OSPFv3 NSSA

- Feature History for OSPFv3 NSSA Option, on page 1
- OSPFv3 NSSA Option, on page 1
- How to configure OSPFv3 NSSA, on page 2
- Configuration example for NSSA for OSPFv3, on page 7
- Additional references, on page 9

Feature History for OSPFv3 NSSA Option

This table provides release and platform support information for the features explained in this module.

These features are available in all the releases subsequent to the one they were introduced in, unless noted otherwise.

Release	Feature Name and Description	Supported Platform
Cisco IOS XE 17.18.1	OSPFv3 NSSA Option: The OSPFv3 NSSA Option is a network optimization feature that allows a central site that runs on OSPF protocol to connect to a remote site that runs on any other routing protocol.	Cisco C9610 Series Smart Switches

OSPFv3 NSSA Option

The OSPFv3 NSSA Option is a network optimization feature that allows a central site that runs on OSPF protocol to connect to a remote site that runs on any other routing protocol.

Without OSPFv3 NSSA option

Without OSPFv3 NSSA option, the following happens:

- OSPF stub areas do not allow external route redistribution (from protocols like RIP).
- Two routing protocols must be maintained (OSPF in the core, RIP at the edge).

- The stub area cannot import external routes from the remote network.

Key components in OSPFv3 NSSA option

- Autonomous System Border Router (ASBR): ASBRs connect an OSPFv3 network and a non-OSPFv3 network (or another OSPFv3 domain if routes are being explicitly redistributed). The ASBR redistributes routes from RIP (or another protocol) into OSPFv3 NSSA as type 7 LSAs.
- Area Border Router (ABR): ABRs connect regular areas to the backbone (Area 0). An ABR converts type 7 LSAs to type 5 LSAs, allowing the rest of the OSPFv3 domain to learn about the external routes.
- Route summarization: Route summarization is the process of combining a group of contiguous network addresses into a single summary address. This reduces the size of routing tables and improves routing efficiency. In OSPFv3, summarization is performed at Area Border Routers (ABRs) and Autonomous System Boundary Routers (ASBRs). If an OSPFv3 area contains multiple contiguous subnets, the ABR can be configured to advertise a single summary route to other areas instead of multiple individual routes.

How to configure OSPFv3 NSSA

The following sections provide configuration information on how to configure OSPFv3 NSSA.

Configure an OSPFv3 NSSA area and its parameters

Perform this procedure to configure an OSPFv3 NSSA area and it parameters.

Procedure

Step 1 **enable**

Example:

```
Device> enable
```

Enables privileged EXEC mode.

Enter your password, if prompted.

Step 2 **configure terminal**

Example:

```
Device# configure terminal
```

Enters global configuration mode.

Step 3 **router ospfv3** *process-id*

Example:

```
Device(config)# router ospf 15
```

Enables OSPFv3 routing and enters router configuration mode.

process-id: The process ID is an internally used, identification parameter that is locally assigned. Each OSPF has a unique process ID.

Process ID can be a positive integer from 1 to 65535.

Step 4 **area** *area-id* **nssa default-information-originate nssa-only**

Example:

```
Device(config-router)# area 1 nssa default-information-originate nssa-only
```

Configures an NSSA area and sets the default advertisement to this NSSA area.

- **default-information-originate**: Select on an ABR to allow importing type 7 LSAs into the NSSA.
- **nssa-only**: Advertises the summary only as a Type 7 LSA into NSSA areas

Step 5 **address-family** {**ipv4** |**ipv6**} [**unicast**]

Example:

```
Device(config-router)# address-family ipv4 unicast
OR
Device(config-router)# address-family ipv6 unicast
```

(Optional) Enables address family configuration mode for Open Shortest Path First version 3 (OSPFv3).

- The **address-family ipv4 unicast** command configures an IPv4 address family.
- The **address-family ipv6 unicast** command configures an IPv6 address family.

Step 6

- For IPv4

  summary-prefix {*ip-prefix* | *ip-address-mask*} [**not-advertise** | [**tag** *tag-value*] [**nssa-only**]]
- For IPv6

  summary-prefix *ipv6-prefix* [**not-advertise** | [**tag** *tag-value*] [**nssa-only**]]

Example:

```
Device(config-router-af)# summary-prefix 10.1.0.0/16 nssa-only
OR
Device(config-router-af)# summary-prefix 2001:DB8::/32 nssa-only
```

Controls the route summarization and filtering during the translation and limits the summary to NSSA areas

- *ip-prefix*: The summarized IPv4 network address and the network mask for the summary route.
- *ipv6-prefix*: The summarized IPv6 network address and the network mask for the summary route.
- **not-advertise**: (Optional) Suppresses the summary route advertisement.
- **tag** *tag*: (Optional) Associates a tag with the summary route for policy or routing decisions.
- **nssa-only**: (Optional) Advertises the summary only as a Type 7 LSA for NSSA areas.

Step 7 **exit**

Example:

```
Device(config-router-af)# exit
```

Exits address-family router configuration mode and returns to the router configuration mode.

Step 8 **redistribute** *protocol* [*process-id*] [*as-number*] [**include-connected** {**level-1** | **level-1-2** | **level-2**} [**metric** *metric-value* | **transparent**] [**metric-type** *type-value*] [**match** {**internal** | **external 1** | **external 2**}] [**subnets**] [**nssa-only**] [**tag** *tag-value*] [**route-map** *map-tag*]

Example:

```
Device(config-router)# redistribute rip subnets
```

Redistributes routes from one routing domain into another routing domain.

- *protocol*: Specifies the source of the routes to be redistributed. Common options include:
  - **connected**: Redistributes directly connected networks that are not already part of the OSPF process.
  - **static**: Redistributes static routes configured on the router.
  - **bgp**: Redistributes routes learned via BGP.
  - **eigrp**: Redistributes routes learned via EIGRP.
  - **rip**: Redistributes routes learned via RIP.
  - **ospf**: Redistributes routes from another OSPF process (e.g., if running multiple OSPF instances or different process IDs).
  - **isis**: Redistributes routes learned via IS-IS.
- *process-id*: (Optional) If the source protocol (e.g., EIGRP, OSPF, RIP) uses a process ID, you specify it here to identify the specific instance from which routes should be taken.
- *as-number*: (Optional) If the source `protocol` (e.g., BGP or EIGRP) uses an Autonomous System (AS) number, you specify it here. This is often used interchangeably with `process-id` for protocols that use AS numbers.
- **include-connected** {**level-1** | **level-1-2** | **level-2**}: (Optional) This parameter is typically specific to IS-IS redistribution into OSPF, or sometimes OSPF into IS-IS, or even OSPF into OSPF (though less common). It refers to the IS-IS routing levels:
  - level-1: Redistributes Level-1 IS-IS routes.
  - level-2: Redistributes Level-2 IS-IS routes.
  - level-1-2: Redistributes both Level-1 and Level-2 IS-IS routes.
  - This part of the syntax suggests a context where IS-IS is the source protocol.
- **metric** {*metric-value* | **transparent**}: This is a critical parameter for redistribution. When routes are redistributed from one protocol to another, they need to be assigned a metric that is meaningful to the destination protocol.
  - *metric-value*: A specific numerical value to assign as the metric. The meaning of this value depends on the destination protocol (e.g., hop count for RIP, cost for OSPF).
  - **transparent**: This attempts to carry the original metric information across, but its effectiveness can vary depending on the protocols involved.
- **metric-type** *type-value*: (Optional) For OSPF, this is crucial for how external routes are handled:

  - type-1 (E1): The cost to the external destination is the sum of the external metric and the internal OSPF cost to reach the ASBR. This is often preferred for more accurate path selection.
  - type-2 (E2): The cost to the external destination is *only* the external metric, regardless of the internal OSPF cost to reach the ASBR. This is the default if not specified. Type 2 routes are always preferred over Type 1 routes if their external metric is the same.

- **match** {**internal** | **external 1** | **external 2**}: This parameter is specific to OSPF when redistributing OSPF routes into another protocol. It allows you to specify which *types* of OSPF routes to redistribute:
  - **internal**: Redistributes OSPF intra-area (Type 1) and inter-area (Type 2) routes.
  - **external 1**: Redistributes OSPF external Type 1 routes.
  - **external 2**: Redistributes OSPF external Type 2 routes.
- **nssa-only**: (Optional) This keyword is specifically used when redistributing routes into a Not-So-Stubby Area (NSSA). When nssa-only is specified, the redistributed routes will be advertised as Type 7 LSAs within the NSSA, and then translated to Type 5 LSAs by the NSSA ABR for propagation to other areas. Without this, redistribution into an NSSA might not be allowed or might behave differently
- **tag** *tag-value*: (Optional) Assigns a numerical tag to the redistributed routes. This tag can be used later in route maps or distribute lists for filtering or policy-based routing.
- **route-map** *map-tag*: (Optional) A very powerful and commonly used parameter. It references a `route-map` that provides granular control over which routes are redistributed and how their attributes (like metric, metric-type, tag) are modified during the redistribution process.
- **subnets**: This parameter is specific to OSPF. By default, OSPF only redistributes classful networks. To redistribute subnetted networks (which is almost always desired in modern networks), you must include the **subnets** keyword. Without it, only major network numbers (e.g., 10.0.0.0/8, 192.168.1.0/24) would be redistributed.

Step 9 **end**

Example:

```
Device(config-router)# end
```

Exits router configuration mode and returns to privileged EXEC mode.

Configure an NSSA ABR as a forced NSSA LSA translator for OSPFv3

Perform this procedure to configure an NSSA ABR as a forced NSSA LSA translator for OSPFv3.

Procedure

Step 1 **enable**

Example:

```
Device> enable
```

Enables privileged EXEC mode.

Enter your password, if prompted.

Step 2 **configure terminal**

Example:

```
Device# configure terminal
```

Enters global configuration mode.

Step 3 **router ospfv3** *process-id*

Example:

```
Device(config)# router ospfv3 15
```

Enables OSPFv3 routing and enters router configuration mode.

process-id: The process ID is an internally used, identification parameter that is locally assigned. Each OSPF has a unique process ID.

Process ID can be a positive integer from 1 to 65535.

Step 4 **area** *area-id* **nssa translate type7 always**

Example:

```
Device(config-router)# area 10 nssa translate type7 always
```

Configures a NSSA ABR device as a forced NSSA LSA translator.

Note
This command can be used if RFC 3101 is disabled and RFC 1587 is used.

- *area-id*: The OSPF area number or address in dotted-decimal format.
- **always**: Manually force this ABR to always be the translator, ensuring consistent LSA translation and avoiding translator role changes due to device restarts or network topology changes.

Step 5 **area** *area-id* **nssa translate type7 suppress-fa**

Example:

```
Device(config-router)# area 10 nssa translate type7 suppress-fa
```

Allows ABR to suppress the forwarding address in translated Type-5 LSA.

Step 6 **end**

Example:

```
Device(config-router)# end
```

Exits router configuration mode and returns to privileged EXEC mode.

Disable RFC 3101 compatibility and enable RFC 1587 compatibility (OSPFv3)

Perform this task to disable RFC 3101 compatibility and enable RFC 1587 compatibility.

Procedure

Step 1 **enable**

Example:

```
Device> enable
```

Enables privileged EXEC mode.

Enter your password, if prompted.

Step 2 **configure terminal**

Example:

```
Device# configure terminal
```

Enters global configuration mode.

Step 3 **router ospfv3** *process-id*

Example:

```
Device(config)# router ospfv3 1
```

Enables OSPFv3 routing and enters router configuration mode.

process-id: The process ID is an internally used, identification parameter that is locally assigned. Each OSPF has a unique process ID.

Process ID can be a positive integer from 1 to 65535.

Step 4 **compatible rfc1587**

Example:

```
Device(config-router)# compatible rfc1587
```

Enables the device to be RFC 1587 compliant.

Step 5 **end**

Example:

```
Device(config-router)# end
```

Enables the device to be RFC 1587 compliant.

Configuration example for NSSA for OSPFv3

Use the **show ospfv3** command to confirm that the device is acting as an ASBR and that the OSPFv3 Area 1 has been configured as a NSSA area.

```
Device#show ospfv3

 OSPFv3 1 address-family ipv4
 Router ID 3.3.3.3
```

```
Supports NSSA (compatible with RFC 1587)
It is an autonomous system boundary router
Redistributing External Routes from,
   static
Router is not originating router-LSAs with maximum metric
Initial SPF schedule delay 5000 msecs
Minimum hold time between two consecutive SPFs 10000 msecs
Maximum wait time between two consecutive SPFs 10000 msecs
Minimum LSA interval 5 secs
Minimum LSA arrival 1000 msecs
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msecs
Retransmission pacing timer 66 msecs
Number of external LSA 0. Checksum Sum 0x000000
Number of areas in this router is 1. 0 normal 0 stub 1 nssa
Graceful restart helper support enabled
Reference bandwidth unit is 100 mbps
RFC1583 compatibility enabled
   Area 1
       Number of interfaces in this area is 1
       It is a NSSA area
       Configured to translate Type-7 LSAs, inactive (RFC3101 support disabled)
       Perform type-7/type-5 LSA translation, suppress forwarding address
       Area has no authentication
       SPF algorithm last executed 00:00:07.160 ago
       SPF algorithm executed 3 times
       Area ranges are
       Number of LSA 3. Checksum Sum 0x0245F0
       Number of opaque link LSA 0. Checksum Sum 0x000000
       Number of DCbitless LSA 0
       Number of indication LSA 0
       Number of DoNotAge LSA 0
       Flood list length 0
```

The table below describes the significant **show ip ospf** display fields and their descriptions.

Table 1: show ospfv3 Field Descriptions

Field	Description
Supports NSSA (compatible with RFC 1587)	Specifies that RFC 1587 is active or that the OSPFv3 NSSA area is RFC 1587 compatible.
Configured to translate Type-7 LSAs, inactive (RFC3101 support disabled)	Specifies that the OSPFv3 NSSA area has an ABR device configured to act as a forced translator of Type 7 LSAs. However, it is inactive because RFC 3101 is disabled.

The output of the router LSA in LSDB displays Nt-Bit if it is set in the header of LSA.

```
Router Link States (Area 1)

  LS age: 94
  Options: (N-Bit, R-bit, DC-Bit, AF-Bit, Nt-Bit)
  LS Type: Router Links
  Link State ID: 0
  Advertising Router: 2.2.2.2
  LS Seq Number: 80000002
  Checksum: 0x8AD5
  Length: 56
  Area Border Router
```

```
AS Boundary Router
Unconditional NSSA translator
Number of Links: 2
```

The “Unconditional NSSA translator” line indicates that the NSSA ASBR device is as a forced NSSA LSA translator.

Additional references

Table 2: Standards and RFCs

Standard/RFC	Title
RFC 3101	The OSPF Not-So-Stubby Area (NSSA) Option
RFC 1587	The OSPF NSSA Option