via 192.168.0.3, 00:26:09, Ethernet1/0
O       192.168.255.2 [110/11] via 192.168.0.2, 00:26:09, Ethernet1/0
C       192.168.255.1 is directly connected, Loopback1
C    192.168.0.0/24 is directly connected, Ethernet1/0
O    192.168.1.0/24 [110/20] via 192.168.0.3, 00:26:09, Ethernet1/0
```

这是R2的show ip route命令的命令输出。

```
r2#show ip route

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

!--- The gateway of last resort is not set. 10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
C 10.255.255.2/32 is directly connected, Loopback0 C 10.0.0.0/24 is directly connected,
Ethernet0/0 O E2    10.0.1.0/24 [110/20] via 192.168.0.1, 00:25:34, Ethernet1/0
O E2    10.255.255.1/32 [110/1] via 192.168.0.1, 00:25:34, Ethernet1/0
O E2    10.255.255.4/32 [110/11] via 192.168.0.1, 00:25:34, Ethernet1/0
     192.168.255.0/32 is subnetted, 3 subnets
O       192.168.255.3 [110/11] via 192.168.0.3, 00:26:45, Ethernet1/0
C       192.168.255.2 is directly connected, Loopback1
O       192.168.255.1 [110/11] via 192.168.0.1, 00:26:45, Ethernet1/0
C    192.168.0.0/24 is directly connected, Ethernet1/0
O    192.168.1.0/24 [110/20] via 192.168.0.3, 00:26:45, Ethernet1/0
```

为什么发生此问题？

在路由器上启用多个OSPF进程时，从软件角度来看，这些进程是独立的。在一个OSPF进程内，OSPF协议始终优先使用内部路由而非外部路由。但是，OSPF不会在进程之间选择任何OSPF路由（例如，在决定应将哪个进程的路由添加到路由表中时，OSPF度量和路由类型不会被考虑在内）。

不同OSPF进程之间不存在交互，而管理距离是决定因素。因此，由于两个OSPF进程的默认管理距离为110，因此尝试安装该路由的第一个进程会将其添加到路由表中。因此，必须配置来自不同OSPF进程的路由的管理距离，以便某些OSPF进程的路由优先于另一进程的路由，而不是作为偶然。

有关管理距离的详细信息，请参阅什么是管理距离。有关Cisco路由器如何选择要放入路由表中的路由的详细信息，请参阅Cisco路由器中的路由选择。

解决方案

解决方案 1

由于我们知道在上述情况下，路由器根据管理距离选择最佳路由，因此防止此行为的逻辑方法是增加OSPF 2中外部路由的管理距离。这样，通过OSPF 1获知的路由始终优先于从OSPF 1重新分发到OSPF 2的外部路由。这是使用子路由器配置命令**distance ospf external *<value>*** 如以下配置所示。

R1

```
hostname r1
!
ip subnet-zero
!
interface Loopback0
 ip address 10.255.255.1 255.255.255.255
!
interface Loopback1
 ip address 192.168.255.1 255.255.255.255
!
interface Ethernet0/0
 ip address 10.0.0.1 255.255.255.0
!
interface Ethernet1/0
 ip address 192.168.0.1 255.255.255.0
!
router ospf 1
 router-id 10.255.255.1
 log-adjacency-changes
 passive-interface Loopback0
 network 10.0.0.0 0.0.0.255 area 0
 network 10.255.255.1 0.0.0.0 area 0
!
router ospf 2
 router-id 192.168.255.1
 log-adjacency-changes
 redistribute ospf 1 subnets match internal
 passive-interface Loopback1
 network 192.168.0.0 0.0.0.255 area 0
 network 192.168.255.1 0.0.0.0 area 0
 distance ospf external 115
!--- Increases the administrative distance of external
!--- routes to 115. ! ip classless ! end
```

R2

```
hostname r2
!
ip subnet-zero
!
interface Loopback0
 ip address 10.255.255.2 255.255.255.255
!
interface Loopback1
 ip address 192.168.255.2 255.255.255.255
!
interface Ethernet0/0
 ip address 10.0.0.2 255.255.255.0
```

```
!
interface Ethernet1/0
 ip address 192.168.0.2 255.255.255.0
!
router ospf 1
 router-id 10.255.255.2
 log-adjacency-changes
 passive-interface Loopback0
 network 10.0.0.0 0.0.0.255 area 0
 network 10.255.255.2 0.0.0.0 area 0
!
router ospf 2
 router-id 192.168.255.2
 log-adjacency-changes
 redistribute ospf 1 subnets match internal
 passive-interface Loopback1
 network 192.168.0.0 0.0.0.255 area 0
 network 192.168.255.2 0.0.0.0 area 0
 distance ospf external 115
!--- Increases the administrative distance of !---
external routes to 115. ! ip classless ! end
```

更改OSPF 2中外部路由的管理距离时产生的路由表如下所示。

这是R1的**show ip route**命令的命令输出。

```
r1#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

!--- The gateway of the last resort is not set. 10.0.0.0/8 is variably subnetted, 5 subnets, 2
masks O 10.255.255.2/32 [110/11] via 10.0.0.2, 00:00:35, Ethernet0/0 C 10.0.0.0/24 is directly
connected, Ethernet0/0 O IA    10.0.1.0/24 [110/20] via 10.0.0.4, 00:00:35, Ethernet0/0
C       10.255.255.1/32 is directly connected, Loopback0
O       10.255.255.4/32 [110/11] via 10.0.0.4, 00:00:35, Ethernet0/0
     192.168.255.0/32 is subnetted, 3 subnets
O       192.168.255.3 [110/11] via 192.168.0.3, 00:00:35, Ethernet1/0
O       192.168.255.2 [110/11] via 192.168.0.2, 00:00:35, Ethernet1/0
C       192.168.255.1 is directly connected, Loopback1
C    192.168.0.0/24 is directly connected, Ethernet1/0
O    192.168.1.0/24 [110/20] via 192.168.0.3, 00:00:35, Ethernet1/0
```

这是R2的**show ip route**命令的命令输出。

```
r2#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

!--- The gateway of the last resort is not set. 10.0.0.0/8 is variably subnetted, 5 subnets, 2
masks C 10.255.255.2/32 is directly connected, Loopback0 C 10.0.0.0/24 is directly connected,
Ethernet0/0 O 10.255.255.1/32 [110/11] via 10.0.0.1, 00:01:28, Ethernet0/0 O IA    10.0.1.0/24
```

```
[110/20] via 10.0.0.4, 00:01:28, Ethernet0/0
O       10.255.255.4/32 [110/11] via 10.0.0.4, 00:01:28, Ethernet0/0
     192.168.255.0/32 is subnetted, 3 subnets
O       192.168.255.3 [110/11] via 192.168.0.3, 00:01:28, Ethernet1/0
C       192.168.255.2 is directly connected, Loopback1
O       192.168.255.1 [110/11] via 192.168.0.1, 00:01:28, Ethernet1/0
C    192.168.0.0/24 is directly connected, Ethernet1/0
O    192.168.1.0/24 [110/20] via 192.168.0.3, 00:01:28, Ethernet1/0
```

请注意，在某些情况下，如果还有从OSPF 2重分发到OSPF 1的路由协议，并且有其他路由协议重分发到OSPF 2(路由信息协议[RIP]、增强型内部网关路由协议(EIGRP)静态等)，这可能导致OSPF 2中对这些外部路由进行次优路由。

解决方案 2

如果实施两个不同OSPF进程的最终原因是过滤某些路由，则Cisco IOS®软件版本12.2(4)T中有一项新功能，称为OSPF ABR第3类LSA过滤，它允许您在ABR中执行路由过滤。

在上例中，作为OSPF 2一部分的链路可以配置为OSPF 1内的另一个区域，而不是配置第二个OSPF进程。然后，您可以使用此新功能在R1和R2中实施所需的路由过滤。有关此功能的详细信息，请参阅OSPF ABR第3类LSA过滤。

相关信息

- OSPF 支持页
- IP 路由协议支持页
- IP 路由 支持页
- 技术支持 - Cisco Systems