



OSPF Limit on Number of Redistributed Routes

- [Feature History for OSPF Limit on Number of Redistributed Routes, on page 1](#)
- [OSPF Limit on Number of Redistributed Routes, on page 1](#)
- [Guidelines to configure OSPF Limit on Number of Redistributed Routes, on page 2](#)
- [Configure an OSPF Limit on the Number of Redistributed Routes, on page 2](#)
- [Configuration examples for OSPF Limit on Number of Redistributed Routes, on page 12](#)

Feature History for OSPF Limit on Number of Redistributed Routes

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	OSPF Limit on Number of Redistributed Routes: OSPF Limit on Number of Redistributed Routes is a network enhancement feature that limits the maximum number of prefixes (routes) that can be redistributed into OSPF from other protocols or OSPF processes.	Cisco C9350 Series Smart Switches Cisco C9610 Series Smart Switches

OSPF Limit on Number of Redistributed Routes

OSPF Limit on Number of Redistributed Routes is a network enhancement feature that limits the maximum number of prefixes (routes) that can be redistributed into OSPF from other protocols or OSPF processes.

Without such a limit, an access switch acting as an Autonomous System Boundary Router (ASBR) could potentially receive a massive number of routes from another routing domain, for example a full BGP routing table from the internet, or a large number of routes from an EIGRP domain, and then redistribute all of them into OSPF.

How OSPF Limit on Number of Redistributed Routes works

When redistribution occurs, for example from BGP, EIGRP, or another OSPF process, the device will only inject up to the specified number of external routes into the OSPF domain. If the configured limit is exceeded, additional routes are not redistributed, helping to protect the OSPF domain from route table overload.

Guidelines to configure OSPF Limit on Number of Redistributed Routes

- This feature is supported only for the IPv6 address family.
- You must have OSPF configured in your network either along with another protocol, or another OSPF process for redistribution.

Configure an OSPF Limit on the Number of Redistributed Routes

Follow the steps in each of these tasks to configure an OSPF limit on the number of redistributed routes

Procedure

-
- | | |
|---------------|--|
| Step 1 | Limit the number of OSPF redistributed routes |
| Step 2 | Request a warning message about the number of routes redistributed into OSPF |
| Step 3 | Limit the number of OSPFv3 redistributed routes |
| Step 4 | Request a warning message about the number of routes redistributed into OSPFv3 |
-

Limit the number of OSPF redistributed routes

This task limits the number of OSPF redistributed routes. If the maximum configured value is reached, no more routes are distributed.

Procedure

-
- | | |
|---------------|--|
| Step 1 | enable

Example:
<code>Device> enable</code>

Enables privileged EXEC mode.

Enter your password, if prompted. |
| Step 2 | configure terminal |

Example:

```
Device# configure terminal
```

Enters global configuration mode.

Step 3 `router ospf process-id [vrf vrf-name]`**Example:**

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

Enables OSPF routing and enters router configuration mode.

- *process-id*: The process ID is an internally used identification parameter that is locally assigned. Each OSPF process has a unique process ID.
Process ID can be a positive integer from 1 to 65535.
- *vrf*: Indicates that the OSPF process is being configured for a specific VRF.
- *vrf-name*: Specifies the name of the VRF for which this OSPF process is being created.

Step 4 `redistribute protocol [process-id] [as-number] [include-connected {level-1 | level-1-2 | level-2}] [metric metric-value] [metric-type type-value] [nssa-only] [tag tag-value] [route-map map-tag]`**Example:**

```
Device(config-router)# redistribute eigrp 10
```

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 (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 (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***: (Optional) Specifies the metric (cost) that OSPF will assign to the redistributed routes. This value is used by OSPF to calculate the path cost to these external destinations. If not specified, a default metric (often 1 or 20, depending on the source protocol and platform) will be used.
- **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.
- **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.

Step 5 **redistribute maximum-prefix *maximum* [*threshold*]**

Example:

```
Device(config-router-af) # redistribute maximum-prefix 100 80
```

Sets a maximum number of IP prefixes that are allowed to be redistributed into OSPF.

- ***maximum***: The maximum value.
The default value is set at 10240 routes.
- ***threshold***: The threshold value defaults to 75 percent.

Note

If the warning-only keyword is configured in this command, no limit is enforced; a warning message is logged.

Step 6 **end**

Example:

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

Returns to privileged EXEC mode.

Request a warning message about the number of routes redistributed into OSPF

This task limits the number of OSPF redistributed routes. If the maximum configured value is reached, a warning message is displayed but routes are still redistributed.

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 ospf *process-id* [*vrf vrf-name*]

Example:

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

Enables OSPF routing and enters router configuration mode.

- *process-id*: The process ID is an internally used identification parameter that is locally assigned. Each OSPF process has a unique process ID.
Process ID can be a positive integer from 1 to 65535.
- *vrf*: Indicates that the OSPF process is being configured for a specific VRF.
- *vrf-name*: Specifies the name of the VRF for which this OSPF process is being created.

Step 4 redistribute *protocol* [*process-id*] [*as-number*] [**include-connected** {*level-1* | *level-1-2* | *level-2*}] [*metric metric-value*] [**metric-type** *type-value*] [**nssa-only**] [**tag** *tag-value*] [**route-map** *map-tag*]

Example:

```
Device(config-router)# redistribute eigrp 10
```

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 (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 (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**: (Optional) Specifies the metric (cost) that OSPF will assign to the redistributed routes. This value is used by OSPF to calculate the path cost to these external destinations. If not specified, a default metric (often 1 or 20, depending on the source protocol and platform) will be used.
- **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.
- **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.

Step 5 **redistribute maximum-prefix maximum [threshold] [warning-only]**

Example:

```
Device(config-router-af) # redistribute maximum-prefix 100 80 warning-only
```

Causes a warning message to be logged when the maximum number of IP prefixes have been redistributed to OSPF.

- Because the **warning-only** keyword is included, no limit is imposed on the number of redistributed prefixes into OSPF.
- The *threshold* value defaults to 75 percent.
- This example causes two warnings:
 - one at 80 percent of 1000 (800 routes redistributed) and
 - another at 1000 routes redistributed.

Step 6 **end**

Example:

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

Returns to privileged EXEC mode.

Limit the number of OSPFv3 redistributed routes

This task limits the number of OSPFv3 redistributed routes. If the maximum configured value is reached, no more routes are distributed.

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 ospf process-id [vrf vrf-name]**

Example:

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

Enables OSPF routing and enters router configuration mode.

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

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

- **vrf**: Indicates that the OSPF process is being configured for a specific VRF.
- *vrf-name*: Specifies the name of the VRF for which this OSPF process is being created.

Step 4 **address-family ipv6 [unicast]**

Example:

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

Enters IPv6 address family configuration mode.

unicast: Specifies the configuration is for IPv6 unicast routing.

Step 5 **redistribute protocol [process-id] [as-number] [include-connected {level-1 | level-1-2 | level-2} [metric metric-value] [metric-type type-value] [nssa-only] [tag tag-value] [route-map map-tag]**

Example:

```
Device(config-router-af)# redistribute eigrp 10
```

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 (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 (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*: (Optional) Specifies the metric (cost) that OSPF will assign to the redistributed routes. This value is used by OSPF to calculate the path cost to these external destinations. If not specified, a default metric (often 1 or 20, depending on the source protocol and platform) will be used.
- **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.
- **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.

Step 6 **redistribute maximum-prefix** *maximum* [*threshold*] [**warning-only**]

Example:

```
Device(config-router-af)# redistribute maximum-prefix 100 80 warning-only
```

Causes a warning message to be logged when the maximum number of IP prefixes have been redistributed to OSPFv3.

- Because the **warning-only** keyword is included, no limit is imposed on the number of redistributed prefixes into OSPF.
- The *threshold* value defaults to 75 percent.
- This example causes two warnings:
 - one at 80 percent of 1000 (800 routes redistributed) and
 - another at 1000 routes redistributed.

Step 7 **exit-address-family**

Example:

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

Exits IPv6 address family configuration mode.

Step 8 **end**

Example:

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

Returns to privileged EXEC mode.

Request a warning message about the number of routes redistributed into OSPFv3

This task limits the number of OSPFv3 redistributed routes. If the maximum configured value is reached, no more routes are distributed.

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 ospf *process-id* [**vrf** *vrf-name*]**

Example:

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

Enables OSPF routing and enters router configuration mode.

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

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

- **vrf**: Indicates that the OSPF process is being configured for a specific VRF.
- **vrf-name**: Specifies the name of the VRF for which this OSPF process is being created.

Step 4 **address-family ipv6 [**unicast**]**

Example:

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

Enters IPv6 address family configuration mode.

unicast: Specifies the configuration is for IPv6 unicast routing.

Step 5 **redistribute *protocol* [*process-id*] [*as-number*] [**include-connected** {**level-1** | **level-1-2** | **level-2**} [*metric metric-value*] [**metric-type** *type-value*] [**nssa-only**] [**tag** *tag-value*] [**route-map** *map-tag*]**

Example:

```
Device(config-router-af)# redistribute eigrp 10
```

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 (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 (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*: (Optional) Specifies the metric (cost) that OSPF will assign to the redistributed routes. This value is used by OSPF to calculate the path cost to these external destinations. If not specified, a default metric (often 1 or 20, depending on the source protocol and platform) will be used.
- **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.
- **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.

Step 6 **redistribute maximum-prefix** *maximum* [*threshold*]

Example:

```
Device(config-router-af) # redistribute maximum-prefix 100 80
```

Sets a maximum number of IP prefixes that are allowed to be redistributed into OSPF.

- *maximum*: The maximum value.
The default value is set at 10240 routes.
- *threshold*: The threshold value defaults to 75 percent.

Note

If the warning-only keyword is configured in this command, no limit is enforced; a warning message is logged.

Step 7 **exit-address-family**

Example:

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

Exits IPv6 address family configuration mode.

Step 8 **end**

Example:

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

Returns to privileged EXEC mode.

Configuration examples for OSPF Limit on Number of Redistributed Routes

The following sections provide configuration examples for OSPF Limit on Number of Redistributed Routes.

Example: OSPF Limit on Number of Redistributed Routes

This example shows how to set a maximum of 1200 prefixes that can be redistributed into the OSPF process 1. Prior to reaching the limit, when the number of prefixes that are redistributed reaches 80 percent of 1200 (960 prefixes), a warning message is logged. Another warning message is logged when the limit is reached and no more routes are redistributed.

```
Device> enable
Device# configure terminal
```

```
Device(config)# router ospf 1
Device(config-router-af)# redistribute static subnets
Device(config-router-af)# redistribute maximum-prefix 1200 80
```

This example shows how to set a maximum of 1200 prefixes that can be redistributed into the OSPFv3 process 1.

```
Device> enable
Device# configure terminal
Device(config)# router ospfv3 1
Device(config-router)# address-family ipv6
Device(config-router-af)# redistribute static subnets
Device(config-router-af)# redistribute maximum-prefix 1200 80
```

Example: Requesting a warning message about the number of redistributed routes

This example shows how to enable two warning messages to be logged, the first if the number of prefixes that are redistributed reaches 85 percent of 600 (510 prefixes), and the second if the number of redistributed routes reaches 600. However, the number of redistributed routes is not limited.

```
Device> enable
Device# configure terminal
Device(config)# router ospf 11
Device(config-router-af)# redistribute eigrp 10 subnets
Device(config-router-af)# redistribute maximum-prefix 600 85 warning-only
```

This example shows how to enable two warnings to be logged for an OSSPv3 process.

```
Device> enable
Device# configure terminal
Device(config)# router ospfv3 11
Device(config-router)# address-family ipv6
Device(config-router-af)# redistribute eigrp 10 subnets
Device(config-router-af)# redistribute maximum-prefix 600 85 warning-only
```

Example: Requesting a warning message about the number of redistributed routes