

Suboptimal Routing When Redistributing Between OSPF Processes

Document ID: 28025

Contents

Introduction

Prerequisites

- Requirements

- Components Used

- Conventions

Problem

Why Does This Problem Occur?

Solutions

- Solution 1

- Solution 2

Related Information

Introduction

This document demonstrates the suboptimal routing problem when redistributing between Open Shortest Path First (OSPF) processes and offers solutions.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software and hardware versions.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.


Conventions

For more information on document conventions, refer to the Cisco Technical Tips Conventions.

Problem

When redistributing between different OSPF processes in multiple points on the network, it is possible to get into situations of suboptimal routing or even worse, a routing loop.

In the topology below we have the processes OSPF 1 and OSPF 2. Router 1 (R1) and router 2 (R2) are redistributing from OSPF 1 into OSPF 2.



The configurations for routers R1 and R2 are shown below.

| R1 |
|---|
| <pre> 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 </pre> |

```

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

```

In the above topology, R4's E1/0 is in Area 1 and E0/0 is in Area 0. Therefore, R4 is an Area Border Router (ABR) advertising the network 10.0.1.0/24 as the inter-area (IA) route to R1 and R2. R1 and R2 are redistributing this information into OSPF 2. The **redistribute** configuration commands are highlighted in the above configurations of R1 and R2. Therefore, both R1 and R2 are going to learn about 10.0.1.0/24 as IA via OSPF 1 and as external type 2 (E2) via OSPF 2 because the external link-state advertisements (LSAs) are propagated throughout the OSPF 2 domain.

Since IA routes are always preferred over E1 or E2 routes, the expectation is to see, in the routing table of R1 and R2, that 10.0.1.0/24 is an IA route with next-hop R4. However, when viewing their routing tables, something different is seen – on R1, 10.0.1.0/24 is an IA route with next-hop R4 but on R2, 10.0.1.0/24 is an E2 route with next-hop R1.

This is the command output of the **show ip route** command for R1.

```
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

```

This is the command output of the **show ip route** command for R2.

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

```

Why Does This Problem Occur?

When enabling multiple OSPF processes on a router, from the software perspective, the processes are independent. OSPF protocol, inside one OSPF process, always prefers the Internal route over the External route. However, OSPF does not do any OSPF route selection between processes (for instance, OSPF metrics and route types are not taken into account, when deciding the route of which process should be installed into the routing table).

There is no interaction between different OSPF processes, and the tie–breaker is the administrative distance. Thus, since both OSPF processes have a default administrative distance of 110, the first process trying to install that route makes it into the routing table. Therefore, administrative distance for routes from different OSPF processes must be configured, so that routes of certain OSPF processes are preferred over routes of another process by human intention, and not as a matter of chance.

For more information on administrative distance, refer to [What is Administrative Distance](#). For more information on how a Cisco router selects which routes to place in the routing table, refer to [Route Selection in Cisco Routers](#).

Solutions

Solution 1

Since we know that in the above case, the routers are selecting the best route based on administrative distance, the logical way to prevent this behavior is to increase the administrative distance of the external routes in OSPF 2. This way, routes learned via OSPF 1 will always be preferred over external routes redistributed from OSPF 1 into OSPF 2. This is done using the sub–router configuration command **distance ospf external <value>** as shown in the configurations below.

```

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
```

The resulting routing table when changing the administrative distance of the external routes in OSPF 2 is shown below.

This is the command output of the **show ip route** command for R1.

```
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

```

This is the command output of the **show ip route** command for R2.

```

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

```

It is important to note that in some cases, when there is also redistribution from OSPF 2 into OSPF 1 and there are other routing protocols being redistributed into OSPF 2 (Routing Information Protocol [RIP], Enhanced Interior Gateway Routing Protocol (EIGRP) statics, and so forth), this can lead to suboptimal routing in OSPF 2 for those external routes.

Solution 2

If the ultimate reason to implement two different OSPF processes is to filter certain routes, there is a new feature in Cisco IOS® Software Release 12.2(4)T called OSPF ABR Type 3 LSA Filtering which allows you to do route filtering in the ABR.

Instead of configuring a second OSPF process, the links that are part of OSPF 2, in the example above, could be configured as another area inside OSPF 1. Then, you can implement the required route filtering in R1 and R2 with this new feature. For more information about this feature, refer to OSPF ABR Type 3 LSA Filtering.

Related Information

- [OSPF Support Page](#)
 - [IP Routed Protocols Support Page](#)
 - [IP Routing Support Page](#)
 - [Technical Support – Cisco Systems](#)
-

