

# OSPF External path selection: External Type-2 (E2) VS NSSA Type-2 (N2)

## Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Background Information](#)

[Abstract from RFC 3101 Section 2.5](#)

[Abstract from RFC 1587 Section 3.5](#)

[Scenario 1](#)

[Network Diagram](#)

[Scenario 2](#)

[Network Diagram](#)

[Related Information](#)

## Introduction

The purpose of this document is to demonstrate the Open Shortest Path First (OSPF) path selection behavior when a router receives both a Type-5 link-state advertisement (LSA) and a Type-7 LSA for a given external network. When redistribution is performed in a non-NSSA area, OSPF will inject a Type-5 LSA into the OSPF domain. Redistribution into an NSSA area creates a special type of LSA referred to as Type-7, which can only exist in an NSSA area.

## Prerequisites

Refer to the network diagram in Figure 1 as you use this document:




Figure 1

In the network diagram, there is both a non-backbone area 1 and a NSSA area 50 connected to

R1. R1 is an Area Border Router (ABR) connected to backbone area 0. Both R2 and R3 are responsible for redistributing the same prefix 192.0.2.100/32 into the OSPF domain.

## Requirements

Cisco recommends that you have knowledge of the OSPF protocol.

## Components Used

The information in this document is based on these software versions:

- Cisco CSR1000V Version 16.4.1

## Background Information

Cisco IOS-XE devices support RFC 3101 for external path calculation. RFC 1587 is obsoleted by RFC 3101 but RFC 1587-specific behavior can still be enabled through configuration. In Cisco IOS Release 15.1(2)S and later releases, the output of the show ip ospf command shows whether the device is using RFC 3101 or RFC 1587.

### Abstract from RFC 3101 Section 2.5

(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.

### Abstract from RFC 1587 Section 3.5

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

## Scenario 1

### Network Diagram




Figure 2

In this scenario, we will be looking at what behavior is observed when using RFC 3101 for external path calculation. We will be interested in prefix 192.0.2.100/32 that is redistributed on both R3 and R2.

Type-1 LSA from R1 is in the below output:

```
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

On R1 we have the following external LSAs in our database:

```
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
```

Now lets check what LSA is preferred on R1:

```
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
```

As we can see in the above output, R1 prefers LSAs Type-7 from R2. This is because we are following RFC 3101, which has the following path calculation preference

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

Note: Please be aware that the following path calculation preference is applicable if the current LSA is functionally the same as an installed LSA. We can verify that the forwarding metric for both LSAs are the same looking at Type-1 LSA of R1.

Now if we clear the P-bit on NSSA Type-7 LSA from R2 we will see that we will prefer Type-5 LSA from R3:

## Abstract from RFC 3101 Section 2.4

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.

Before we proceed with clearing the P-bit on R2, here is output of type-7 LSA from R2

```
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
```

The P-bit can be cleared when an NSSA border router originates both a Type-5 LSA and a Type-7 LSA for the same network.

```
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
```

Here are some important characteristics about the above output mentioned below:

- Bit P—This bit is used in order to tell the NSSA ABR whether to translate type 7 into type 5.
- No Type 7/5 translation means bit P = 0.
- Type 7/5 translation means bit P = 1.
- If bit P = 0, then the NSSA ABR must not translate this LSA into Type 5. This happens when NSSA ASBR is also an NSSA ABR.
- If bit P = 1, then the NSSA ABR must translate this type 7 LSA into a type 5 LSA. If there are multiple NSSA ABRs, the one with the highest router ID does this.

Now when we check on R1 we see that we prefer Type-5 over Type-7 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

```

## Scenario 2

### Network Diagram




Figure 3

In this scenario, we will be looking at what behavior is observed when using RFC 1587 for external path calculation. RFC 3101 compliance is automatically enabled on IOS-XE devices. To replace RFC 3101 compatibility with RFC 1587 compatibility for route selection in not-so-stubby area (NSSA) Area Border Routers (ABRs), use the **compatible rfc1587** command in router configuration mode or address family configuration mode. To restore RFC 3101 compatibility, use the **no** form of this command.

We will be interested in prefix 192.0.2.100/32 which is redistributed on both R3 and R2. First we must enable RFC 1587 compatibility on R1

```

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)

```

Once we have enabled Compatibility RFC 1587 on R1 we can check what paths are in our database and what LSA is preferred:

```

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
```

Now lets check what LSA is preferred on R1:

```
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
```

The Type-5 LSA is preferred.

In the above output, you might have also noticed that R1 is not translating Type-7 to Type-5, this is

because only Type-7 routes that have been added to the routing table are candidates for translation.

## Related Information

- [Technical Support - Cisco Systems](#)
- [RFC 3101](#)
- [RFC 1587](#)