OSPF外部路径选择：外部2类(E2)VS NSSA 2类(N2)

目录

简介

本文档旨在演示当路由器收到给定外部网络的5类链路状态通告(LSA)和7类LSA时的开放最短路径优先(OSPF)路径选择行为。 在非NSSA区域中执行重分发时，OSPF会将第5类LSA注入OSPF域。重分发到NSSA区域会创建一种特殊类型的LSA，称为Type-7，它只能存在于NSSA区域中。

先决条件

使用本文档时，请参阅图1中的网络图：

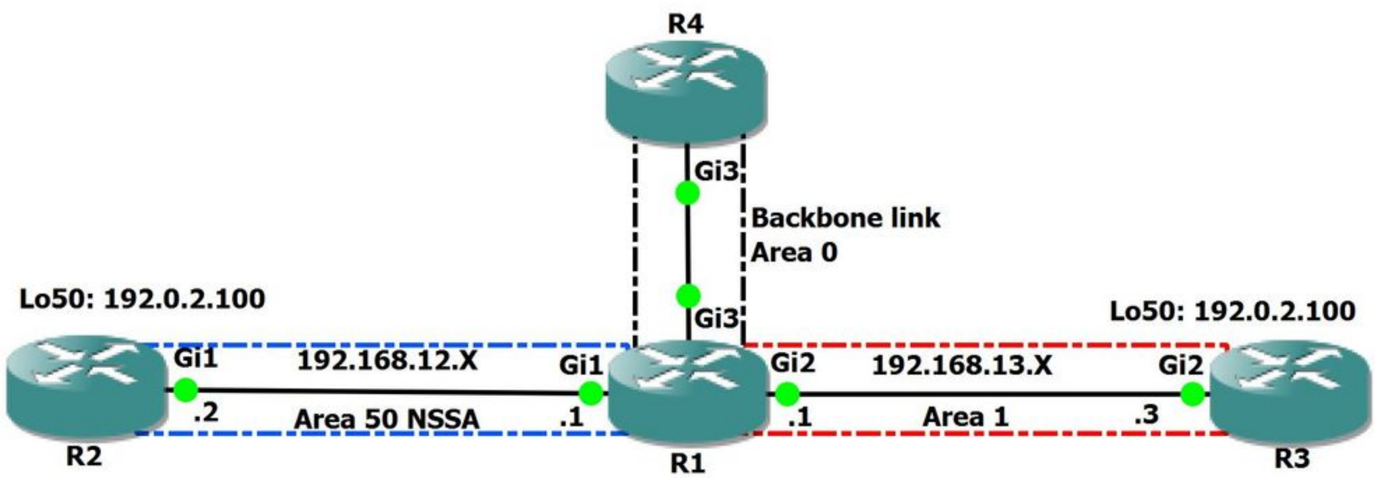

图 1

在网络图中，非主干区域1和NSSA区域50都连接到R1。R1是连接到主干区域0的区域边界路由器(ABR)。R2和R3负责将相同的前缀192.0.2.100/32重分布到OSPF域。

要求

Cisco建议您了解OSPF协议。

使用的组件

本文档中的信息基于以下软件版本：

- 思科CSR1000V版本16.4.1

背景信息

Cisco IOS-XE设备支持RFC 3101进行外部路径计算。RFC 1587已被RFC 3101淘汰，但RFC 1587特定行为仍可通过配置启用。在Cisco IOS版本15.1(2)S及更高版本中，show ip ospf命令的输出显示设备是使用RFC 3101还是RFC 1587。

摘自RFC 3101第2.5节

```
(e) If the current LSA is functionally the same as an
 installed LSA (i.e., same destination, cost and non-zero
 forwarding address) then apply the following priorities in
 deciding which LSA is preferred:

1. A Type-7 LSA with the P-bit set.
2. A Type-5 LSA.
3. The LSA with the higher router ID.
```

摘自RFC 1587第3.5节

```
5. Otherwise, compare the cost of this new AS external path
         to the ones present in the table. Note that type-5 and
         type-7 routes are directly comparable. Type-1 external
         paths are always shorter than Type-2 external paths.
         Type-1 external paths are compared by looking at the sum
         of the distance to the forwarding address/ASBR and the
         advertised Type-1 paths (X+Y). Type-2 external paths are
         compared by looking at the advertised Type-2 metrics,
         and then if necessary, the distance to the forwarding
         address/ASBR.
         When a type-5 LSA and a type-7 LSA are found to have the
         same type and an equal distance, the following priorities
         apply (listed from highest to lowest) for breaking the tie.

                 a. Any type 5 LSA.
                 b. A type-7 LSA with the P-bit set and the forwarding
                    address non-zero.
                 c. Any other type-7 LSA.

         If the new path is shorter, it replaces the present paths
         in the routing table entry. If the new path is the same
         cost, it is added to the routing table entry's list of
         paths
```

场景 1

网络图

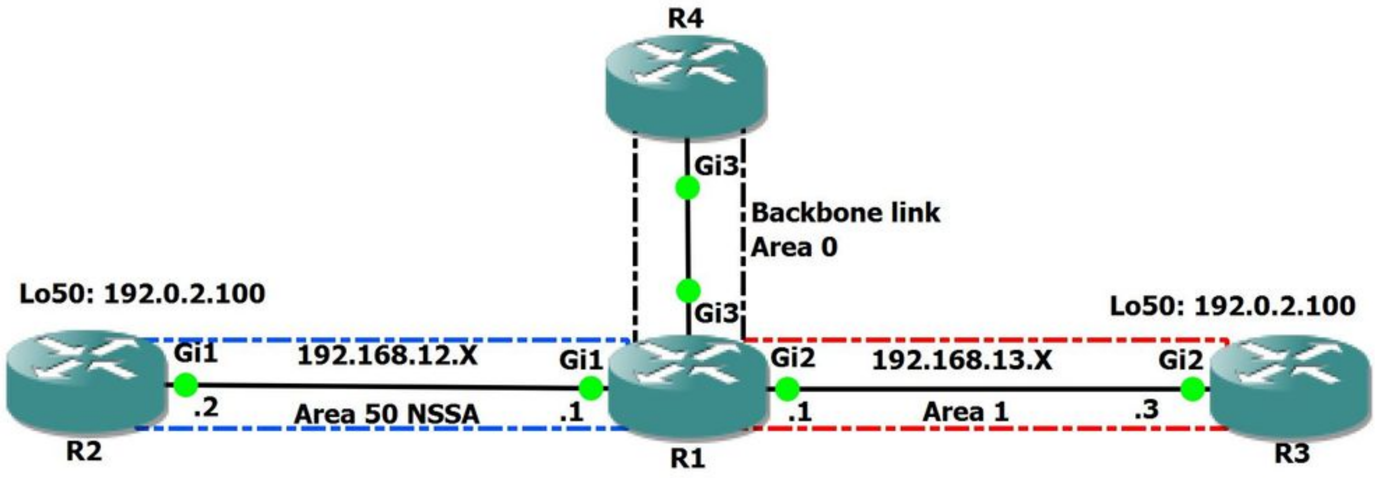

图 2

在此场景中，我们将观察使用RFC 3101进行外部路径计算时观察到的行为。我们对在R3和R2上重分发的前缀192.0.2.100/32感兴趣。

R1的第1类LSA在以下输出中：

```
R1#show ip ospf database router 1.1.1.1

            OSPF Router with ID (1.1.1.1) (Process ID 1)

                Router Link States (Area 0)

  LS age: 51
  Options: (No TOS-capability, DC)
  LS Type: Router Links
  Link State ID: 1.1.1.1
  Advertising Router: 1.1.1.1
  LS Seq Number: 80000007
  Checksum: 0x3BD6
  Length: 48
  Area Border Router
  AS Boundary Router
  Number of Links: 2

    Link connected to: another Router (point-to-point)
     (Link ID) Neighboring Router ID: 4.4.4.4
     (Link Data) Router Interface address: 192.168.14.1
      Number of MTID metrics: 0
       TOS 0 Metrics: 1

    Link connected to: a Stub Network
     (Link ID) Network/subnet number: 192.168.14.0
     (Link Data) Network Mask: 255.255.255.0
      Number of MTID metrics: 0
       TOS 0 Metrics: 1
```

```
                Router Link States (Area 1)

  LS age: 562
  Options: (No TOS-capability, DC)
  LS Type: Router Links
  Link State ID: 1.1.1.1
  Advertising Router: 1.1.1.1
  LS Seq Number: 8000000C
  Checksum: 0xEC26
  Length: 48
  Area Border Router
  AS Boundary Router
  Number of Links: 2

    Link connected to: another Router (point-to-point)
     (Link ID) Neighboring Router ID: 3.3.3.3
     (Link Data) Router Interface address: 192.168.13.1
      Number of MTID metrics: 0
       TOS 0 Metrics: 1

    Link connected to: a Stub Network
     (Link ID) Network/subnet number: 192.168.13.0
     (Link Data) Network Mask: 255.255.255.0
      Number of MTID metrics: 0
       TOS 0 Metrics: 1



                Router Link States (Area 50)

  LS age: 562
  Options: (No TOS-capability, DC)
  LS Type: Router Links
  Link State ID: 1.1.1.1
  Advertising Router: 1.1.1.1
  LS Seq Number: 80000012
  Checksum: 0x42CA
  Length: 48
  Area Border Router
  AS Boundary Router
  Number of Links: 2

    Link connected to: another Router (point-to-point)
     (Link ID) Neighboring Router ID: 2.2.2.2
     (Link Data) Router Interface address: 192.168.12.1
      Number of MTID metrics: 0
       TOS 0 Metrics: 1

    Link connected to: a Stub Network
     (Link ID) Network/subnet number: 192.168.12.0
     (Link Data) Network Mask: 255.255.255.0
      Number of MTID metrics: 0
       TOS 0 Metrics: 1
```

在R1上，我们的数据库中有以下外部LSA:

```
R1#show ip ospf database external

            OSPF Router with ID (1.1.1.1) (Process ID 1)
                Type-5 AS External Link States
```

```
  LS age: 706
  Options: (No TOS-capability, DC, Upward)
  LS Type: AS External Link
  Link State ID: 192.0.2.100 (External Network Number )
  Advertising Router: 1.1.1.1
  LS Seq Number: 80000001
  Checksum: 0xE617
  Length: 36
  Network Mask: /32
        Metric Type: 2 (Larger than any link state path)
        MTID: 0
        Metric: 20
        Forward Address: 192.168.12.2
        External Route Tag: 0

  LS age: 600
  Options: (No TOS-capability, DC, Upward)
  LS Type: AS External Link
  Link State ID: 192.0.2.100 (External Network Number )
  Advertising Router: 3.3.3.3
  LS Seq Number: 80000002
  Checksum: 0xBFAC
  Length: 36
  Network Mask: /32
        Metric Type: 2 (Larger than any link state path)
        MTID: 0
        Metric: 20
        Forward Address: 0.0.0.0
        External Route Tag: 0


R1#show ip ospf database nssa-external

            OSPF Router with ID (1.1.1.1) (Process ID 1)
                Type-7 AS External Link States (Area 50)

  LS age: 865
  Options: (No TOS-capability, Type 7/5 translation, DC, Upward)
  LS Type: AS External Link
  Link State ID: 192.0.2.100 (External Network Number )
  Advertising Router: 2.2.2.2
  LS Seq Number: 80000002
  Checksum: 0x32BC
  Length: 36
  Network Mask: /32
        Metric Type: 2 (Larger than any link state path)
        MTID: 0
        Metric: 20
        Forward Address: 192.168.12.2
        External Route Tag: 0
```

现在，让我们检查R1上首选的LSA:

```
R1#show ip ospf rib 192.0.2.100

            OSPF Router with ID (1.1.1.1) (Process ID 1)
                Base Topology (MTID 0)
```

```
OSPF local RIB
Codes: * - Best, > - Installed in global RIB
LSA: type/LSID/originator

*>  192.0.2.100/32, NSSA2, cost 20, fwd cost 1, tag 0, area 50
     SPF Instance 38, age 00:04:51
      contributing LSA: 7/192.0.2.100/2.2.2.2 (area 50)
      contributing LSA: 5/192.0.2.100/3.3.3.3
     Flags: RIB, HiPrio, ViaFwAddr, IntraNonBB, NSSA P-bit
      via 192.168.12.2, GigabitEthernet1 label 1048578
       Flags: RIB
       LSA: 7/192.0.2.100/2.2.2.2
```

如上输出所示，R1首选R2的LSA第7类。这是因为我们遵循的是RFC 3101，它具有以下路径计算首选项

1.具有P位集的第7类LSA。
2.第5类LSA。
3.路由器ID更高的LSA。

注意：请注意，如果当前LSA的功能与安装的LSA相同，则以下路径计算首选项适用。我们可以通过查看R1的第1类LSA来检验两个LSA的转发度量是否相同。

现在，如果我们从R2清除NSSA第7类LSA上的P位，我们将看到我们更喜欢从R3处获取第5类LSA:

摘自RFC 3101第2.4节

```
An NSSA internal AS boundary router must set the P-bit in the LSA
   header's option field of any Type-7 LSA whose network it wants
   advertised into the OSPF domain's full transit topology.  The LSAs of
   these networks must have a valid non-zero forwarding address.  If the
   P-bit is clear the LSA is not translated into a Type-5 LSA by NSSA
   border routers.

   When an NSSA border router originates both a Type-5 LSA and a Type-7
   LSA for the same network, then the P-bit must be clear in the Type-7
   LSA so that it isn't translated into a Type-5 LSA by another NSSA
   border router.
```

在继续清除R2上的P位之前，以下是R2的第7类LSA输出

```
R2#show ip ospf database nssa-external

            OSPF Router with ID (2.2.2.2) (Process ID 1)

                Type-7 AS External Link States (Area 50)

  LS age: 1215
  Options: (No TOS-capability, Type 7/5 translation, DC, Upward)
  LS Type: AS External Link
  Link State ID: 192.0.2.100 (External Network Number )
  Advertising Router: 2.2.2.2
  LS Seq Number: 80000002
  Checksum: 0x32BC
  Length: 36
  Network Mask: /32
        Metric Type: 2 (Larger than any link state path)
```

```
        MTID: 0
        Metric: 20
        Forward Address: 192.168.12.2
        External Route Tag: 0
```

当NSSA边界路由器为同一网络同时发起第5类LSA和第7类LSA时，可以清除P位。

```
R2#show ip ospf database nssa-external

            OSPF Router with ID (2.2.2.2) (Process ID 1)

                Type-7 AS External Link States (Area 50)

  LS age: 44
  Options: (No TOS-capability, No Type 7/5 translation, DC, Upward)
  LS Type: AS External Link
  Link State ID: 192.0.2.100 (External Network Number )
  Advertising Router: 2.2.2.2
  LS Seq Number: 80000003
  Checksum: 0xBFAD
  Length: 36
  Network Mask: /32
        Metric Type: 2 (Larger than any link state path)
        MTID: 0
        Metric: 20
        Forward Address: 0.0.0.0
        External Route Tag: 0
```

以下是上述输出的一些重要特征：

- 位 P — 此位用于告知 NSSA ABR 是否将类型 7 转换为类型 5。
- 无类型 7/5 转换表示位 P = 0。
- 类型 7/5 转换表示位 P = 1。
- 如果位P = 0，则NSSA ABR不得将此LSA转换为第5类。当NSSA ASBR也是NSSA ABR时，会发生这种情况。
- 如果位 P = 1，则 NSSA ABR 必须将此类型 7 LSA 转换为类型 5 LSA。如果有多个NSSA ABR，则具有最高路由器ID的NSSA ABR会执行此操作。

现在，当我们检查R1时，我们发现我们更喜欢第5类，而不是第7类LSA。

```
R1#show ip ospf rib 192.0.2.100

            OSPF Router with ID (1.1.1.1) (Process ID 1)

                Base Topology (MTID 0)

OSPF local RIB
Codes: * - Best, > - Installed in global RIB
LSA: type/LSID/originator

*>  192.0.2.100/32, Ext2, cost 20, fwd cost 1, tag 0
     SPF Instance 39, age 00:03:32
      contributing LSA: 7/192.0.2.100/2.2.2.2 (area 50)
      contributing LSA: 5/192.0.2.100/3.3.3.3
     Flags: RIB, HiPrio, IntraNonBB
      via 192.168.13.3, GigabitEthernet2 label 1048578
       Flags: RIB
       LSA: 5/192.0.2.100/3.3.3.3
```

场景 2

网络图

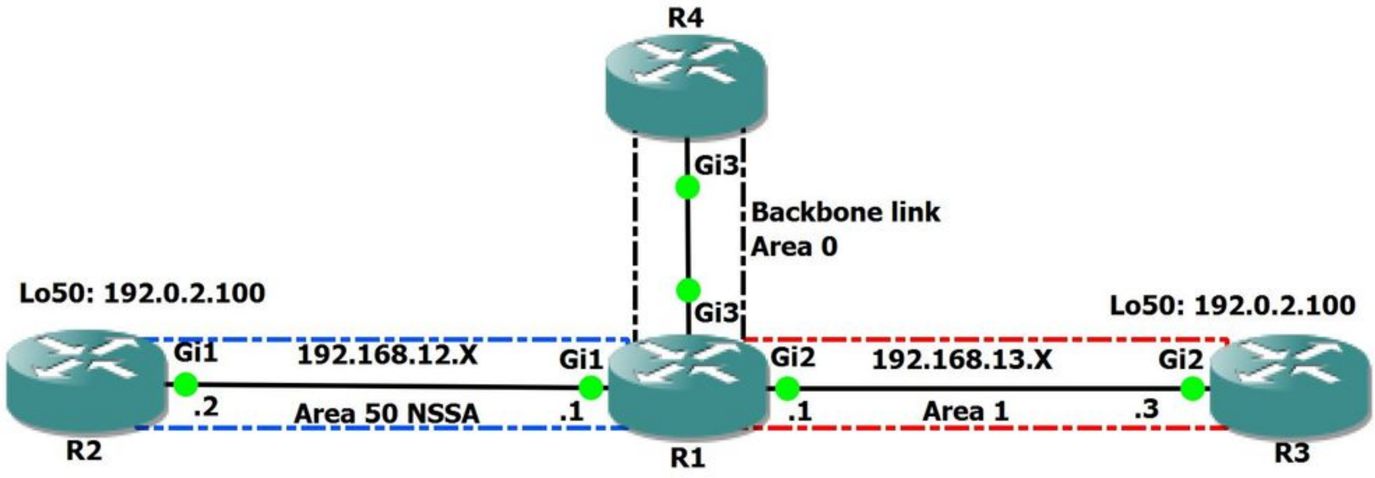

图 3

在此场景中，我们将观察使用RFC 1587进行外部路径计算时观察到的行为。RFC 3101合规性在IOS-XE设备上自动启用。要将RFC 3101兼容性替换为RFC 1587兼容性，以便在非末节区域(NSSA)区域边界路由器(ABR)中进行路由选择，请在路由器配置模式或地址系列配置模式下使用**compatiblerfc1587** 命令。要恢复RFC 3101兼容性，请使**用此**命令的no形式。

我们将对前缀192.0.2.100/32感兴趣，该前缀在R3和R2上重分发。首先，我们必须在R1上启用RFC 1587兼容性

```
R1#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)#router ospf 1
R1(config-router)#compatible rfc1587

R1#show ip ospf  | in RFC
 Supports NSSA (compatible with RFC 1587)
```

在R1上启用兼容性RFC 1587后，我们可以检查数据库中有哪些路径以及首选的LSA:

```
R1#show ip ospf database external

            OSPF Router with ID (1.1.1.1) (Process ID 1)
                Type-5 AS External Link States

  LS age: 115
  Options: (No TOS-capability, DC, Upward)
  LS Type: AS External Link
  Link State ID: 192.0.2.100 (External Network Number )
  Advertising Router: 3.3.3.3
  LS Seq Number: 80000003
  Checksum: 0xBDAD
  Length: 36
  Network Mask: /32
```

```
        Metric Type: 2 (Larger than any link state path)
        MTID: 0
        Metric: 20
        Forward Address: 0.0.0.0
        External Route Tag: 0


R1#show ip ospf database nssa-external

            OSPF Router with ID (1.1.1.1) (Process ID 1)
                Type-7 AS External Link States (Area 50)

  LS age: 48
  Options: (No TOS-capability, Type 7/5 translation, DC, Upward)
  LS Type: AS External Link
  Link State ID: 192.0.2.100 (External Network Number )
  Advertising Router: 2.2.2.2
  LS Seq Number: 80000005
  Checksum: 0x2CBF
  Length: 36
  Network Mask: /32
        Metric Type: 2 (Larger than any link state path)
        MTID: 0
        Metric: 20
        Forward Address: 192.168.12.2
        External Route Tag: 0
```

现在，让我们检查LSA是什么 在R1上首选：

```
R1#show ip ospf rib 192.0.2.100

            OSPF Router with ID (1.1.1.1) (Process ID 1)
                Base Topology (MTID 0)

OSPF local RIB
Codes: * - Best, > - Installed in global RIB
LSA: type/LSID/originator

*>  192.0.2.100/32, Ext2, cost 20, fwd cost 1, tag 0
     SPF Instance 44, age 00:01:56
      contributing LSA: 7/192.0.2.100/2.2.2.2 (area 50)
      contributing LSA: 5/192.0.2.100/3.3.3.3
     Flags: RIB, HiPrio, IntraNonBB, PartialSPF
      via 192.168.13.3, GigabitEthernet2 label 1048578
       Flags: RIB
       LSA: 5/192.0.2.100/3.3.3.3
```

首选第5类LSA。

在上述输出中，您可能还注意到R1没有将第7类路由转换为第5类路由，这是因为只有添加到路由表中的第7类路由才是转换的候选路由。

相关信息

- 技术支持 - Cisco Systems
- RFC 3101
- RFC 1587