

OSPF Commands

All commands applicable for the Cisco NCS 5500 Series Router are also supported on the Cisco NCS 540 Series Router that is introduced from Cisco IOS XR Release 6.3.2. References to earlier releases in Command History tables apply to only the Cisco NCS 5500 Series Router.

Note

- Starting with Cisco IOS XR Release 6.6.25, all commands applicable for the Cisco NCS 5500 Series Router are also supported on the Cisco NCS 560 Series Routers.
- Starting with Cisco IOS XR Release 6.3.2, all commands applicable for the Cisco NCS 5500 Series Router are also supported on the Cisco NCS 540 Series Router.
- References to releases before Cisco IOS XR Release 6.3.2 apply to only the Cisco NCS 5500 Series Router.
- Cisco IOS XR Software Release 7.0.1 specific updates are not applicable for the following variants of Cisco NCS 540 Series Routers:
  - N540-28Z4C-SYS-A
  - N540-28Z4C-SYS-D
  - N540X-16Z4G8Q2C-A
  - N540X-16Z4G8Q2C-D
  - N540X-16Z8Q2C-D
  - N540-12Z20G-SYS-A
  - N540-12Z20G-SYS-D
  - N540X-12Z16G-SYS-A
  - N540X-12Z16G-SYS-D

This module describes the commands used to configure and monitor the Open Shortest Path First (OSPF) routing protocol.

For detailed information about OSPF concepts, configuration tasks, and examples, see the Implementing OSPF on Cisco NCS 5500 Series Routers module in the *Routing Configuration Guide for Cisco NCS 5500 Series Routers*.

- address-family (OSPF), on page 5
- adjacency stagger (OSPF), on page 6
- area (OSPF), on page 8
- authentication (OSPF), on page 10
- authentication-key (OSPF), on page 12
- auto-cost (OSPF), on page 14
- capability opaque disable, on page 15
- clear ospf process, on page 16
- clear ospf redistribution, on page 18
- clear ospf routes, on page 19
- clear ospf statistics, on page 20
- cost (OSPF), on page 22
- cost-fallback (OSPF), on page 24
- database-filter all out (OSPF), on page 26
- dead-interval (OSPF), on page 27
- default-cost (OSPF), on page 29
- default-information originate (OSPF), on page 30
- default-metric (OSPF), on page 32
- demand-circuit (OSPF), on page 33
- disable-dn-bit-check, on page 35
- distance (OSPF), on page 36
- distance ospf, on page 38
- distribute-list, on page 40
- domain-id (OSPF), on page 42
- domain-tag, on page 44
- exchange-timer, on page 45
- fast-reroute (OSPFv2), on page 47
- fast-reroute per-link exclude interface, on page 49
- fast-reroute per-prefix exclude interface (OSPFv2), on page 50
- fast-reroute per-prefix use-candidate-only (OSPFv2), on page 51
- flood-reduction (OSPF), on page 52
- hello-interval (OSPF), on page 54
- ignore lsa mospf, on page 56
- interface (OSPF), on page 57
- log adjacency changes (OSPF), on page 59
- loopback stub-network , on page 60
- max external-lsa, on page 61
- max-lsa, on page 63
- max-metric, on page 65
- maximum interfaces (OSPF), on page 68
- maximum paths (OSPF), on page 69
- maximum redistributed-prefixes (OSPF), on page 70
- message-digest-key, on page 72

- mpls ldp auto-config (OSPF), on page 75
- mpls ldp sync (OSPF), on page 76
- mtu-ignore (OSPF), on page 77
- multi-area-interface, on page 79
- neighbor (OSPF), on page 81
- neighbor database-filter all out, on page 83
- network (OSPF), on page 84
- nsf (OSPF), on page 86
- nsf flush-delay-time (OSPF), on page 88
- nsf interval (OSPF), on page 89
- nsf lifetime (OSPF), on page 90
- nssa (OSPF), on page 91
- ospf name-lookup, on page 93
- packet-size (OSPF), on page 94
- passive (OSPF), on page 96
- priority (OSPF), on page 98
- protocol shutdown, on page 100
- queue dispatch flush-lsa, on page 101
- queue dispatch incoming, on page 102
- queue dispatch rate-limited-lsa, on page 103
- queue dispatch spf-lsa-limit, on page 104
- queue limit, on page 105
- range (OSPF), on page 106
- redistribute (OSPF), on page 108
- retransmit-interval (OSPF), on page 113
- route-policy (OSPF), on page 115
- router-id (OSPF), on page 116
- router ospf (OSPF), on page 118
- security ttl (OSPF), on page 119
- sham-link, on page 121
- show ospf, on page 122
- show ospf border-routers, on page 126
- show ospf database, on page 128
- show ospf flood-list, on page 143
- show ospf interface, on page 145
- show ospf message-queue, on page 147
- show ospf neighbor, on page 149
- show ospf request-list, on page 158
- show ospf retransmission-list, on page 160
- show ospf routes, on page 162
- show ospf sham-links, on page 167
- show ospf summary-prefix, on page 169
- show ospf virtual-links, on page 171
- show protocols (OSPF), on page 173
- snmp context (OSPF), on page 175
- snmp trap (OSPF), on page 177

- snmp trap rate-limit (OSPF), on page 178
- spf prefix-priority (OSPFv2), on page 179
- stub (OSPF), on page 181
- summary-prefix (OSPF), on page 182
- timers lsa group-pacing, on page 184
- timers lsa min-arrival, on page 185
- timers throttle lsa all (OSPF), on page 186
- timers throttle spf (OSPF), on page 188
- transmit-delay (OSPF), on page 190
- virtual-link (OSPF), on page 192
- vrf (OSPF), on page 194

address-family (OSPF)

To enter address family configuration mode for Open Shortest Path First (OSPF), use the **address-family** command in the appropriate mode. To disable address family configuration mode, use the **no** form of this command.

address-family ipv4 [**unicast**]
no address-family ipv4 [**unicast**]

Syntax Description

ipv4	Specifies IP Version 4 (IPv4) address prefixes.
unicast	(Optional) Specifies unicast address prefixes.

Command Default

An address family is not specified.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

OSPF version 2 automatically provides routing services for IPv4 unicast topologies, so this command is redundant.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure the OSPF router process with IPv4 unicast address prefixes:

```
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# address-family ipv4 unicast
```

adjacency stagger (OSPF)

To configure staggering of OSPF adjacency during reload, process restart, and process clear, use the **adjacency stagger** command in XR Config mode. To turn off adjacency staggering, either use the **disable** keyword or use the **no** form of this command.

adjacency stagger {**disable** | *initial-num-nbr* *max-num-nbr*}
no adjacency stagger

Syntax Description

disable	Disables adjacency staggering.
initial-num-nbr	The initial number of simultaneous neighbors allowed to form adjacency to FULL in any area to bring up to FULL after a router reload, OSPF process restart, or OSPF process clear. Range is 1-65535. Default is 2.
max-num-nbr	The subsequent number of simultaneous neighbors allowed to form adjacency, per OSPF instance, after the initial set of OSPF neighbors have become FULL. Range is 1-65535. Default is 64.

Command Default

OSPF adjacency staggering is enabled.

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Staggering of the OSPF adjacency during reload, process restart (without NSR or graceful-restart), and process clear reduces the overall adjacency convergence time.

Initially, allow 2 (configurable) neighbors to form adjacency to FULL per area. After the first adjacency reaches FULL, up to 64 (configurable) neighbors can form adjacency simultaneously for the OSPF instance (all areas). However, areas without any FULL adjacency is restricted by the initial area limit.

Note Adjacency stagger and OSPF nonstop forwarding (NSF) are mutually exclusive. Adjacency stagger will not be activated if **nsf** is configured under router ospf configuration.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure adjacency stagger for a 2 neighbors initially and for a maximum of 3 neighbors:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# adjacency stagger 2 3
```

area (OSPF)

To configure an Open Shortest Path First (OSPF) area, use the **area** command in the appropriate mode. To terminate an OSPF area, use the **no** form of this command.

area *area-id*
no area *area-id*

Syntax Description

area-id	Identifier of an OSPF area. The *area-id* argument can be specified as either a decimal value or an IP address (dotted decimal) format. Range is 0 to 4294967295.

Command Default

No OSPF area is defined.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **area** command to explicitly configure an area. Commands configured under the area configuration mode (such as the **interface** [OSPF] and **authentication** commands), are automatically bound to that area.

To modify or remove the area, the *area-id* argument format must be the same as the format used when creating the area. Otherwise, even if the actual 32-bit value matches, the area is not matched. For example, if you create an area with an *area-id* of 10 it would not match an *area-id* of 0.0.0.10.

Note To remove the specified area from the router configuration, use the **no area** *area-id* command. The **no area** *area-id* command removes the area and all area options, such as **authentication** , **default-cost** , **nssa** , **range** , **stub** , **virtual-link** , and **interface.**

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure area 0 and HundredGigE interface 0/2/0/0. HundredGigE interface 0/2/0/0 is bound to area 0 automatically.

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/2/0/0
```

authentication (OSPF)

To enable plain text, Message Digest 5 (MD5) authentication, or null authentication for an Open Shortest Path First (OSPF) interface, use the **authentication** command in the appropriate mode. To remove such authentication, use the **no** form of this command.

authentication [**message-digest** [**keychain** *keychain*] | **null**]
no authentication

Syntax Description

message-digest	(Optional) Specifies that MD5 is used.
keychain *keychain*	(Optional) Specifies a keychain name.
null	(Optional) Specifies that no authentication is used. Useful for overriding password or MD5 authentication if configured for an area.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the authentication parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the authentication parameter specified for the process.

If this command is not specified at any level, then the interface does not use authentication.

If no keyword is specified, plain text authentication is used.

Command Modes

Interface configuration

Area configuration

Router configuration

Virtual-link configuration

VRF configuration

Multi-area interface configuration

Sham-link configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **authentication** command to specify an authentication type for the interface, which overrides the authentication specified for the area to which this interface belongs. If this command is not included in the configuration file, the authentication configured in the area to which the interface belongs is assumed (as specified by the area **authentication** command).

The authentication type and password must be the same for all OSPF interfaces that are to communicate with each other through OSPF. If you specified plain text authentication, use the **authentication-key** command to specify the plain text password.

If you enable MD5 authentication with the **message-digest** keyword, you must configure a key with the **message-digest-key** interface command.

To manage the rollover of keys and enhance MD5 authentication for OSPF, you can configure a container of keys called a keychain with each key comprising the following attributes: generate/accept time, key identification, and authentication algorithm. The keychain management feature is always enabled.

Note Changes to the system clock will impact the validity of the keys in the existing configuration.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set authentication for areas 0 and 1 of OSPF routing process 201. Authentication keys are also provided.

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 201
RP/0/RP0/CPU0:router(config-ospf)# router-id 10.1.1.1
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# authentication
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/1/0/1
RP/0/RP0/CPU0:router(config-ospf-ar-if)# authentication-key mykey
RP/0/RP0/CPU0:router(config-ospf-ar-if)# exit
RP/0/RP0/CPU0:router(config-ospf)# area 1
RP/0/RP0/CPU0:router(config-ospf-ar)# authentication
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/1/0/0
RP/0/RP0/CPU0:router(config-ospf-ar-if)# authentication-key mykey1
```

The following example shows how to configure use of an authentication keychain:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 201
RP/0/RP0/CPU0:router(config-ospf)# router-id 10.1.1.1
RP/0/RP0/CPU0:router(config-ospf)# authentication message-digest keychain mykeychain
```

authentication-key (OSPF)

To assign a password to be used by neighboring routers that are using the Open Shortest Path First (OSPF) simple password authentication, use the **authentication-key** command in the appropriate mode. To remove a previously assigned OSPF password, use the **no** form of this command.

authentication-key [**clear** | **encrypted**] *password*
no authentication-key

Syntax Description

clear	(Optional) Specifies that the key be clear text.
encrypted	(Optional) Specifies that the key be encrypted using a two-way algorithm.
password	Any contiguous string up to 8 characters in length that can be entered from the keyboard. For example, *mypswd2*.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the OSPF password parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the OSPF password parameter specified for the process.

If this command is not specified at any level, then no password is specified.

Clear is the default if the **clear** or **encrypted** keyword is not specified.

Command Modes

Interface configuration

Area configuration

Router configuration

Virtual-link configuration

VRF configuration

Multi-area configuration

Sham-link configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The password created by this command is inserted directly into the OSPF header when the Cisco IOS XR software originates routing protocol packets. A separate password can be assigned to each network on an individual interface basis. All neighboring routers on the same network must have the same password to be able to exchange OSPF information.

The **authentication-key** command must be used with the **authentication** command. If the **authentication** command is not configured, the password provided by the **authentication-key** command is ignored and no authentication is adopted by the OSPF interface.

Note The **authentication-key** command cannot be used with the **authentication** command when the **message-digest** or **null** keyword is configured.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure an authentication password as the string yourpass:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 201
RP/0/RP0/CPU0:router(config-ospf)# authentication-key yourpass
```

auto-cost (OSPF)

To control how the Open Shortest Path First (OSPF) protocol calculates default metrics for the interface, use the **auto-cost** command in the appropriate mode. To revert to the default reference bandwidth, use the **no** form of this command.

auto-cost {**reference-bandwidth** *mbps* | **disable**}
no auto-cost {**reference-bandwidth** | **disable**}

Syntax Description

reference-bandwidth	*mbps*	Specifies a rate in Mbps (bandwidth). Range is 1 to 4294967.
disable		Assigns a cost based on interface type.

Command Default

mbps : 100 Mbps

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

By default OSPF calculates the OSPF metric for an interface according to the bandwidth of the interface.

The OSPF metric is calculated as the *mbps* value divided by bandwidth, with *mbps* equal to 108 by default.

If you have multiple links with high bandwidth (such as OC-192), you might want to use a larger number to differentiate the cost on those links. That is, the metric calculated using the default *mbps* value is the same for all high-bandwidth links.

Recommended usage of cost configuration for OSPF interfaces with high bandwidth is to be consistent: Either explicitly configure (by using the **cost** command) or choose the default (by using the **auto-cost** command).

The value set by the **cost** command overrides the cost resulting from the **auto-cost** command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the reference value for the auto cost calculation to 1000 Mbps:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# auto-cost reference-bandwidth 1000
```

capability opaque disable

To prevent Multiprotocol Label Switching traffic engineering (MPLS TE) topology information flooded to the network through opaque LSAs, use the **capability opaque disable** command in the appropriate mode. To restore MPLS TE topology information flooded through opaque LSAs to the network, use the **no** form of the command.

capability opaque disable
no capability opaque disable

Command Default

Opaque LSAs are allowed.

Command Modes

XR Config mode

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The **capability opaque disable** command prevents flooded MPLS TE information (Types 1 and 4) through opaque LSAs of all scope (Types 9, 10, and 11).

Control opaque LSA support capability must be enabled for OSPF to support MPLS TE.

The MPLS TE topology information is flooded to the area through opaque LSAs by default.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to prevent OSPF from supporting opaque services:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# capability opaque disable
```

clear ospf process

To reset an Open Shortest Path First (OSPF) router process without stopping and restarting it, use the **clear ospf process** command in XR EXEC mode.

clear ospf [*process-name* [**vrf** {*vrf-name* | **all**}]] **process**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPF processes are reset.
vrf	(Optional) An OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF instance to be reset.
all	(Optional) Resets all OSPF VRF instances.

Command Default

No default behavior or value

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

When the OSPF router process is reset, OSPF releases all resources allocated, cleans up the internal database, and shuts down and restarts all interfaces that belong to the process.

Note The **clear ospf process** command may change the router ID unless the OSPF router ID is explicitly configured through the router-id (OSPF), on page 116 command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to reset all OSPF processes:

```
RP/0/RP0/CPU0:router# clear ospf process
```

The following example shows how to reset the OSPF 1 process:

```
RP/0/RP0/CPU0:router# clear ospf 1 process
```

clear ospf redistribution

To clear all routes redistributed from other protocols out of the Open Shortest Path First (OSPF) routing table, use the **clear ospf redistribution** command in XR EXEC mode.

clear ospf [*process-name* [**vrf** {*vrf-name* | **all**}]] **redistribution**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPF routes are cleared.
vrf	(Optional) OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF instance to be reset.
all	(Optional) Resets all OSPF VRF instances.

Command Default

No default behavior or value

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **clear ospf redistribution** command to cause the routing table to be read again. OSPF regenerates and sends Type 5 and Type 7 link-state advertisements (LSAs) to its neighbors. If an unexpected route has appeared in the OSPF redistribution, using this command corrects the issue.

Note Use of this command can cause a significant number of LSAs to flood the network. We recommend that you use this command with caution.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to clear all redistributed routes across all processes from other protocols:

```
RP/0/RP0/CPU0:router# clear ospf redistribution
```

clear ospf routes

To clear all Open Shortest Path First (OSPF) routes from the OSPF routing table, use the **clear ospf routes** command in XR EXEC mode.

clear ospf [*process-name* [**vrf** {*vrf-name* | **all**}]] **routes**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPF routes are cleared.
vrf	(Optional) OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF instance to be reset.
all	(Optional) Resets all OSPF VRF instances.

Command Default

No default behavior or value

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to clear all OSPF routes from the OSPF routing table and recompute valid routes. When the OSPF routing table is cleared, OSPF routes in the global routing table are also recalculated.

```
RP/0/RP0/CPU0:router# clear ospf routes
```

clear ospf statistics

To clear the Open Shortest Path First (OSPF) statistics of neighbor state transitions, use the **clear ospf statistics** command in XR EXEC mode.

clear ospf [*process-name* [**vrf** {*vrf-name* | **all**}]] **statistics** [**neighbor** [*type interface-path-id*] [*ip-address*]]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPF statistics of neighbor state transitions are cleared.
vrf	(Optional) OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF instance to be reset.
all	(Optional) Resets all OSPF VRF instances.
neighbor	(Optional) Clears the state transition counters of the specified neighbor only.
type	(Optional) Interface type. For more information, use the question mark (?) online help function.
interface-path-id	(Optional) Physical interface or virtual interface. Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.
ip-address	(Optional) IP address of a specified neighbor for whom you want to clear the state transition counter.

Command Default

No default behavior or value

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **clear ospf statistics** command to reset OSPF counters. Reset is useful to detect changes in counter values.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to reset the OSPF transition state counters for all neighbors on Packet-over-SONET/SDH (POS) interface 0/2/0/0:

```
RP/0/RP0/CPU0:router# clear ospf statistics neighbor POS 0/2/0/0
```

cost (OSPF)

To explicitly specify the interface (network) for Open Shortest Path First (OSPF) path calculation, use the **cost** command in the appropriate mode. To remove the cost, use the **no** form of this command.

cost *cost*
no cost

Syntax Description

cost	Unsigned integer value expressed as the link-state metric. Range is 1 to 65535.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the cost parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the cost parameter specified for the process.

If this command is not specified at any level, then the cost is calculated by the **auto-cost** command.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Multi-area configuration

Sham-link configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The link-state metric is advertised as the link cost in the router link advertisement. Cisco IOS XR software does not support type of service (ToS), so you can assign only one cost for each interface.

In general, the path cost is calculated using the following formula:

108 / bandwidth (the default auto cost is set to 100 Mbps)

This calculation is the default reference bandwidth used by the auto-costing calculation which establishes the interface auto-cost The **auto-cost** command can set this reference bandwidth to some other value. The **cost** command is used to override the auto-costing calculated default value for interfaces.

Using this formula, the default path cost is 1 for any interface that has a link bandwidth of 100 Mbps or higher. If this value does not suit the network, configure the reference bandwidth for auto calculating costs based on the link bandwidth.

The value set by the **cost** command overrides the cost resulting from the **auto-cost (OSPF)** command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the cost value to 65 for HundredGigE interface 0/1/0/1:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/1/0/1
RP/0/RP0/CPU0:router(config-ospf-ar-if)# cost 65
```

cost-fallback (OSPF)

To apply higher cost than the normal interface cost when the cumulative bandwidth of a bundle interface goes below the threshold specified and to revert to the original cost if the cumulative bandwidth goes above the configured threshold, use the **cost-fallback** command. To remove the cost-fallback, use the **no** form of this command.

cost-fallback cost threshold bandwidth
no cost-fallback

Syntax Description

cost **threshold**	Unsigned integer value expressed as the link-state metric. Range is 1 to 65535, but typically, cost-fallback value is supposed to be set to a value higher than the normal cost.
bandwidth	Unsigned integer value expressed in Mbits per second. Range is 1 to 4294967.

Command Default

If this command is not specified in interface configuration mode, the currently effective interface cost takes effect even when the cumulative bandwidth goes down below the maximum bandwidth. Unlike the interface cost command, this cost-fallback command is available only under interface configuration mode; it is not available in area or process level. Unlike other interface specific parameters, no inheritance will take place from area or process level if this command is not specified at interface level.

Command Modes

Interface configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The fallback cost must be set to a higher value than the normal interface cost. The motivation of setting the fallback cost is to cost out an interface or disfavor an interface without shutting it down when its cumulative bandwidth goes below the user specified threshold, so that the traffic can take an alternative path. The normal interface cost will take over when the cumulative bandwidth reaches or exceeds user-specified threshold.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the cost-fallback value:

The following example shows how to set the cost-fallback value for Bundle-Ether:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 100
RP/0/RP0/CPU0:router(config-ospf)# router-id 2.2.2.2
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface bundle-Ether
RP/0/RP0/CPU0:router(config-ospf-ar-if)# cost-fallback 1000 threshold 300
```

database-filter all out (OSPF)

To filter outgoing link-state advertisements (LSAs) to an Open Shortest Path First (OSPF) interface, use the **database-filter all out** command in the appropriate mode. To restore the forwarding of LSAs to the interface, use the **disable** form of the command.

database-filter all out [**disable** | **enable**]

Syntax Description

disable	(Optional) Disables filtering.
enable	(Optional) Enables filtering.

Command Default

The database filter is disabled.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Multi-area configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Use the **database-file all out** command to perform the same function that the neighbor database-filter all out, on page 83 command performs on a neighbor basis.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to prevent flooding of OSPF LSAs to broadcast, nonbroadcast, and point-to-point networks reachable through HundredGigE interface 0/1/0/1:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/1/0/1
RP/0/RP0/CPU0:router(config-ospf-ar-if)# database-filter all out
```

dead-interval (OSPF)

To set the interval after which a neighbor is declared dead when no hello packets are observed, use the **dead-interval** command in the appropriate mode. To return to the default time, use the **no** form of this command.

dead-interval *seconds*
no dead-interval

Syntax Description

seconds	Integer that specifies the interval (in seconds). Range is 1 to 65535. The value must be the same for all nodes on the network.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the dead interval parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the dead interval parameter specified for the process.

If this command is not specified at any level, then the dead interval is four times the interval set by the **hello-interval** (OSPF) command.

Command Modes

Interface configuration

Area configuration

Router configuration

Virtual-link configuration

VRF configuration

Multi-area configuration

Sham-link configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The dead interval value must be the same for all routers and access servers on a specific network.

If the hello interval is configured, the dead interval value must be larger than the hello interval value. The dead interval value is usually configured four times larger than the hello interval value.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the OSPF dead interval to 40 seconds:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/1/0/1
RP/0/RP0/CPU0:router(config-ospf-ar-if)# dead-interval 40
```

default-cost (OSPF)

To specify a cost for the default summary route sent into a stub area or not-so-stubby area (NSSA), use the **default-cost** command in area configuration mode. To remove the assigned default route cost, use the **no** form of this command.

default-cost *cost*
no default-cost *cost*

Syntax Description

cost	Cost for the default summary route used for a stub or NSSA area. The acceptable value is a 24-bit number.

Command Default

cost : 1

Command Modes

Area configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **default-cost** command only on an Area Border Router (ABR) attached to a stub or an NSSA area.

In all routers and access servers attached to the stub area, the area should be configured as a stub area using the **stub** command in the area submode. Use the **default-cost** command only on an ABR attached to the stub area. The **default-cost** command provides the metric for the summary default route generated by the ABR into the stub area.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to assign a default cost of 20 to a stub area. The HundredGigE interface 0/4/0/3 is also configured in the stub area):

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 201
RP/0/RP0/CPU0:router(config-ospf)# area 10.15.0.0
RP/0/RP0/CPU0:router(config-ospf-ar)# stub
RP/0/RP0/CPU0:router(config-ospf-ar)# default-cost 20
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/4/0/3
```

default-information originate (OSPF)

To generate a default external route into an Open Shortest Path First (OSPF) routing domain, use the **default-information originate** command in the appropriate mode. To disable this feature, use the **no** form of this command.

default-information originate [**always**] [**metric** *metric-value*] [**metric-type** *type-value*] [**route-policy** *policy-name*]
no default-information originate

Syntax Description

always	(Optional) Always advertises the default route regardless of whether the routing table has a default route.
metric *metric-value*	(Optional) Specifies the metric used for generating the default route. The default metric value is 1. Range is 1 to 16777214.
metric-type *type-value*	(Optional) Specifies the external link type associated with the default route advertised into the OSPF routing domain. It can be one of the following values: **1**—Type 1 external route **2**—Type 2 external route
route-policy *policy-name*	(Optional) Specifies that a routing policy be used and the routing policy name.

Command Default

When you do not use this command in router configuration mode, no default external route is generated into an OSPF routing domain.

metric-value : 1

type-value : 2

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Whenever you use the **redistribute** or **default-information originate** command to redistribute routes into an OSPF routing domain, the software automatically becomes an Autonomous System Boundary Router (ASBR). However, an ASBR does not, by default, generate a default route into the OSPF routing domain. The software still must have a default route for itself before it generates one, except when you have specified the **always** keyword.

The **default-information originate** route-policy attach point conditionally injects the default route 0.0.0.0/0 into the OSPF link-state database, and is done by evaluating the attached policy. If any routes specified in the policy exist in the global RIB, then the default route is inserted into the link-state database. If there is no match condition specified in the policy, the policy passes and the default route is generated into the link-state database.

For information about the default-information originate attach point, see the *OSPF Policy Attach Points* section in the *Implementing Routing Policy* chapter in *Routing Configuration Guide for Cisco NCS 5500 Series RoutersRouting Configuration Guide for Cisco NCS 540 Series RoutersRouting Configuration Guide*.

For information about routing policies, see the *Routing Policy Commands* chapter in the *Routing Command Reference Guide*.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to specify a metric of 100 for the default route redistributed into the OSPF routing domain and an external metric type of Type 1:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router ospf 109
RP/0/RP0/CPU0:router(config-ospf)#redistribute igrp 108 metric 100
RP/0/RP0/CPU0:router(config-ospf)#default-information originate metric 100 metric-type 1
```

default-metric (OSPF)

To set default metric values for routes redistributed from another protocol into the Open Shortest Path First (OSPF) protocol, use the **default-metric** command in the appropriate mode. To return to the default state, use the **no** form of this command.

default-metric *value*
no default-metric *value*

Syntax Description

value	Default metric value appropriate for the specified routing protocol. Range is 1 to 16777214.

Command Default

Built-in, automatic metric translations, as appropriate for each routing protocol.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **default-metric** command with the **redistribute** command to cause the current routing protocol to use the same metric value for all redistributed routes. A default metric helps solve the problem of redistributing routes with incompatible metrics. Whenever metrics do not convert, use a default metric to provide a reasonable substitute and enable the redistribution to proceed.

The default-metric value configured in OSPF configuration does not apply to connected routes that are redistributed to OSPF using the **redistribute connected** command. To set a non-default metric for connected routes, configure OSPF with the **redistribute connected metric** *metric-value* command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to advertise Intermediate System-to-Intermediate System (IS-IS) protocol-derived routes into OSPF and assign a metric of 10:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# default-metric 10
RP/0/RP0/CPU0:router(config-ospf)# redistribute isis IS-IS_isp
```

demand-circuit (OSPF)

To configure the Open Shortest Path First (OSPF) protocol to treat the interface as an OSPF demand circuit, use the **demand-circuit** command in the appropriate mode. To remove the demand circuit designation from the interface, use the **no** form of this command.

demand-circuit [**disable** | **enable**]
no demand-circuit

Syntax Description

disable	(Optional) Disables the interface as an OSPF demand circuit.
enable	(Optional) Enables the interface as an OSPF demand circuit.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the demand circuit parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the demand circuit parameter specified for the process.

If this command is not specified at any level, then the circuit is not a demand circuit.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

On point-to-point interfaces, only one end of the demand circuit must be configured with this command. Periodic hello messages are suppressed and periodic refreshes of link-state advertisements (LSAs) do not flood the demand circuit. Use the **demand-circuit** command to allow the underlying data link layer to be closed when the topology is stable. In point-to-multipoint topology, only the multipoint end must be configured with this command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the configuration for an OSPF demand circuit:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
```

```
RP/0/RP0/CPU0:router(config-ospf)# demand-circuit
```

disable-dn-bit-check

To specify that down bits should be ignored, use the **disable-dn-bit-check** command in VPN routing and forwarding (VRF) configuration mode. To specify that down bits should be considered, use the **no** form of this command.

disable-dn-bit-check
no disable-dn-bit-check

Command Default

Down bits are considered.

Command Modes

VRF configuration mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to specify that down bits be ignored:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# vrf v1
RP/0/RP0/CPU0:router(config-ospf-vrf)# disable-dn-bit-check
```

distance (OSPF)

To define an administrative distance, use the **distance** command in an appropriate configuration mode. To remove the **distance** command from the configuration file and restore the system to its default condition in which the software removes a distance definition, use the **no** form of this command.

distance *weight* [*ip-address wildcard-mask* [*access-list-name*]]
no distance *weight ip-address wildcard-mask* [*access-list-name*]

Syntax Description

weight	Administrative distance. Range is 10 to 255. Used alone, the *weight* argument specifies a default administrative distance that the software uses when no other specification exists for a routing information source. Routes with a distance of 255 are not installed in the routing table. lists the default administrative distances. Table 1: Default Administrative Distances, on page 37
ip-address	(Optional) IP address in four-part, dotted-decimal notation.
wildcard-mask	(Optional) Wildcard mask in four-part, dotted decimal format. A bit set to 1 in the *mask* argument instructs the software to ignore the corresponding bit in the address value.
access-list-name	(Optional) Name of an IP access list to be applied to incoming routing updates.

Command Default

If this command is not specified, then the administrative distance is the default, as specified in Table 1: Default Administrative Distances, on page 37.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

An administrative distance is an integer from 10 to 255. In general, the higher the value, the lower the trust rating. An administrative distance of 255 means that the routing information source cannot be trusted at all and should be ignored. Weight values are subjective; no quantitative method exists for choosing weight values.

If an access list is used with this command, it is applied when a network is being inserted into the routing table. This behavior allows you to filter networks based on the IP prefix supplying the routing information. For example, you could filter possibly incorrect routing information from networking devices not under your administrative control.

The order in which you enter **distance** commands can affect the assigned administrative distances in unexpected ways (see the “Examples” section for further clarification).

This table lists default administrative distances.

Table 1: Default Administrative Distances

Route Source	Default Distance
Connected interface	0
Static route out on interface	0
State route to next-hop	1
External BGP	20
OSPF	110
IS-IS	115
RIP version 1 and 2	120
Internal BGP	200
Unknown	255

Task ID

Task ID	Operations
ospf	read, write

Examples

In the following example, the **router ospf** command sets up OSPF routing instance1. The first **distance** command sets the default administrative distance to 255, which instructs the software to ignore all routing updates from networking devices for which an explicit distance has not been set. The second **distance** command sets the administrative distance for all networking devices on the Class C network 192.168.40.0 0.0.0.255 to 90.

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# distance 255
RP/0/RP0/CPU0:router(config-ospf)# distance 90 192.168.40.0 0.0.0.255
```

distance ospf

To define Open Shortest Path First (OSPF) route administrative distances based on route type, use the **distance ospf** command in router configuration mode. To restore the default value, use the **no** form of this command.

distance ospf {**intra-area** | **inter-area** | **external**} *distance*
no distance ospf

Syntax Description

intra-area | **inter-area** | **external**	Sets the type of area. It can be one of the following values: **intra-area** —All routes within an area. **inter-area** —All routes from one area to another area. **external** —All routes from other routing domains, learned by redistribution. Any combination of the above areas is allowed.
distance	Route administrative distance.

Command Default

distance : 110

Command Modes

Router configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

You must specify one of the keywords.

Use the **distance ospf** command to perform the same function as the **distance** command used with an access list. However, the **distance ospf** command sets a distance for an entire group of routes, rather than a specific route that passes an access list.

A common reason to use the **distance ospf** command is when you have multiple OSPF processes with mutual redistribution, and you want to prefer internal routes from one over external routes from the other.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to change the external distance to 200, making the route less reliable:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# redistribute ospf 2
RP/0/RP0/CPU0:router(config-ospf)# distance ospf external 200
RP/0/RP0/CPU0:router(config-ospf)# exit
```

```
RP/0/RP0/CPU0:router(config)# router ospf 2
RP/0/RP0/CPU0:router(config-ospf)# redistribute ospf 1
RP/0/RP0/CPU0:router(config-ospf)# distance ospf external 200
```

distribute-list

To filter networks received or transmitted in Open Shortest Path First (OSPF) updates, use the **distribute-list** command in the appropriate mode. To change or cancel the filter, use the **no** form of this command.

distribute-list {*access-list-name* {**in** | **out** [**bgp** *number* | **connected** | **ospf** *instance* | **static**]} | **route-policy** *route-policy-name* **in**}
no distribute-list {*access-list-name* {**in** | **out**} | **route-policy** *route-policy-name* **in**}

Syntax Description

access-list-name	Standard IP access list name. The list defines which networks are to be received and which are to be suppressed in routing updates.
in	Applies the access list or route-policy to incoming routing updates.
out	Applies the access list to outgoing routing updates. The **out** keyword is available only in router configuration mode.
bgp	(Optional) Applies the access list to BGP routes.
connected	(Optional) Applies the access list to connected routes.
ospf	(Optional) Applies the access list to OSPF routes (not the current OSPF process).
static	(Optional) Applies the access list to statically configured routes.
route-policy *route-policy-name*	Specifies the route-policy to filter OSPF prefixes.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the distribute list parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the distribute list parameter specified for the process.

If this command is not specified at any level, then the distribute list is disabled.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Multi-area configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **distribute-list** command to limit which OSPF routes are installed on this router. The **distribute-list** command does not affect the OSPF protocol itself.

The **distribute-list in** is configurable at instance (process), area, and interface levels. Regular OSPF configuration inheritance applies. Configuration is inherited from instance > area > interface levels.

Use the **route-policy** *route-policy-name* keyword and argument to allow use of route policies to filter OSPF prefixes.

Note Either an access-list, or a route-policy can be used in a single command, not both. Configuring the command with access-list removes the route-policy configuration, and vice versa.

The "if tag..." statements can be used in **distribute-list in** *route-policy*. The matching on route tag supports operators "eq/ge/is/le". Operator "in" is not supported.

Note When distribute-list ACL or route-policy is modified, OSPF removes all routes from the RIB, and repopulates the RIB. Due to an expected delay in the repopulation of the RIB, there could be traffic disruption.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to prevent OSPF routes from the 172.17.10.0 network from being installed if they are learned in area 0:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# ipv4 access-list 3
RP/0/RP0/CPU0:router(config-ipv4-acl)# deny 172.17.10.0 0.0.0.255
RP/0/RP0/CPU0:router(config-ipv4-acl)# permit any any
!
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# distribute-list 3 in
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/1/0/3
```

domain-id (OSPF)

To specify the Open Shortest Path First (OSPF) VPN routing and forwarding (VRF) domain ID, use the **domain-id** command in VRF configuration mode. To remove an OSPF VRF domain ID, use the **no** form of this command.

domain-id [**secondary**] **type** [**0005** | **0105** | **0205** | **8005**] **value** *value*
no domain-id [**secondary**] **type** [**0005** | **0105** | **0205** | **8005**] **value** *value*

Syntax Description

secondary		(Optional) OSPF secondary domain ID.
type		Primary OSPF domain ID in hex format.
value	*value*	OSPF domain ID value in hex format (six octets).

Command Default

No domain ID is specified.

Command Modes

VRF configuration mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

An OSPF domain id must be explicitly configured. The OSPF domain ID helps OSPF determine how to translate a prefix received through Border Gateway Protocol (BGP) from the remote provider edge (PE). If the domain IDs match, OSPF generates a Type 3 link state advertisement (LSA). If the domain IDs do not match, OSPF generates a Type 5 LSA.

There is only one primary domain ID. There can be multiple secondary domain IDs.

Note When an IOS XR router and an IOS router are configured as peers, the two Domain IDs must match. Manually configure the IOS XR Domain ID value to match the IOS default Domain ID value. This ensures that the routes have route code "OIA" because they are learned as inter-area routes. If the Domain IDs do not match, the routes have route code, "O-E2" because they are learned as external routes. Use the **show ip ospf** command to get the OSPF Domain ID from the IOS router. Then, set the IOS XR Domain ID to the same value using the **domain-id** command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to specify a domain ID:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf o1
RP/0/RP0/CPU0:router(config-ospf)# vrf v1
RP/0/RP0/CPU0:router(config-ospf-vrf)# domain-id type 0105 value AABBCCDDEEFF
```

domain-tag

To specify the Open Shortest Path First (OSPF) VPN routing and forwarding (VRF) domain tag, use the **domain-tag** command in VRF configuration mode. To remove an OSPF VRF domain tag, use the **no** form of this command.

domain-tag *tag*
no domain-tag

Syntax Description

tag	OSPF domain tag as a 32-bit value. The valid range is 0 to 4294967295.

Command Default

No OSPF VRF domain tag is specified.

Command Modes

VRF configuration mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The domain tag is added to any Type 5 link state advertisements (LSAs) generated as a result of VPN-IP routes received from Border Gateway Protocol (BGP). The domain-tag is derived from BGP autonomous system number (ASN).

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to specify the domain tag:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf o1
RP/0/RP0/CPU0:router(config-ospf)# domain-tag 234
```

exchange-timer

To set a time limit to prevent an indefinite EXCHANGE or LOADING phase during adjacency formation, use the **exchange-timer** command in router configuration mode.

exchange-timer *exchange-time-in-minutes* **hold-time** *hold-time-in-minutes* **recovery-count** *recovery-count*

Syntax Description

exchange-time-in-minutes	Duration in minutes, after which the OSPF interface state is brought down if adjacencies remain in the exchange/loading phase on that interface. The range is from 1 to 35791394. The default value is 10.
hold-time	Specifies the option to set the time after which a recovery attempt will be made.
hold-time-in-minutes	Duration in minutes until which the OSPF interface state is kept in the down state before attempting recovery once it was made down due to exchange-timeout. The range is from 1 to 35791394. The default value is 5 minutes.
recovery-count	Specifies the number of attempts the router makes to bring up the adjacency toward the ASBR before keeping the adjacency shut indefinitely.
recovery-count	The maximum number of retries that a router undertakes to restore the OSPF state of an interface before it is considered to be permanently down. Once this limit is surpassed, the OSPF state of the interface remains down until it is manually reset, as outlined in the subsequent recovery section. The recovery count value resets to the default value of 300 when the interface adjacency reaches the FULL state with the neighbor. The range is from 1 to 4294967294. The default value is 300.

Command Default

The exchange timer operates using its default values.

Command Modes

Router Configuration

Command History

Release	Modification
Release 7.10.1	This command was introduced.

Usage Guidelines

- To adjust the default parameters of the exchange timer, you can use the **exchange-timer** command. We recommend configuring this command with values that suit your specific scale requirements.

  By default, the exchange timer is enabled with the following default values:

  - Default Exchange Time: 10 minutes

- Default Hold Time: 5 minutes
- Default Recovery Count: 300

These default values have been chosen to ensure that the interface remains operational for an extended period before being permanently brought down. The total duration for attempting recovery is calculated as follows:

```
(Exchange Time + Hold Time) * Recovery Count = 3.125 days
```

- To recover a permanently disabled interface:

Option 1:

1. Shut down the permanently disabled interface (**shutdown** command).
2. Bring up the interface (**no shutdown** command).

Option 2:

1. Remove the interface from OSPF.
2. Add the interface again to OSPF.

Task ID

Task ID	Operation
OSPF	read, write

Example

The following example shows how to set a time limit to prevent an indefinite EXCHANGE or LOADING phase during adjacency formation.

```
Router(config)# router ospf 1
Router(config-ospf)# exchange-timer 357913 hold-time 35791 recovery-count 4294967294
```

fast-reroute (OSPFv2)

To enable IP fast reroute loop-free alternate (LFA) computation, use the **fast-reroute** command in the appropriate OSPF configuration mode. To disable the IP fast reroute loop-free alternate computation, use the **no** form of this command.

To disable loop-free alternate computation that is enabled on a higher level, use the **fast-reroute** command with **disable** keyword.

fast-reroute {**per-link** | **per-prefix**} [**disable**]
no fast-reroute

Syntax Description

per-link	Enables per-link loop-free alternate computation.
per-prefix	Enables per-prefix loop-free alternate computation.
disable	(Optional) Disables loop-free alternate computation that was enabled on a higher level.

Command Default

IP fast-reroute LFA computation is disabled.

Command Modes

Area configuration

Interface configuration

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Only one mode of computation can be configured on an interface - per-link or per-prefix. Different modes of computations can be enabled on different interfaces; one set of interface using per-link and other set using per-prefix computation. Based on the outgoing interface of the primary path, per-link or per-prefix backup path will be computed.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to enable per-link computation of loop-free alternates under interface POS 0/3/0/0:

```
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/3/0/0
```

```
RP/0/RP0/CPU0:router(config-ospf-ar-if)# fast-reroute per-link
```

This example shows how to enable per-prefix computation of loop-free alternates under area 0:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router ospf 1
RP/0/RP0/CPU0:router(config-ospf)#area 0
RP/0/RP0/CPU0:router(config-ospf-ar)#fast-reroute per-prefix
```

This example shows how to disable computation of loop-free alternates that was configured under area 0:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router ospf 1
RP/0/RP0/CPU0:router(config-ospf)#area 0
RP/0/RP0/CPU0:router(config-ospf-ar)#fast-reroute per-prefix
RP/0/RP0/CPU0:router(config-ospf-ar)#interface HundredGigE 0/3/0/0
RP/0/RP0/CPU0:router(config-ospf-ar-if)#fast-reroute disable
```

fast-reroute per-link exclude interface

To excludes specified interface to be used as a backup during (IPFRR) loop-free alternate (LFA) computation, use the **fast-reroute per-link exclude interface** command, in the appropriate OSPF configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-link exclude interface *type interface-path-id*
no fast-reroute per-link exclude interface *type interface-path-id*

Syntax Description

type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.

Command Default

No interfaces are excluded.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to exclude an interface from IP fast reroute loop-free alternate (LFA) computation:

```
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf-ar-if)# fast-reroute per-link exclude interface HundredGigE
 0/3/2/1
```

fast-reroute per-prefix exclude interface (OSPFv2)

To exclude interface to be used as a backup path from fast-reroute loop-free alternate per-prefix computation, use the **fast-reroute per-prefix exclude interface** command in the appropriate OSPF configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-prefix exclude interface *type interface-path-id*
no fast-reroute per-prefix exclude interface *type interface-path-id*

Syntax Description

type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.

Command Default

No interfaces are excluded.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Backup paths via the excluded interfaces will not be computed.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to exclude interface POS0/6/0/1 from being used as a backup path:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router ospf 100
RP/0/RP0/CPU0:router(config-ospf)#fast-reroute per-prefix exclude interface HundredGigE
0/6/0/1
```

fast-reroute per-prefix use-candidate-only (OSPFv2)

To restrict the backup interfaces to those that are present on the LFA candidate list, use the **fast-reroute per-prefix use-candidate-only** command in router OSPF configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-prefix use-candidate-only [**enable** | **disable**]
fast-reroute per-prefix use-candidate-only

Syntax Description

enable	Enables backup selection from candidate-list only.
disable	Disables backup selection from candidate-list only.

Command Default

Disabled.

Command Modes

Router OSPF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to restrict the backup interfaces to those that are present on the LFA candidate list:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router ospf 100
RP/0/RP0/CPU0:router(config-ospf)#fast-reroute per-prefix use-candidate-only
```

flood-reduction (OSPF)

To suppress the unnecessary flooding of link-state advertisements (LSAs) in stable topologies, use the **flood-reduction** command in the appropriate mode. To remove this functionality from the configuration, use the **no** form of this command.

flood-reduction [**enable** | **disable**]
no flood-reduction [**enable** | **disable**]

Syntax Description

enable	(Optional) Turns on this functionality at a specific level.
disable	(Optional) Turns off this functionality at a specific level.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the flood reduction parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the flood reduction parameter specified for the process.

If this command is not specified at any level, then flood reduction is disabled.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

All routers supporting the OSPF demand circuit are compatible and can interact with routers supporting flooding reduction.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to reduce the flooding of unnecessary LSAs for area 0:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/1/0/3
RP/0/RP0/CPU0:router(config-ospf-ar-if)# flood-reduction
```

hello-interval (OSPF)

To specify the interval between consecutive hello packets that are sent on the Open Shortest Path First (OSPF) interface, use the **hello-interval** command in the appropriate mode. To return to the default time, use the **no** form of this command.

hello-interval *seconds*
no hello-interval

Syntax Description

seconds	Interval (in seconds). The value must be the same for all nodes on a specific network. Range is 1 to 65535.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the hello interval parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the hello interval parameter specified for the process.

If this command is not specified at any level, then the hello interval is 10 seconds (broadcast) or 30 seconds (non-broadcast).

Command Modes

Interface configuration

Area configuration

Router configuration

Virtual-link configuration

VRF configuration

Multi-area configuration

Sham-link configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The hello interval value is advertised in the hello packets. The shorter the hello interval, the faster topological changes are detected, but more routing traffic occurs. This value must be the same for all routers and access servers on a specific network.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the interval between hello packets to 15 seconds:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/1/0/1
RP/0/RP0/CPU0:router(config-ospf-ar-if)# hello-interval 15
```

ignore lsa mospf

To suppress the sending of syslog messages when the router receives link-state advertisement (LSA) Type 6 multicast Open Shortest Path First (MOSPF) packets, which are unsupported, use the **ignore lsa mospf** command in an appropriate configuration mode. To restore the sending of syslog messages, use the **no** form of this command.

ignore lsa mospf
no ignore lsa mospf

This command has no keywords or arguments.

Command Default

When you do not specify this command in router configuration mode, each MOSPF packet received by the router causes the router to send a syslog message.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Cisco routers do not support LSA Type 6 (MOSPF), and they generate syslog messages if they receive such packets. If the router is receiving many MOSPF packets, you might want to configure the router to ignore the packets and thus prevent a large number of syslog messages.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure the router to suppress the sending of syslog messages when it receives MOSPF packets:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# ignore lsa mospf
```

interface (OSPF)

To define the interfaces on which the Open Shortest Path First (OSPF) protocol runs, use the **interface** command in area configuration mode. To disable OSPF routing for interfaces, use the **interface** form of this command.

interface *type interface-path-id*
no interface *type interface-path-id*

Syntax Description

type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtualinterface. **Note** se the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.

Command Default

When you do not specify this command in configuration mode, OSPF routing for interfaces is not enabled.

Command Modes

Area configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **interface** command to associate a specific interface with an area. The interface remains associated with the area even when the IP address of the interface changes.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how the OSPF routing process 109 defines four OSPF areas (0, 2, 3, and 10.9.50.0), and associates an interface with each area:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 4/0/0/3
!
RP/0/RP0/CPU0:router(config-ospf)# area 2
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/1/0/3
!
RP/0/RP0/CPU0:router(config-ospf)# area 3
```

```
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 3/0/0/2
!
RP/0/RP0/CPU0:router(config-ospf)# area 10.9.50.0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 3/0/0/1
```

log adjacency changes (OSPF)

To configure the router to send a syslog message when the state of an Open Shortest Path First (OSPF) neighbor changes, use the **log adjacency changes** command in XR Config mode. To turn off this function, use the **disable** keyword. To log all state changes, use the **detail** keyword.

log adjacency changes {**detail** | **disable**}

Syntax Description

detail	Provides all (DOWN, INIT, 2WAY, EXSTART, EXCHANGE, LOADING, FULL) adjacency state changes.
disable	Disables sending adjacency change messages.

Command Default

The router sends a syslog message when the state of an OSPF neighbor changes.

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **log adjacency changes** command to display high-level changes to the state of the peer relationship. Configure this command if you want to know about OSPF neighbor changes.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure the software to send a syslog message for any OSPF neighbor state changes:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
RP/0/RP0/CPU0:router(config-ospf)# log adjacency changes detail
```

loopback stub-network

To enable advertising loopback as stub networks, use the **loopback stub-network** command in an appropriate configuration mode. To disable advertising loopback as stubnetworks, use the **no** form of this command.

loopback stub-network [**enable** | **disable**]
no loopback stub-network

Syntax Description

enable	(Optional) Enables advertising loopbacks as stub networks.
disable	(Optional) Disables advertising loopbacks as stub networks.

Command Default

By default, OSPF advertises loopbacks as stub hosts.

Command Modes

OSPF interface configuration

OSPF router configuration

OSPF area configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

In the interface submode, the command can be enabled only on loopback interfaces.

Task ID

Task ID	Operation
ospf	read, write

Examples

The following example shows how to enable advertising loopback as a stub network, under OSPF interface configuration:

```
RP/0/RP0/CPU0:router(config)#router ospf 100
RP/0/RP0/CPU0:router(config-ospf)#loopback stub-network enable
```

max external-lsa

To specify the maximum number of external LSAs that the router will accept from a neighbor, use the **max-external-lsa** command in router configuration mode.

max-external-lsa *lsa-limit* *logging-threshold-percentage* [**suppress-neighbor** | **warning-only**]

Syntax Description

lsa-limit	Specifies the maximum number of external LSAs that a router will accept from a neighbor. If the neighbor sends more than the specified number of LSAs, the adjacency will be brought down and kept down for a specified interval. The range is from 1 to 4294967294. The default value is 50,000.
logging-threshold-percentage	Specifies the percentage of the total number of external LSAs from an ASBR at which a warning log is given. The default value is 75.
suppress-neighbor	Specifies that the link towards an ASBR be brought down if the ASBR is the immediate neighbor. This option is enabled by default.
warning-only	Specifies the option to overwrite the suppress-neighbor behavior. If the **warning-only** keyword is configured instead of **suppress-neighbor** keyword, the link to the ASBR is not brought down when the limit of external LSAs is reached. Only alert logs is given.

Command Default

If the command is not configured, the ASBR does not monitor unsuccessful adjacencies towards other ASBRs caused by the ASBR reaching the configured maximum limit for external Link State Advertisements (LSAs). Additionally, the ASBR does not incorporate a counter that, upon reaching the specified number of attempts to establish adjacency, causes the neighbor to cease further efforts, triggering permanent adjacency downtime.

Command Modes

Router Configuration

Command History

Release	Modification
Release 7.10.1	This command was introduced.

Usage Guidelines

We recommend configuring this command with values that suit your specific scale requirements.

Task ID

Task ID	Operation
OSPF	read, write

Example

The following example show how to configure the maximum number of external LSAs, the warning log percentage, and the link-down option for an OSPF router.

```
Router(config)# router ospf 1
Router(config-ospf)# area 0
Router(config-ospf-ar)# max-external-lsa 50000 suppress-neighbor
```

max-lsa

To limit the number of nonself-generated link-state advertisements (LSAs) that an Open Shortest Path First (OSPF) routing process can keep in the OSPF link-state database (LSDB), use the **max-lsa** command in XR Config mode. To remove the limit of non self-generated LSAs that an OSPF routing process can keep in the OSPF LSDB, use the **no** form of this command.

max-lsa *max* [*threshold*] [**warning-only**] [**ignore-time** *value*] [**ignore-count** *value*] [**reset-time** *value*]
no max-lsa *max* [*threshold*] [**warning-only**] [**ignore-time** *value*] [**ignore-count** *value*] [**reset-time** *value*]

Syntax Description

max	Maximum number of nonself-generated LSAs the OSPF process can keep in the OSPF LSBD. The default is 500000 LSAs.
threshold	(Optional) The percentage of the maximum LSA number, as specified by the maximum-number argument, at which a warning message is logged. The default is 75 percent.
warning-only	(Optional) Specifies that only a warning message is sent when the maximum limit for LSAs is exceeded. Disabled by default.
ignore-time *value*	(Optional) Specifies the time, in minutes, to ignore all neighbors after the maximum limit of LSAs has been exceeded. The default is 5 minutes.
ignore-count *value*	(Optional) Specifies the number of times the OSPF process can consecutively be placed into the ignore state. The default is 5 times.
reset-time *value*	(Optional) Specifies the time, in minutes, after which the ignore count is reset to zero. The default is 2 times **ignore-time** .

Command Default

Disabled

Command Modes

XR Config mode

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

This command allows you to protect the OSPF routing process from the large number of received LSAs that can result from a misconfiguration on another router in the OSPF domain (for example, the redistribution of a large number of IP prefixes to OSPF).

When this feature is enabled, the router keeps count of the number of all received (nonself-generated) LSAs. When the configured *threshold* value is reached, an error message is logged. When the configured *max* number of received LSAs is exceeded, the router stops accepting new LSAs.

If the count of received LSAs is higher than the configured *max* number after one minute, the OSPF process disables all adjacencies in the given context and clears the OSPF database. This state is called the ignore state. In this state, all OSPF packets received on all interfaces belonging to the OSPF instance are ignored and no OSPF packets are generated on its interfaces. The OSPF process remains in the ignore state for the duration of the configured **ignore-time** . When the **ignore-time** expires, the OSPF process returns to normal operation and starts building adjacencies on all its interfaces.

To prevent the OSPF instance from endlessly oscillating between its normal state and the ignore state, as a result of the LSA count immediately exceeding the *max* number again after it returns from the ignore state, the OSPF instance keeps a count of how many times it has been in the ignore state. This counter is called the **ignore-count** . If the **ignore-count** exceeds its configured value, the OSPF instance remains in the ignore state permanently.

To return the OSPF instance to its normal state, you must issue the **clear ip ospf** command. The **ignore-count** is reset to zero if the LSA count does not exceed the *max* number again during the time configured by the **reset-time** keyword.

If you use the **warning-only** keyword, the OSPF instance never enters the ignore state. When LSA count exceeds the *max* number, the OSPF process logs an error message and the OSPF instance continues in its normal state operation.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure the OSPF instance to accept 12000 nonself-generated LSAs in the global routing table, and 1000 nonself-generated LSAs in VRF V1.

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 0
RP/0/RP0/CPU0:router(config-ospf)# max-lsa 12000
RP/0/RP0/CPU0:router(config-ospf)# vrf V1
RP/0/RP0/CPU0:router(config-ospf)# max-lsa 1000
```

The following example shows how to display the current status of the OSPF instance:

```
RP/0/RP0/CPU0:router# show ospf 0

  Routing Process "ospf 0" with ID 10.0.0.2
  NSR (Non-stop routing) is Disabled
  Supports only single TOS(TOS0) routes
  Supports opaque LSA
  It is an area border router
  Maximum number of non self-generated LSA allowed 12000
     Current number of non self-generated LSA 1
     Threshold for warning message 75%
     Ignore-time 5 minutes, reset-time 10 minutes
     Ignore-count allowed 5, current ignore-count 0
```

max-metric

To configure the Open Shortest Path First (OSPF) protocol to signal other networking devices not to prefer the local router as an intermediate hop in their shortest path first (SPF) calculations, use the **max-metric** command in XR Config mode. To disable this function, use the **no** form of this command.

max-metric router-lsa [**external-lsa** *overriding metric*] [**include-stub**] [**on-proc-migration**] [**on-proc-restart**] [**on-startup**] [**on-switchover**] [**wait-for-bgp**] [**summary-lsa**]
no max-metric router-lsa

Syntax Description

router-lsa	Always originates router link-state advertisements (LSAs) with the maximum metric.
external-lsa *overriding metric*	(Optional) Overrides the external-lsa metric with the max-metric value.The *overriding metric* argument specifies the number of in-summary-LSAs. The range is 1 to 16777215>. The default is 16711680.
include-stub	(Optional) Advertises stub links in router-LSA with the max-metric value (0xFFFF).
on-proc-migration *time*	(Optional) Sets the maximum metric temporarily after a process migration to originate router-LSAs with the max-metric value. The *time* range is 5 to 86400 seconds.
on-proc-restart *time*	(Optional) Sets the maximum metric temporarily after a process restart to originate router-LSAs with the max-metric value. The *time* range is 5 to 86400 seconds.
on-startup *time*	(Optional) Sets the maximum metric temporarily after a reboot to originate router-LSAs with the max-metric value. The *time* range is 5 to 86400 seconds.
on-switchover *time*	(Optional) Sets the maximum metric temporarily after a switchover to originate router-LSAs with the max-metric value. The *time* range is 5 to 86400 seconds. **Note** OSPF will not populate maximum metric on the router's generated LSAs, when the OSPF routing process is configured to support Nonstop Routing (NSR) or Nonstop Forwarding/Graceful restart (NSF/GR).
wait-for-bgp	(Optional) Causes OSPF to originate router LSAs with the maximum metric and allows Border Gateway Protocol (BGP) to decide when to start originating router LSAs with a normal metric instead of the maximum metric.
summary-lsa	(Optional) specifies the number of in summary-LSAs. The range is 1 to 16777215. The default is 16711680.

Command Default

Router LSAs are originated with normal link metrics.

overriding-metric :16711680

Command Modes

XR Config mode

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **max-metric** command to cause the software to originate router LSAs with router link metrics set to LSInfinity (0XFFFF). This feature can be useful in Internet backbone routers that run both OSPF and BGP because OSPF converges more quickly than BGP and may begin attracting traffic before BGP has converged, resulting in dropped traffic.

If this command is configured, the router advertises its locally generated router LSAs with a metric of 0XFFFF. This action allows the router to converge but not attract transit traffic if there are better, alternative paths around this router. After the specified *announce-time* value or notification from BGP has expired, the router advertises the local router LSAs with the normal metric (interface cost).

If this command is configured with the **on-startup** keyword, then the maximum metric is temporarily set only after reboot is initiated. If this command is configured without the **on-startup** keyword, then the maximum metric is permanently used until the configuration is removed.

If the **include-stub** keyword is enabled, the stub-links in the router LSA will be sent with the max-metric. If the **summary-lsa** keyword is enabled, all self-generated summary LSAs will have a metric set to 0xFF0000, unless the metric value is specified with the max-metric value parameter. If the **external-lsa** keyword is enabled, all self-generated external LSAs will have a metric set to 0xFF0000, unless the metric value is specified with the max-metric value parameter.

This command might be useful when you want to connect a router to an OSPF network, but do not want real traffic flowing through it if there are better, alternative paths. If there are no alternative paths, this router still accepts transit traffic as before.

Some cases where this command might be useful are as follows:

- During a router reload, you prefer that OSPF wait for BGP to converge before accepting transit traffic. If there are no alternative paths, the router still accepts transit traffic.
- A router is in critical condition (for example, it has a very high CPU load or does not have enough memory to store all LSAs or build the routing table).
- When you want to gracefully introduce or remove a router to or from the network.
- When you have a test router in a lab, connected to a production network.



Note For older OSPF implementations (RFC 1247), router links in received router LSAs with a metric and cost of LSInfinity are not used during SPF calculations. Hence, no transit traffic is set to the routers originating such router LSAs.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure OSPF to originate router LSAs with the maximum metric until BGP indicates that it has converged:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
RP/0/RP0/CPU0:router(config-ospf)# max-metric router-lsa on-startup wait-for-bgp
```

maximum interfaces (OSPF)

To limit the number of interfaces that can be configured for an Open Shortest Path First (OSPF) process, use the **maximum interfaces** command in the appropriate mode. To return to the default limit, use the **no** form of this command.

maximum interfaces *number-interfaces*
no maximum interfaces

Syntax Description

number-interfaces	Number of interfaces. Range is 1 to 4294967295.

Command Default

If the command is not specified, the default is 1024.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **maximum interface** command to increase or decrease the limit on the number of interfaces configured for an OSPF process.

You cannot configure a limit lower than the number of interfaces currently configured for the OSPF process. To lower the limit, remove interfaces from the OSPF configuration until the number of configured interfaces is at or below the desired limit. You may then apply the new, lower limit.

Task ID

Task ID	Operations
ospf	read, write

Examples

This example shows how to configure a maximum interface limit of 1500 on a router:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
RP/0/RP0/CPU0:router(config-ospf)# maximum interfaces 1500
```

maximum paths (OSPF)

To control the maximum number of parallel routes that the Open Shortest Path First (OSPF) protocol can support, use the **maximum paths** command in an appropriate configuration mode. To remove the **maximum paths** command from the configuration file and restore the system to its default condition with respect to the routing protocol, use the **no** form of this command.

maximum paths *maximum-routes-number*
no maximum paths

Syntax Description

maximum-routes-number	Maximum number of parallel routes that OSPF can install in a routing table. Range is 1 to 64. **Note** The maximum number of paths that can be configured is 64.

Command Default

The default value for maximum-paths depends on the platform supported maximum-paths value. 32 paths

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The default value for maximum-path depends on the platform supported maximum-path value. Any custom value you define for the maximum-path parameter must be within the maximum value supported by platform. The configuration will be rejected if the value you have specified is more then what the platform supports.

When the maximum number of parallel routes is reduced, all existing paths are pruned and paths reinstalled at the new maximum number. During this route-reduction period, you may experience some packet loss for a few seconds. This may impact route traffic.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to allow a maximum of two paths to a destination:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
RP/0/RP0/CPU0:router(config-ospf)# maximum paths 2
```

maximum redistributed-prefixes (OSPF)

To limit the aggregate number of prefixes that can be redistributed into an Open Shortest Path First (OSPF) process, use the **maximum redistributed-prefix** command in the appropriate mode. To return to the default limit, use the **no** form of this command.

maximum redistributed-prefixes *maximum* [*threshold-value*] [**warning-only**]
no maximum redistributed-prefixes

Syntax Description

maximum	Number of routes. Range is 1 to 4294967295.
threshold-value	(Optional) Threshold value (as a percentage) at which to generate a warning message. Range is 1 to 100.
warning-only	(Optional) Gives only a warning when the limit is exceeded.

Command Default

If the command is not specified, the default is 10000.

The threshold value defaults to 75 percent.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.
Release 7.9.1	• Support for the maximum redistributed-prefixes limit was added for the prefixes redistributed as Type-3 LSAs. • Support for counting only the number of redistributed prefixes instead of the number of LSAs was added. (If multiple prefixes are aggregated in an LSA using the summary-prefix command, individual prefixes will be counted)

Usage Guidelines

- Use the **maximum redistributed-prefixes** command to increase or decrease the maximum number of prefixes (also referred to as routes) redistributed for an OSPF process.

  If the maximum value specified is less than the existing number of redistributed prefixes, then the prefixes are reset and the redistribution is restarted.

- Starting Cisco IOS XR Software Release 7.9.1, when the **summary-prefix**command is used to create aggregate LSAs for the routes redistributed into OSPF, each individual redistributed prefix mapped under the subnet advertised in the LSA is counted instead of the number of LSAs. For example, if 5 redistributed prefixes are aggregated by one LSA using the **summary-prefix** command, only the number of redistributed prefixes, which is 5 in this case, is counted and not the number of LSAs, which is 1.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure a maximum number of routes that can be redistributed for an OSPF routing process:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
RP/0/RP0/CPU0:router(config-ospf)# maximum redistributed-prefixes 15000
```

message-digest-key

To specify a key used with Open Shortest Path First (OSPF) Message Digest 5 (MD5) authentication, use the **message-digest-key** command in the appropriate mode. To remove an old MD5 key, use the **no** form of this command.

message-digest-key *key-id* **md5** {*key* | **clear** *key* | **encrypted** *key*}
no message-digest-key *key-id*

Syntax Description

key-id	Key number. Range is 1 to 255.
md5	Enables OSPF MD5 authentication.
key	Alphanumeric string of up to 16 characters.
clear	Specifies that the key be clear text.
encrypted	Specifies that the key be encrypted using a two-way algorithm.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the message digest key parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the message digest key parameter specified for the process.

If this command is not specified at any level, then OSPF MD5 authentication is disabled.

Command Modes

Interface configuration

Area configuration

XR Config mode

Virtual-link configuration

VRF configuration

Multi-area configuration

Sham-link configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Usually, one key individual interface is used to generate authentication information when packets are sent and to authenticate incoming packets. The same key identifier on the neighbor router must have the same *key* value.

For authentication to be enabled, you must configure the **message-digest-key** command together with the **authentication** command and its **message-digest** keyword. Both the **message-digest-key** and **authentication** commands can be inherited from a higher configuration level.

The process of changing keys is as follows. Suppose the current configuration is:

```
interface HundredGigE 0/3/0/2
 message-digest-key 100 md5 OLD
```

You change the configuration to the following:

```
interface HundredGigE 0/3/0/2
 message-digest-key 101 md5 NEW
```

The system assumes its neighbors do not have the new key yet, so it begins a rollover process. It sends multiple copies of the same packet, each authenticated by different keys. In this example, the system sends out two copies of the same packet—the first one authenticated by key 100 and the second one authenticated by key 101.

Rollover allows neighboring routers to continue communication while the network administrator is updating them with the new key. Rollover stops after the local system finds that all its neighbors know the new key. The system detects that a neighbor has the new key when it receives packets from the neighbor authenticated by the new key.

After all neighbors have been updated with the new key, the old key should be removed. In this example, you would enter the following:

```
interface ethernet 1
 no ospf message-digest-key 100
```

Then, only key 101 is used for authentication on interface 1.

We recommend that you not keep more than one key individual interface. Every time you add a new key, you should remove the old key to prevent the local system from continuing to communicate with a hostile system that knows the old key. Removing the old key also reduces overhead during rollover.

Note The MD5 key is always stored in encrypted format on the router. The **clear** and **encrypted** keywords inform the router whether the value that is entered is encrypted or unencrypted.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set a new key 19 with the password *8ry4222* :

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
RP/0/RP0/CPU0:router(config-ospf)# area 0
```

```
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/1/0/1
RP/0/RP0/CPU0:router(config-ospf-ar-if)# message-digest-key 19 md5 8ry4222
```

mpls ldp auto-config (OSPF)

To enable Label Distribution Protocol (LDP)-Interior Gateway Protocol (IGP) interface automatic configuration, use the **mpls ldp auto-config** command in the appropriate mode. To disable LDP-IGP interface automatic configuration, use the **no** form of this command.

mpls ldp auto-config
no mpls ldp auto-config

Command Default

LDP-IGP interface automatic configuration is disabled for OSPF.

Command Modes

Interface configuration

Area configuration

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to enable LDP-IGP interface automatic configuration:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf o1
RP/0/RP0/CPU0:router(config-ospf)# mpls ldp auto-config
```

mpls ldp sync (OSPF)

To enable Label Distribution Protocol (LDP)-Interior Gateway Protocol (IGP) synchronization, use the **mpls ldp sync** command in the appropriate mode. To disable LDP-IGP synchronization, use the **no** form of this command.

mpls ldp sync [**disable**]
no mpls ldp sync

Syntax Description

disable	(Optional) Disables MPLS LDP synchronization from within the OSPF interface and area configuration submodes only. For the OSPF router configuration mode, use the **no** form of the command.

Command Default

LDP-IGP synchronization is disabled for OSPF.

Command Modes

Interface configuration

Area configuration

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to enable LDP-IGP synchronization:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf o1
RP/0/RP0/CPU0:router(config-ospf)# mpls ldp sync
```

mtu-ignore (OSPF)

To prevent Open Shortest Path First (OSPF) from checking whether neighbors are using the same maximum transmission unit (MTU) on a common interface when exchanging database descriptor (DBD) packets, use the **mtu-ignore** command in the appropriate mode. To reset to default, use the **no** form of this command.

mtu-ignore [**disable** | **enable**]
no mtu-ignore

Syntax Description

disable	(Optional) Enables checking for whether OSPF neighbors are using the MTU on a common interface.
enable	(Optional) Disables checking for whether OSPF neighbors are using the MTU on a common interface.

Command Default

The default is **mtu-ignore** with no keywords, which disables MTU checking.

If this command is not specified in interface configuration mode,

then the interface adopts the MTU ignore parameter specified by the area.

If this command is not specified in area configuration mode,

then the interface adopts the MTU ignore parameter specified for the process.

If this command is not specified at any level,

then OSPF checks the MTU received from neighbors when exchanging DBD packets.

Command Modes

Interface configuration

Area configuration

XR Config mode

VRF configuration

Multi-area configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

OSPF checks whether OSPF neighbors are using the same MTU on a common interface. This check is performed when neighbors exchange DBD packets. If the receiving MTU in the DBD packet is higher than the MTU configured on the incoming interface, OSPF adjacency is not established.

The keywords, **disable** and **enable** , do not need to be used. If no keywords are used, the **mtu-ignore** command disables MTU checking. You can then use the **no mtu-ignore** command to activate MTU checking.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to disable MTU mismatch detection on receiving DBD packets:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/1/0/3
RP/0/RP0/CPU0:router(config-ospf-ar-if)# mtu-ignore
```

multi-area-interface

To enable multiple adjacencies for different Open Shortest Path First (OSPF) areas and enter multi-area interface configuration mode, use the **multi-area-interface** command in the area configuration mode. To reset to the default, use the **no** form of this command.

multi-area-interface *type interface-path-id*
no multi-area-interface *type interface-path-id*

Syntax Description

type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.

Command Default

An OSPF network is enabled for one area only.

Command Modes

Area configuration

Command History

Release	Modification
Release 24.2.1	The command is extended to support loopback interfaces.
Release 6.0	This command was introduced.

Usage Guidelines

Use the **multi-area-interface** command to enable area border routers (ABRs) to establish multiple adjacencies for different OSPF areas.

Each multiple area adjacency is announced as a point-to-point unnumbered link in the configured area. This point-to-point link provides a topological path for that area. The first or primary adjacency using the link advertises the link consistent with draft-ietf-ospf-multi-area-adj-06.txt.

You can configure multi-area adjacency on any interface where only two OSF speakers are attached. In the case of native broadcast networks, the interface must be configured as an OPSF point-to-point type using the **network point-to-point** command to enable the interface for a multi-area adjacency.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to enable multiple area adjacency for OSPF 109:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/1/0/3
RP/0/RP0/CPU0:router(config-ospf-ar-if)# area 1
RP/0/RP0/CPU0:router(config-ospf-ar)# multi-area-interface HundredGigE 0/1/0/3
RP/0/RP0/CPU0:router(config-ospf-ar-mif)# ?

    authentication       Enable authentication
    authentication-key   Authentication password (key)
    commit               Commit the configuration changes to running
    cost                 Interface cost
    database-filter      Filter OSPF LSA during synchronization and flooding
    dead-interval        Interval after which a neighbor is declared dead
    describe             Describe a command without taking real actions
    distribute-list      Filter networks in routing updates
    do                   Run an exec command
    exit                 Exit from this submode
    hello-interval       Time between HELLO packets
    message-digest-key   Message digest authentication password (key)
    mtu-ignore           Enable/Disable ignoring of MTU in DBD packets
    no                   Negate a command or set its defaults
    packet-size          Customize size of OSPF packets upto MTU
    pwd                  Commands used to reach current submode
    retransmit-interval  Time between retransmitting lost link state advertisements
    root                 Exit to the global configuration mode
    show                 Show contents of configuration
    transmit-delay       Estimated time needed to send link-state update packet
RP/0/RP0/CPU0:router(config-ospf-ar-mif)#
```

The following example shows how to configure, associate, and reuse a loopback IPv4 interface address for all areas in an ABR:

```
Router(config)#interface Loopback 0
Router(config-if)#ipv4 address 192.0.2.2/32
Router(config-if)#commit

Router(config)#router ospf 100
Router(config-ospf)#router-id 10.10.10.1
Router(config-ospf)#area 1
Router(config-ospf-ar)#interface Loopback 0
Router(config-ospf-ar-if)#commit

Router(config)#router ospf 100
Router(config-ospf)#router-id 10.10.10.1
Router(config-ospf)#area 0
Router(config-ospf-ar)#multi-area-interface Loopback 0
Router(config-ospf-ar-mif)#commit
```

neighbor (OSPF)

To configure Open Shortest Path First (OSPF) routers interconnecting to nonbroadcast networks, use the **neighbor** command in interface configuration mode. To remove a configuration, use the **no** form of this command.

neighbor *ip-address* [**cost** *number*] [**priority** *number*] [**poll-interval** *seconds*]
no neighbor *ip-address* [**cost** *number*] [**priority** *number*] [**poll-interval** *seconds*]

Syntax Description

ip-address	Interface IP address of the neighbor.
cost *number*	(Optional) Assigns a cost to the neighbor, in the form of an integer from 1 to 65535. Neighbors with no specific cost configured assume the cost of the interface, based on the **cost** command. On point-to-multipoint interfaces, **cost** *number* is the only keyword and argument combination that works. The **cost** keyword does not apply to nonbroadcast multiaccess (NBMA) networks.
priority *number*	(Optional) Specifies an 8-bit number indicating the router priority value of the nonbroadcast neighbor associated with the IP address specified. The **priority** keyword does not apply to point-to-multipoint interfaces.
poll-interval *seconds*	(Optional) Specifies an unsigned integer value (in seconds) reflecting the poll interval. RFC 1247 recommends that this value be much larger than the hello interval. The **poll-interval** keyword does not apply to point-to-multipoint interfaces.

Command Default

No configuration is specified.

priority *number* : 0

poll-interval *seconds* : 120 seconds (2 minutes)

Command Modes

Interface configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

You must include one neighbor entry in the software configuration for each known nonbroadcast network neighbor. The neighbor address must be on the primary address of the interface.

If a neighboring router has become inactive (hello packets have not been received for the router dead interval period), it may still be necessary to send hello packets to the dead neighbor. These hello packets are sent at a reduced rate called the *poll interval*.

When the router starts up, it sends only hello packets to those routers with nonzero priority; that is, routers that are eligible to become designated routers (DRs) and backup designated routers (BDRs). After the DR and BDR are selected, the DR and BDR start sending hello packets to all neighbors to form adjacencies.

To filter all outgoing OSPF link-state advertisement (LSA) packets for the neighbor, use the **neighbor database-filter all out** command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to declare a router at address 172.16.3.4 on a nonbroadcast network, with a priority of 1 and a poll interval of 180 seconds:

```
RP/0/RP0/CPU0:router(config-ospf-ar-if)# neighbor 172.16.3.4 priority 1 poll-interval 180
```

The following example illustrates a network with nonbroadcast:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# interface HundredGigE 1/0/0/3
RP/0/RP0/CPU0:router(config-if)# ip address 172.16.3.10 255.255.255.0

RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 1/0/0/3
RP/0/RP0/CPU0:router(config-ospf-ar-if)# network nonbroadcast
RP/0/RP0/CPU0:router(config-ospf-ar-if)# neighbor 172.16.3.4 priority 1 poll-interval 180

RP/0/RP0/CPU0:router(config-ospf-ar-if)# neighbor 172.16.3.5 cost 10 priority 1
poll-interval 180
RP/0/RP0/CPU0:router(config-ospf-ar-if)# neighbor 172.16.3.6 cost 15 priority 1
poll-interval 180
RP/0/RP0/CPU0:router(config-ospf-ar-if)# neighbor 172.16.3.7 priority 1 poll-interval 180
```

neighbor database-filter all out

To filter all outgoing link-state advertisements (LSAs) to an Open Shortest Path First (OSPF) neighbor, use the **neighbor database-filter all out** command in interface configuration mode. To restore the forwarding of LSAs to the neighbor, use the **no** form of this command.

neighbor *ip-address* **database-filter all out**
no neighbor *ip-address* **database-filter all out**

Syntax Description

ip-address	IP address of the neighbor to which outgoing LSAs are blocked.

Command Default

Instead of all outgoing LSAs being filtered to the neighbor, they are flooded to the neighbor.

Command Modes

Interface configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **neighbor database-filter all out** command to filter all outgoing OSPF LSA packets during synchronization and flooding for point-to-multipoint neighbors on nonbroadcast networks. More neighbor options are available with the **neighbor** command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to prevent flooding of OSPF LSAs from point-to-multipoint networks to the neighbor at IP address 10.2.3.4:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 1/0/0/3
RP/0/RP0/CPU0:router(config-ospf-ar-if)# neighbor 10.2.3.4 database-filter all out
```

network (OSPF)

To configure the Open Shortest Path First (OSPF) network type to a type other than the default for a given medium, use the **network** command in the appropriate mode. To return to the default value, use the **no** form of this command.

network {**broadcast** | **non-broadcast** | {**point-to-multipoint** [**non-broadcast**] | **point-to-point**}}
no network

Syntax Description

broadcast	Sets the network type to broadcast.
non-broadcast	Sets the network type to nonbroadcast multiaccess (NBMA).
point-to-multipoint	Sets the network type to point-to-multipoint.
non-broadcast	(Optional) Sets the point-to-multipoint network to be nonbroadcast. If you use this keyword, the **neighbor** command is required.
point-to-point	Sets the network type to point-to-point.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the network parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the network parameter specified for the process.

If this command is not specified at any level, then the OSPF network type is the default of the given medium.

TenGigEthernet interfaces default to broadcast.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **network** command to configure broadcast networks as NBMA networks when, for example, routers in your network do not support multicast addressing.

Configuring NBMA networks as either broadcast or nonbroadcast assumes that there are virtual circuits from every router to every router or fully meshed network. However, there are other configurations where this assumption is not true; for example, a partially meshed network. In these cases, you can configure the OSPF network type as a point-to-multipoint network. Routing between two routers that are not directly connected go through the router that has virtual circuits to both routers. You need not configure neighbors when using this command.

If this command is issued on an interface that does not allow it, this command is ignored.

OSPF has two features related to point-to-multipoint networks. One feature applies to broadcast networks; the other feature applies to nonbroadcast networks:

- On point-to-multipoint, broadcast networks, you can use the **neighbor** command, and you must specify a cost to that neighbor.
- On point-to-multipoint, nonbroadcast networks, you must use the **neighbor** command to identify neighbors. Assigning a cost to a neighbor is optional.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure the OSPF network as a nonbroadcast network:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/1/0/3
RP/0/RP0/CPU0:router(config-ospf-ar-if)# network non-broadcast
RP/0/RP0/CPU0:router(config-ospf-ar-if)# neighbor 172.16.3.4 priority 1 poll-interval 180
```

nsf (OSPF)

To configure nonstop forwarding (NSF) for the Open Shortest Path First (OSPF) protocol, use the **nsf** command in the appropriate mode. To remove this command from the configuration file and restore the system to its default condition, use the **no** form of this command.

nsf {**cisco** [**enforce global**] | **ietf** [**helper disable**]}
no nsf {**cisco** [**enforce global**] | **ietf** [**helper disable**]}

Syntax Description

cisco	Enables Cisco Nonstop Forwarding.
enforce global	(Optional) Cancels NSF restart when non-NSF network device neighbors are detected.
ietf	Enables Internet Engineering Task Force (IETF) graceful restart.
helper disable	(Optional) Disables router helper support.

Command Default

NSF is disabled.

Command Modes

XR Config mode

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The NSF feature allows for the forwarding of data packets to continue along known routes while routing protocol information (such as OSPF) is being restored following a switchover.

Use the **nsf** command if the router is expected to perform NSF during restart. To experience the full benefits of this feature, configure all neighboring routers with NSF.

When this command is used without the optional **cisco enforce global** keywords and non-NSF neighbors are detected, the NSF restart mechanism aborts on the interfaces of those neighbors and functions properly on others.

When this command is used with the optional **cisco enforce global** keywords and non-NSF neighbors are detected, NSF restart is canceled for the entire OSPF process.

IETF graceful restart provides an NSF mechanism to allow data traffic to flow seamlessly with no packet drops during the transient period when OSPF attempts to recover after a process restart or RP failover, within the guidelines of RFC 3623.

By default, neighbors in helper mode listen to both the NSF Cisco- and NSF IETF-type LSAs. The **nsf** command enables one type of mechanism that would undergo an RP failover or, anticipating an OSPF process restart. If the **cisco** or **ietf** keyword is not entered, NSF is not enabled, irrespective of neighbors in listening mode for both NSF Cisco and NSF IETF.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to cancel NSF restart for the entire OSPF process if non-NSF neighbors are detected on any network interface during restart:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# nsf cisco enforce global
```

nsf flush-delay-time (OSPF)

To configure the maximum time allowed for nonstop forwarding (NSF) external route queries for the Open Shortest Path First (OSPF) protocol, use the **nsf flush-delay-time** command in the appropriate mode. To remove this command from the configuration file and restore the system to its default condition, use the **no** form of this command.

nsf flush-delay-time *seconds*
no nsf flush-delay-time *seconds*

Syntax Description

seconds	Length of time (in seconds) allowed for NSF external route queries. Range is 1 to 3600 seconds.

Command Default

seconds : 300

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure the maximum time for NSF to learn external routes for OSPF at 60 seconds:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# nsf flush-delay-time 60
```

nsf interval (OSPF)

To configure the minimum time between consecutive nonstop forwarding (NSF) restart attempts for the Open Shortest Path First (OSPF) protocol, use the **nsf interval** command in the appropriate mode. To remove this command from the configuration file and restore the system to its default condition, use the **no** form of this command.

nsf interval *seconds*
no nsf interval *seconds*

Syntax Description

seconds	Length of time (in seconds) between consecutive restart attempts. Range is 90 to 3600 seconds.

Command Default

seconds : 90

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

When you use the **nsf interval** command, the OSPF process must be up for at least 90 seconds before OSPF attempts to perform an NSF restart.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure the minimum time between consecutive NSF restart attempts at 120 seconds:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:routerr(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# nsf interval 120
```

nsf lifetime (OSPF)

To configure the maximum time that routes are held in the Routing Information Base (RIB) following an Open Shortest Path First (OSPF) process restart, use the **nsf lifetime** command in the appropriate mode. To remove this command from the configuration file and restore the system to its default condition, use the **no** form of this command.

nsf lifetime *seconds*
no nsf lifetime *seconds*

Syntax Description

seconds	The length of time (in seconds) that routes are held in the RIB. Range is 90 to 3600 seconds.

Command Default

seconds : 95

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

When you use this command, the OSPF process must reconverge within the maximum length of time configured. If the convergence exceeds this length of time, routes are purged from RIB and nonstop forwarding (NSF) restart may fail.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure the maximum lifetime for OSPF NSF at 120 seconds:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# nsf lifetime 120
```

nssa (OSPF)

To configure an area as a not-so-stubby area (NSSA), use the **nssa** command in area configuration mode. To remove the NSSA distinction from the area, use the **no** form of this command.

nssa [**no-redistribution**] [**default-information-originate** [**metric** *metric-value*] [**metric-type** *type-value*]] [**no-summary**]
no nssa

Syntax Description

no-redistribution	(Optional) Imports routes only into the normal areas, but not into the NSSA area, by the **redistribute** command when the router is an NSSA Area Border Router (ABR).
default-information-originate	(Optional) Generates a Type 7 default into the NSSA area. This keyword takes effect only on an NSSA ABR or NSSA Autonomous System Boundary Router (ASBR).
metric *metric-value*	(Optional) Specifies the metric used for generating the default route. If you omit a value and do not specify a value using the **defaultmetric** command, the default metric value is 10. Range is 1 to 16777214.
metric-type *type-value*	(Optional) Specifies the external link type associated with the default route advertised into the OSPF routing domain. It can be one of the following values: **1**—Type 1 external route **2**—Type 2 external route
no-summary	(Optional) Prevents an ABR from sending summary link advertisements into the NSSA.

Command Default

No NSSA area is defined.

Command Modes

Area configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

An NSSA does not flood Type 5 external LSAs from the core into the area, but can import autonomous system external routes in a limited fashion within the area.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure area 1 as an NSSA area:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 1
RP/0/RP0/CPU0:router(config-ospf-ar)# nssa
```

ospf name-lookup

To configure the Open Shortest Path First (OSPF) protocol to look up Domain Name System (DNS) names, use the **ospf name-lookup** command in XR Config mode. To disable this function, use the **no** form of this command.

ospf name-lookup
no ospf name-lookup

Command Default

Routers are displayed by router ID or neighbor ID.

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **ospf name-lookup** command to easily identify a router when executing all OSPF **show** command displays. The router is displayed by name rather than by its router ID or neighbor ID.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure OSPF to identify a router by name:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# ospf name-lookup
```

packet-size (OSPF)

To configure the size of Open Shortest Path First (OSPF) packets up to the size specified by the maximum transmission unit (MTU), use the **packet-size** command in the appropriate configuration mode. To disable this function and reestablish the default packet size, use the **no** form of this command.

packet-size *bytes*
no packet-size

Syntax Description

bytes	Size, in bytes. Range is 576 to 10000 bytes.

Command Default

If the command is not specified, the default packet size is either the interface IP MTU size (if that is lower than 9000 bytes) or 9000 bytes.

Command Modes

XR Config mode

Area configuration

Interface configuration

VRF configuration

Multi-area configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **packet-size** command to customize the size of OSPF packets. The OSPF protocol compares the packet size and the MTU size and uses the lower packet size value.

If the command is not configured, the default packet size is equal to the interface IP MTU size (if that is lower than 9000 bytes) or 9000 bytes. For example, if the interface IP MTU size is 1500 bytes, OSPF uses packet size of 1500 bytes on the interface because the byte size is lower than 9000 bytes. If the interface IP MTU size is 9500 bytes, OSPF uses packet size of 9000 bytes on the interface because the byte size exceeds 9000 bytes. The interface IP MTU size depends on the interface and the platform. In most cases, the default interface IP MTU value will be lower than 9000 bytes.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure the packet size on an interface:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 0
```

```
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 1/0/0/2
RP/0/RP0/CPU0:router(config-ospf-ar-if)# packet-size 3500
```

passive (OSPF)

To suppress the sending of Open Shortest Path First (OSPF) protocol operation on an interface, use the **passive** command in the appropriate mode. To remove the passive configuration, use the **no** form of this command.

passive [**disable** | **enable**]
no passive

Syntax Description

disable	(Optional) Sends OSPF updates.
enable	(Optional) Disables sending OSPF updates.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the passive parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the passive parameter specified for the process.

If this command is not specified at any level, then the passive parameter is disabled and OSPF updates are sent on the interface.

Command Modes

Interface configuration

Area configuration

XR Config mode

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

OSPF routing information is neither sent nor received through the specified interface. The interface appears as a stub network in the OSPF router (Type 1) link-state advertisement (LSA).

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows that GigabitEthernet interface 1/0/0/2 reduces OSPF updates because passive mode is enabled; however, HundredGigE interface 0/1/0/3 receives normal OSPF traffic flow:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 1/0/0/2
```

```
RP/0/RP0/CPU0:router(config-ospf-ar-if)# passive
RP/0/RP0/CPU0:router(config-ospf-ar-if)# exit
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 1/0/0/3
RP/0/RP0/CPU0:router(config-ospf-ar-if)# end
```

priority (OSPF)

To set the router priority for an interface, which helps determine the designated router for an Open Shortest Path First (OSPF) link, use the **priority** command in the appropriate mode. To return to the default value, use the **no** form of this command.

priority *value*
no priority *value*

Syntax Description

value	8-bit unsigned integer indicating the router priority value. Range is 0 to 255.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the priority parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the priority parameter specified for the process.

If this command is not specified at any level, then the default priority is 1.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

When two routers attached to a network both attempt to become the designated router, the one with the higher router priority takes precedence. If there is a tie, the router with the higher router ID takes precedence. A router with a router priority set to zero is ineligible to become the designated router or backup designated router. Router priority is configured only for interfaces to multiaccess networks (in other words, not point-to-point networks).

This priority value is used when you configure the Open Shortest Path First (OSPF) protocol for nonbroadcast networks using the **neighbor** command for OSPF.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows that priority is set through the **priority** and **neighbor** commands for Routers A and B and that the neighbor priority value must reflect that of the neighbor router:

```
Router A
  RP/0/RP0/CPU0:router# configure
  RP/0/RP0/CPU0:router(config)# interface HundredGigE 0/1/0/1
  RP/0/RP0/CPU0:router(config-if)# ipv4 address 10.0.0.2 255.255.255.0
  RP/0/RP0/CPU0:router(config-if)# exit
  RP/0/RP0/CPU0:router(config)# router ospf 1
  RP/0/RP0/CPU0:router(config-ospf)# area 0
  RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/1/0/1
  RP/0/RP0/CPU0:router(config-ospf-ar-if)# network non-broadcast
  RP/0/RP0/CPU0:router(config-ospf-ar-if)# priority 4
  RP/0/RP0/CPU0:router(config-ospf-ar-if)# neighbor 10.0.0.1 priority 6

Router B
  RP/0/RP0/CPU0:router# configure
  RP/0/RP0/CPU0:router(config)# interface HundredGigE POS 0/2/0/1
  RP/0/RP0/CPU0:router(config-if)# ipv4 address 10.0.0.1 255.255.255.0
  RP/0/RP0/CPU0:router(config-if)# exit
  RP/0/RP0/CPU0:router(config)# router ospf 1
  RP/0/RP0/CPU0:router(config-ospf)# area 0
  RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/2/0/1
  RP/0/RP0/CPU0:router(config-ospf-ar-if)# network non-broadcast
  RP/0/RP0/CPU0:router(config-ospf-ar-if)# priority 6
  RP/0/RP0/CPU0:router(config-ospf-ar-if)# neighbor 10.0.0.2 priority 4
```

protocol shutdown

To disable an instance of the Open Shortest Path First (OSPF) protocol so that it cannot form an adjacency on any interface, use the **protocol shutdown** command in the XR Config mode. To reenable the OSPF protocol, use the **no** form of this command.

protocol shutdown
no protocol shutdown

Command Default

No default behavior or values

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **protocol shutdown** command to disable the OSPF protocol for a specific routing instance without removing any existing OSPF configuration parameters.

The OSPF protocol continues to run on the router and you can use the current OSPF configuration, but OSPF does not form any adjacencies on any interface.

This command is similar to performing the **no router ospf** command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to disable the OSPF 1 instance:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospf)# protocol shutdown
```

queue dispatch flush-lsa

To change the number of LSAs scheduled (rate-limited) for flushing, that are processed in each iteration, use the **queue dispatch flush-lsa** command in XR Config mode. To return to the system default value, use the **no** form of this command.

queue dispatch flush-lsa *count*
no queue dispatch flush-lsa

Syntax Description

count	Maximum number of LSAs flushed per run. Range is 30 to 3000.

Command Default

The default LSAs flushed per run is 150 (when the count is not configured).

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to limit the number of LSAs flushed per run to 30:

```
RP/0/RP0/CPU0:router(config-ospf)# queue dispatch flush-lsa 30
```

Use the show ospf message-queue, on page 147 command to see the queue dispatch values, peak lengths, and limits.

queue dispatch incoming

To limit the number of incoming packets (LSAUpdates, LSAcks, DBDs, LSRequests, and Hellos that trigger a change state) processed, use the **queue dispatch incoming** command in XR Config mode. To return to the system default value, use the **no** form of this command.

queue dispatch incoming *count*
no queue dispatch incoming

Syntax Description

count	Maximum number of continuous events processed. Range is 30 to 3000.

Command Default

The default incoming count is 300 packets (when the count is not configured).

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how limit the number of incoming packets processed to 500:

```
RP/0/RP0/CPU0:router(config-ospf)# queue dispatch incoming 500
```

Use the show ospf message-queue, on page 147 command to see the queue dispatch values, peak lengths, and limits.

queue dispatch rate-limited-lsa

To set the maximum number of rate-limited link-state advertisement (LSA) (re-)originations processed per run, use the **queue dispatch rate-limited-lsa** command in XR Config mode. To return to the system default value, use the **no** form of this command.

queue dispatch rate-limited-lsa *count*
no queue dispatch rate-limited-lsa

Syntax Description

count	Maximum number of rate-limited LSAs processed per run. Range is 30 to 3000.

Command Default

The default number of rate-limited LSAs processed per run is 300 (when this count is not configured).

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the maximum number of rate-limited LSA (re-)originations processed per run to 300:

```
RP/0/RP0/CPU0:router(config-ospf)# queue dispatch rate-limited-lsa 300
```

queue dispatch spf-lsa-limit

To change the maximum number of Type 3-4 and Type 5-7 link-state advertisements (LSAs) processed per shortest path first (SPF) iteration within a single SPF run, use the **queue dispatch spf-lsa-limit** command in XR Config mode. To return to the system default value, use the **no** form of this command.

queue dispatch spf-lsa-limit *count*
no queue dispatch spf-lsa-limit

Syntax Description

count	Maximum number of continuous Type 3-4 and Type 5-7 LSAs processed per SPF in each scheduled iteration within a single SPF run. Range is 30 to 3000.

Command Default

The default number of Type 3-4 and Type 5-7 processed per run is 150 LSAs (when this command is not configured).

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to limit the number of continuous Type 3-4 and Type 5-7 LSAs processed by SPF per scheduling run, to 100:

```
RP/0/RP0/CPU0:router(config-ospf)# queue dispatch spf-lsa-limit 100
```

queue limit

To set the high watermark for incoming events by priority, use the **queue limit** in XR Config mode. To return to the system default values, use the **no** form of this command.

queue limit {**high** | **medium** | **low**} *count*
no queue limit {**high** | **medium** | **low**}

Syntax Description

high	High watermark for incoming high-priority events (state-changing Hellos).
medium	High watermark for incoming medium-priority events (LSA ACK).
low	High watermark for incoming low-priority events (DBD/LSUpd/LSReq).
count	Maximum number of events per queue. Events are dropped when the priority queue size exceeds this value. Range is 1000 to 30000.

Command Default

High watermark: 9500 (when the corresponding configuration is not present).

Medium watermark: 9000 (when the corresponding configuration is not present).

Low watermark: 8000 (when the corresponding configuration is not present).

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Always keep the limits in the following order of priority:

Limit for High > Limit for Medium > Limit for Low

Task ID

Task ID	Operations
ospf	read, write

Examples

The following examples show how to set the maximum number of events per queue:

```
RP/0/RP0/CPU0:router(config-ospf)# queue limit high 11000
RP/0/RP0/CPU0:router(config-ospf)# queue limit medium 10000
RP/0/RP0/CPU0:router(config-ospf)# queue limit low 9000
```

range (OSPF)

To consolidate and summarize routes at an area boundary, use the **range** command in area configuration mode. To disable this function, use the **no** form of this command.

range *ip-address mask* [**advertise** | **not-advertise**]
no range *ip-address mask* [**advertise** | **not-advertise**]

Syntax Description

ip-address	IP address in four-part, dotted-decimal notation.
mask	IP address mask.
advertise	(Optional) Sets the address range status to advertise and generates a Type 3 summary link-state advertisement (LSA).
not-advertise	(Optional) Sets the address range status to DoNotAdvertise. The Type 3 summary LSA is suppressed and the component networks remain hidden from other networks.

Command Default

When this command is not specified for Area Border Routers (ABRs), routes at an area boundary are not consolidated or summarized.

Advertise is the default.

Command Modes

Area configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **range** command only with Area Border Router (ABRs). Use the command to consolidate or summarize routes for an area. The result is that a single summary route is advertised to other areas by the ABR. Routing information is condensed at area boundaries. External to the area, a single route is advertised for each address range. This process is called *route summarization* .

Multiple range configurations specifying the **range** command can be configured. Thus, the OSPF protocol can summarize addresses for many different sets of address ranges.

The summarized route uses the maximum cost of the routes assumed in the range.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows area 36.0.0.0 consisting of interfaces whose IP addresses have "10.31.x.x" as the first two octets. The **range** command summarizes interfaces. Instead of advertising eight networks individually, the single route 10.31.0.0 255.255.0.0 is advertised:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 201
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar-if)# interface HundredGigE 0/3/0/2
!
RP/0/RP0/CPU0:router(config-ospf)# area 36.0.0.0
RP/0/RP0/CPU0:router(config-ospf-ar)# range 10.31.0.0 255.255.0.0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/1/0/0
RP/0/RP0/CPU0:router(config-ospf-ar-if)# interface HundredGigE 0/1/0/0
RP/0/RP0/CPU0:router(config-ospf-ar-if)# interface HundredGigE 0/1/0/1
RP/0/RP0/CPU0:router(config-ospf-ar-if)# interface HundredGigE 0/1/0/2
RP/0/RP0/CPU0:router(config-ospf-ar-if)# interface HundredGigE 0/1/0/3
RP/0/RP0/CPU0:router(config-ospf-ar-if)# interface HundredGigE 0/2/0/0
RP/0/RP0/CPU0:router(config-ospf-ar-if)# interface HundredGigE 0/2/0/1
RP/0/RP0/CPU0:router(config-ospf-ar-if)# interface HundredGigE 0/2/0/2
RP/0/RP0/CPU0:router(config-ospf-ar-if)# interface HundredGigE 0/2/0/3
RP/0/RP0/CPU0:router(config-ospf-ar-if)# end
```

redistribute (OSPF)

To redistribute routes from one routing domain into Open Shortest Path First (OSPF), use the **redistribute** command in the appropriate mode. To remove the **redistribute** command from the configuration file and restore the system to its default condition in which the software does not redistribute routes, use the **no** form of this command.

Border Gateway Protocol (BGP)
redistribute bgp *process-id* [**preserve-med**] [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]
no redistribute bgp *process-id* [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]

Local Interface Routes
redistribute connected [**instance** *instance-name*] [**instance IPCP**][**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]
no redistribute connected [**instance** *instance-name*] [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]

Directed-attached gateway redundancy (DAGR)
redistribute dagr [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]
no redistribute dagr [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]

Intermediate System-to-Intermediate System (IS-IS)
redistribute isis *process-id* [**level-1** | **level-2** | **level-1-2**] [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]
no redistribute isis *process-id* [**level-1** | **level-2** | **level-1-2**] [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]

Open Shortest Path First (OSPF)
redistribute ospf *process-id* [**match** {**external** [**1** | **2**] | **internal** | **nssa-external** [**1** | **2**]}] [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]
no redistribute ospf *process-id* [**match** {**external** [**1** | **2**] | **internal** | **nssa-external** [**1** | **2**]}] [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]

Routing Information Protocol (RIP)
redistribute rip [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]
no redistribute rip [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]

IP Static Routes
redistribute static [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]
no redistribute static [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]

Lsa-type summary
redistribute *protocol* [**lsa-type**]
summary

Syntax Description

bgp	Distributes routes from the BGP protocol.
process-id	For the **bgp** keyword, an autonomous system number has the following ranges: • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535. For the **isis** keyword, an IS-IS instance name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string. For the **ospf** keyword, an OSPF instance name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.
preserve-med	(Optional) Preserves the Multi Exit Discriminator (MED) of BGP routes.
metric *metric-value*	(Optional) Specifies the metric used for the redistributed route. Range is 1 to 16777214. Use a value consistent with the source protocol.
metric-type { **1** | **2** }	(Optional) Specifies the external link type associated with the route advertised into the OSPF routing domain. It can be one of two values: • **1** —Type 1 external route • **2** —Type 2 external route
tag *tag-value*	(Optional) Specifies the value attached to each external route. This value is not used by the OSPF protocol itself, but is carried in the external LSAs. Range is 0 to 4294967295.
route-policy *policy-name*	(Optional) Specifies the identifier of a configured policy. A policy is used to filter the importation of routes from this source routing protocol to OSPF.
connected	Distributes routes that are established automatically by virtue of having enabled IP on an interface.
instance	Connected instance.
instance-name	Name of the connected instance.
instance IPCP	Distributes routes from IPCP protocols.
isis	Distributes routes from the IS-IS protocol.
level-1	(Optional) Redistributes Level 1 routes into other IP routing protocols independently.

l **evel-1-2**	(Optional) Distributes both Level 1 and Level 2 routes into other IP routing protocols.
level-2	(Optional) Distributes Level 2 routes into other IP routing protocols independently.
ospf	Distributes routes from the OSPF protocol.
match { **internal** | **external [1** | **2**] | **nssa-external** [**1** | **2]** }	(Optional) Specifies the criteria by which OSPF routes are redistributed into other routing domains. It can be one or more of the following: • **internal** —Routes that are internal to a specific autonomous system (intra- and inter-area OSPF routes). • **external** [**1 | 2**]—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 or Type 2 external routes. • **nssa-external** [**1 | 2**]—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 or Type 2 not-so-stubby area (NSSA) external routes. For the **external** and **nssa-external** options, if a type is not specified, then both Type 1 and Type 2 are assumed. If no match is specified, the default is no filtering.
rip	Distributes routes from the RIP protocol.
static	Distributes IP static routes.
dagr	Distributes routes from the directed-attached gateway redundancy (DAGR).
lsa-type	LSA type for redistributed routes.

Command Default

Route redistribution is disabled.

metric *metric-value:* Default is 20 for routes from all protocols except BGP routes, for which the default is 1.

metric-type : Type 2 external route.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Note When redistributing routes (into OSPF) using both command keywords for setting or matching of attributes and a route policy, the routes are run through the route policy first, followed by the keyword matching and setting.

Redistributed routing information should always be filtered by the **policy** *policy-name* keyword and argument. This filtering ensures that only those routes intended by the administrator are redistributed into OSPF.

For information about routing policies, see the *Routing Policy Commands on Router Routing Command Reference Guide*.

Whenever you use the **redistribute** or default-information originate (OSPF), on page 30 command to redistribute routes into an OSPF routing domain, the router automatically becomes an ASBR. However, an ASBR does not, by default, generate a default route into the OSPF routing domain.

When routes are redistributed between OSPF processes, no OSPF metrics are preserved.

When routes are redistributed into OSPF and no metric is specified with the **metric** keyword, OSPF uses 20 as the default metric for routes from all protocols except BGP routes, which get a metric of 1.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to cause BGP routes to be redistributed into an OSPF domain:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 110
RP/0/RP0/CPU0:router(config-ospf)# redistribute bgp 100
```

The following example shows how to redistribute the specified IS-IS process routes into an OSPF domain. The IS-IS routes are redistributed with a metric of 100.

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
RP/0/RP0/CPU0:router(config-ospf)# redistribute isis 108 metric 100
```

In the following example, network 10.0.0.0 appears as an external link-state advertisement (LSA) in OSPF 1:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# interface HundredGigE 0/1/0/1
RP/0/RP0/CPU0:router(config-if)# ip address 10.0.0.0 255.0.0.0
!
RP/0/RP0/CPU0:router(config)# interface HundredGigE 0/2/0/2
RP/0/RP0/CPU0:router(config)# ip address 10.99.0.0 255.0.0.0
!
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# redistribute ospf 2
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/2/0/2
!
RP/0/RP0/CPU0:router(config)# router ospf 2
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/1/0/1
```

The following example displays the configuration of the **redistribute** *protocol* **lsa-type summary** command to redistribute the routes from particular protocol as Type-3 summary LSAs:

```
Router# config
Router(config)#router ospf 1
Router(config-ospf)#redistribute static lsa-type summary
Router(config-ospf)#summary-in
Router(config-ospf)#commit
```

retransmit-interval (OSPF)

To specify the time between link-state advertisement (LSA) retransmissions for adjacencies belonging to the Open Shortest Path First (OSPF) interface, use the **retransmit-interval** command in the appropriate mode. To return to the default value, use the **no** form of this command.

retransmit-interval *seconds*
no retransmit-interval

Syntax Description

seconds	Time (in seconds) between retransmissions. It must be greater than the expected round-trip delay between any two routers on the attached network. Range is 1 to 65535 seconds.

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the retransmit interval parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the retransmit interval parameter specified for the process.

If this command is not specified at any level, then the default retransmit interval is 5 seconds.

Command Modes

Interface configuration

Area configuration

XR Config mode

Virtual-link configuration

VRF configuration

Multi-area configuration

Sham-link configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the router receives no acknowledgment, it resends the LSA.

The setting of this parameter should be conservative, or needless retransmission results. The value should be larger for serial lines and virtual links.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the retransmit interval value to 8 seconds in interface configuration mode:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 201
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/2/0/1
RP/0/RP0/CPU0:router(config-ospf-ar-if)# retransmit-interval 8
```

route-policy (OSPF)

To specify a routing policy to filter Type 3 link-state advertisements (LSA), use the **route-policy** command in area configuration mode. To disable the routing policy, use the **no** form of this command.

route-policy *route-policy-name* {**in** | **out**}
no route-policy *route-policy-name* {**in** | **out**}

Syntax Description

route-policy-name	Name of route policy.
in	Applies policy to inbound routes.
out	Applies policy to outbound routes.

Command Default

No policy is applied.

Command Modes

Area configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **route-policy** command to specify an OSPF routing policy for an inbound or outbound route. The policy can be used to filter routes or modify route attributes.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to specify an OSPF route policy for inbound routes in area 0:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-area)# route-policy area0_in in
```

router-id (OSPF)

To configure a router ID for the Open Shortest Path First (OSPF) process, use the **router-id** command in the appropriate mode. To cause the software to use the default method of determining the router ID, use the **no** form of this command after clearing or restarting the OSPF process.

router-id router-id
no router-id router-id

Syntax Description

router-id	32-bit router ID value specified in four-part, dotted-decimal notation.

Command Default

If this command is not configured, the router ID is the highest IP version 4 (IPv4) address for an interface on the router, with any loopback interface taking precedence.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

It is good practice to use the **router-id** command to explicitly specify a unique 32-bit numeric value for the router ID. This action ensures that OSPF can function regardless of the interface address configuration. Clear the OSPF process using the **clear ospf process** command or restart the OSPF process for the **no router-id** command to take effect.

OSPF attempts to obtain a router ID in the following ways (in order of preference):

1. By default, when the OSPF process initializes, it checks if there is a router-id in the checkpointing database.
2. The 32-bit numeric value specified by the OSPF **router-id** command in router configuration mode. (This value can be any 32-bit value. It is not restricted to the IPv4 addresses assigned to interfaces on this router, and need not be a routable IPv4 address.)
3. The ITAL selected router-id.
4. The primary IPv4 address of an interface over which this OSPF process is running. The first interface address in the OSPF interface is selected.



Note Unlike OSPF version 3, OSPF version 2 is guaranteed to have at least one interface with an IPv4 address configured.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to assign the IP address of 172.20.10.10 to the OSPF process 109:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
RP/0/RP0/CPU0:router(config-ospf)# router-id 172.20.10.10
```

router ospf (OSPF)

To configure an Open Shortest Path First (OSPF) routing process, use the **router ospf** command in XR Config mode. To terminate an OSPF routing process, use the **no** form of this command.

router ospf *process-name*
no router ospf *process-name*

Syntax Description

process-name	Name that uniquely identifies an OSPF routing process. The process name is any alphanumeric string no longer than 40 characters without spaces.

Command Default

No OSPF routing process is defined.

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

You can specify multiple OSPF routing processes in each router. Up to 10 processes can be configured. The recommendation is not to exceed 4 OSPF processes.

All OSPF configuration commands must be configured under an OSPF routing process. For example, two of these commands are the **default-metric** command and the **router-id** command.

Task ID

Task ID	Operations
ospf	read, write
rib	read, write

Examples

The following example shows how to instantiate an OSPF routing process called 109:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
```

security ttl (OSPF)

To set the security time-to-live (TTL) value in the IP header for Open Shortest Path First (OSPF) packets, use the **security ttl** command in the appropriate configuration mode. To remove this command from the configuration file and restore the system to its default condition, use the **no** form of this command.

security ttl [**hops** *hops-number*]
no security ttl

Syntax Description

hops *hops-number*	IP hops. Maximum number of hops allowed. Range is 1 to 254 hops.

Command Default

hops-number : 1

Command Modes

XR Config mode

Area configuration

Interface configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The **security ttl** command is used for the Generalized TTL Security Mechanism (GTSM) feature to prevent network attacks.

During the act of receiving Link State Advertisement (LSA) from neighbors, network attacks can occur because there are no checks that unicast or multicast packets are originating from a neighbor that is one hop away or multiple hops away over virtual links.

For virtual links, OSPF packets travel multiple hops across the network; hence, the TTL value can be decremented several times. For these type of links, a minimum TTL value must be allowed and accepted for multiple-hop packets.

To filter network attacks originating from invalid sources traveling over multiple hops, the GTSM, RFC 3682, is used to prevent the attacks. GTSM filters link-local addresses and allows for only one-hop neighbor adjacencies through the configuration of TTL value 255. The TTL value in the IP header is set to when OSPF packets are originated and checked on the received OSPF packets against the default GTSM TTL value 255 or the user configured GTSM TTL value, blocking unauthorized OSPF packets originated from TTL hops away.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to set the security TTL for an interface:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/6/0/3
RP/0/RP0/CPU0:router(config-ospf-ar-if)# security ttl 2
```

sham-link

To configure an Open Shortest Path First OSPF sham link between two provider edge routers, use the **sham-link** command in VRF area configuration mode. To terminate an (OSPF) sham link, use the **no** form of this command.

sham-link *source-address destination-address*
no sham-link

Syntax Description

source-address	IP address of the local (source) sham-link endpoint specified in four-part, dotted-decimal notation.
destination-address	IP address of the remote (destination) sham-link endpoint specified in four-part, dotted-decimal notation.

Command Default

No sham link is configured.

Command Modes

VRF area configuration.

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **sham-link** command to configure a point-to-point connection between two provider edge (PE) routers creating an interconnect between two VPN sites (VPN backbone). Sham links are configured on PE provider edge (PE) routers in a Multiprotocol Label Switching (MPLS) VPN backbone.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure an OSPF sham link:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 109
RP/0/RP0/CPU0:router(config_ospf)# vrf vrf_a
RP/0/RP0/CPU0:router(config_ospf_vrf)# area 0
RP/0/RP0/CPU0:router(config_ospf_vrf_ar)# sham-link 192.168.40.0 172.16.30.0
RP/0/RP0/CPU0:router(config_ospf_vrf_ar_sl)# cost 23
```

show ospf

To display general information about Open Shortest Path First (OSPF) routing processes, use the **show ospf** command in XR EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [**summary**]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
vrf *vrf-name* **all**	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved values of the *vrf-name* argument.
summary	(Optional) Displays OSPF summary information.

Command Default

IPv4 and unicast address prefixes

Command Modes

XR EXEC mode

OSPFv3

Command History

Release	Modification
Release 7.9.1	This command is modified to display the maximum number of redistributed prefixes.
Release 6.0	This command was introduced.

Usage Guidelines

Use the **show ospf** command to provide basic information about the OSPF processes running on the router. Additional options provide in-depth information.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospf** command:

```
RP/0/RP0/CPU0:router#show ospf

Routing Process "ospf 1" with ID 1.1.1.1
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 It is an area border router
 Initial SPF schedule delay 5000 msecs
 Minimum hold time between two consecutive SPFs 10000 msecs
 Maximum wait time between two consecutive SPFs 10000 msecs
 Initial LSA throttle delay 500 msecs
```

```
Minimum hold time for LSA throttle 5000 msecs
Maximum wait time for LSA throttle 5000 msecs
Minimum LSA interval 5000 msecs. Minimum LSA arrival 1 secs
Maximum number of configured interfaces 255
Number of external LSA 0. Checksum Sum 00000000
Number of opaque AS LSA 0. Checksum Sum 00000000
Number of DCbitless external and opaque AS LSA 0
Number of DoNotAge external and opaque AS LSA 0
Number of areas in this router is 2. 2 normal 0 stub 0 nssa
External flood list length 0
Non-Stop Forwarding enabled
   Area BACKBONE(0) (Inactive)
       Number of interfaces in this area is 2
       SPF algorithm executed 8 times
       Number of LSA 2. Checksum Sum 0x01ba83
       Number of opaque link LSA 0. Checksum Sum 00000000
       Number of DCbitless LSA 0
       Number of indication LSA 0
       Number of DoNotAge LSA 0
       Flood list length 0
   Area 1
       Number of interfaces in this area is 1
       SPF algorithm executed 9 times
       Number of LSA 2. Checksum Sum 0x0153ea
       Number of opaque link LSA 0. Checksum Sum 00000000
       Number of DCbitless LSA 0
       Number of indication LSA 0
       Number of DoNotAge LSA 0
       Flood list length 0
```

This table describes the significant fields shown in the display.

Table 2: show ospf Field Descriptions

Field	Description
Routing Process “ospf 201” with ID 172.22.110.200	OSPF process name.
Supports only	Number of types of service supported (Type 0 only).
It is	Types are internal, area border, or autonomous system boundary.
Redistributing External Routes from	Lists of redistributed routes, by protocol.
SPF schedule delay	Delay time of SPF calculations.
Minimum LSA interval	Minimum interval between LSAs.
Minimum LSA arrival	Minimum elapsed time between accepting an update for the same link-state advertisement (LSA).
external LSA	Total number of Type 5 LSAs in the LSDB.
opaque LSA	Total number of Type 10 LSAs in the LSDB.
DCbitless...AS LSA	Total number of Demand Circuit Type 5 and Type 11 LSAs.

Field	Description
DoNotAge...AS LSA	Total number of Type 5 and Type 11 LSAs with the DoNotAge bit set.
Number of areas	Number of areas in router, area addresses, and so on.
Area BACKBONE	Backbone is area 0.

This sample output from the **show ospf vrf** *vrf_name* command displays the VRF Lite status:

```
RP/0/RP0/CPU0:router#show ospf vrf vrf1

 VRF vrf1 in Routing Process "ospf 100" with ID 1.1.1.1
 NSR (Non-stop routing) is Disabled
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 It is an area border router
 VRF Lite is enabled
 Router is not originating router-LSAs with maximum metric
 Initial SPF schedule delay 50 msecs
 Minimum hold time between two consecutive SPFs 200 msecs
 Maximum wait time between two consecutive SPFs 5000 msecs
 Initial LSA throttle delay 50 msecs
 Minimum hold time for LSA throttle 200 msecs
 Maximum wait time for LSA throttle 5000 msecs
 Minimum LSA interval 200 msecs. Minimum LSA arrival 100 msecs
 LSA refresh interval 1800 seconds
 Flood pacing interval 33 msecs. Retransmission pacing interval 66 msecs
 Adjacency stagger enabled; initial (per area): 2, maximum: 64
    Number of neighbors forming: 0, 2 full
 Maximum number of configured interfaces 1024
 Number of external LSA 0. Checksum Sum 00000000
 Number of opaque AS LSA 0. Checksum Sum 00000000
 Number of DCbitless external and opaque AS LSA 0
 Number of DoNotAge external and opaque AS LSA 0
 Number of areas in this router is 2. 2 normal 0 stub 0 nssa
 External flood list length 0
 SNMP trap is disabled
    Area BACKBONE(0)
        Number of interfaces in this area is 1
        SPF algorithm executed 4 times
        Number of LSA 16.  Checksum Sum 0x071c6a
        Number of opaque link LSA 0.  Checksum Sum 00000000
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0
        Number of LFA enabled interfaces 0, LFA revision 0
        Number of Per Prefix LFA enabled interfaces 0
        Number of neighbors forming in staggered mode 0, 1 full
    Area 1
        Number of interfaces in this area is 4
        SPF algorithm executed 5 times
        Number of LSA 14.  Checksum Sum 0x066d93
        Number of opaque link LSA 0.  Checksum Sum 00000000
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0
        Number of LFA enabled interfaces 0, LFA revision 0
        Number of Per Prefix LFA enabled interfaces 0
```

```
        Number of neighbors forming in staggered mode 0, 1 full
```

The **show ospf** command displays the maximum number of redistributed prefix is limited to 1000.

```
Router #show ospf
Thu Dec  8 18:16:48.332 IST

 Routing Process "ospf 1" with ID 192.168.0.1
 Role: Primary Active
 NSR (Non-stop routing) is Enabled
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 It is an autonomous system boundary router
 Maximum number of non self-generated LSA allowed 1000
    Current number of non self-generated LSA 804
    Threshold for warning message 60%
    Ignore-time 1 minutes, reset-time 2 minutes
    Ignore-count allowed 2, current ignore-count 0
 Redistributing External Routes from,
    static
 Maximum number of redistributed prefixes 1000
    Threshold for warning message 70%
    Current number of redistributed prefixes 100
```

show ospf border-routers

To display the internal Open Shortest Path First (OSPF) routing table entries to an Area Border Router (ABR) and Autonomous System Boundary Router (ASBR), use the **show ospf border-routers** command in XR EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] **border-routers** [*router-id*]

Syntax Description

process-name	(Optional) OSPF process name. If this argument is included, only information for the specified routing process is included.
vrf *vrf-name* **all**	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved vrf-names.
router-id	(Optional) Router ID associated with the border router. The value of the *router-id* argument can be any 32-bit router ID value specified in four-part, dotted-decimal notation. No default exists.

Command Default

IPv4 and unicast address prefixes

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **show ospf border-routers** command to list all OSPF border routers visible to the specified processes and to ascertain the OSPF topology of the router.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospf border-routers** command:

```
RP/0/RP0/CPU0:router# show ospf border-routers

OSPF 1 Internal Routing Table

Codes: i - Intra-area route, I - Inter-area route

i 172.31.97.53 [1] via 172.16.1.53, GigabitEthernet 3/0/0/0, ABR/ASBR , Area 0, SPF 3
```

This table describes the significant fields shown in the display.

Table 3: show ospf border-routers Field Descriptions

Field	Description
i	Type of this route; i indicates an intra-area route, I an interarea route.
172.31.97.53	Router ID of destination.
[1]	Cost of using this route.
172.16.1.53	Next-Next hop toward the destination.
GigabitEthernet 3/0/0/0	Packets destined for 172.16.1.53 are sent over GigabitEthernet interface 3/0/0/0.
ABR/ASBR	Router type of the destination; it is either an Area Border Router (ABR) or Autonomous System Boundary Router (ASBR) or both.
Area 0	Area ID of the area from which this route is learned.
SPF 3	Internal number of the shortest path first (SPF) calculation that installs this route.

show ospf database

To display lists of information related to the Open Shortest Path First (OSPF) database for a specific router, use the **show ospf database** command in XR EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database**
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**adv-router** *ip-address*]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**asbr-summary**] [*link-state-id*]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**asbr-summary**] [*link-state-id*] [**internal**] [**adv-router** [*ip-address*]]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**asbr-summary**] [*link-state-id*] [**internal**] [**self-originate**]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**database-summary**]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**external**] [*link-state-id*]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**external**] [*link-state-id*] [**internal**] [**adv-router** [*ip-address*]]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**external**] [*link-state-id*] [**internal**] [**self-originate**]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**network**] [*link-state-id*]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**network**] [*link-state-id*] [**internal**] [**adv-router** [*ip-address*]]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**network**] [*link-state-id*] [**internal**] [**self-originate**]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**nssa-external**] [*link-state-id*]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**nssa-external**] [*link-state-id*] [**internal**] [**adv-router** [*ip-address*]]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**nssa-external**] [*link-state-id*] [**internal**] [**self-originate**]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**opaque-area**] [*link-state-id*]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**opaque-area**] [*link-state-id*] [**internal**] [**adv-router**] [*ip-address*]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**opaque-area**] [*link-state-id*] [**internal**] [**self-originate**]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**opaque-as**] [*link-state-id*]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**opaque-as**] [*link-state-id*] [**internal**] [**adv-router** [*ip-address*]]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**opaque-as**] [*link-state-id*] [**internal**] [**self-originate**]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**opaque-link**] [*link-state-id*]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**opaque-link**] [*link-state-id*] [**internal**] [**adv-router** [*ip-address*]]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**opaque-link**] [*link-state-id*] [**internal**] [**self-originate**]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**router**] [*link-state-id*]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**router**] [**internal**] [**adv-router** [*ip-address*]]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**router**] [**internal**] [**self-originate**] [*link-state-id*]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**self-originate**]

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**summary**] [*link-state-id*]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**summary**] [*link-state-id*] [**internal**] [**adv-router** [*ip-address*]]
show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **database** [**summary**] [*link-state-id*] [**internal**] [**self-originate**] [*link-state-id*]
show ospf [*process-name*] **database** [**database-summary**] [**detail**]
show ospf [*process-name*] **database** [**database-summary**] [**adv-router** *router ID*]

Syntax Description

process-name	(Optional) OSPF process name that uniquely identifies an OSPF routing process. The process name is any alphanumeric string no longer than 40 characters. If this argument is included, only information for the specified routing process is included.
vrf	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF. The *vrf-name* argument can be specified as an arbitrary string. The strings “default” and “all” are reserved VRF names.
all	(Optional) Specifies all OSPF VRF instances.
area-id	(Optional) Area number used to define the particular area.
adv-router *ip-address*	(Optional) Displays all LSAs of the specified router.
asbr-summary	(Optional) Displays information only about the Autonomous System Boundary Router (ASBR) summary LSAs.
link-state-id	(Optional) Portion of the Internet environment that is being described by the advertisement. The value entered depends on the link-state type of the advertisement. It must be entered in the form of an IP address. When the link-state advertisement (LSA) is describing a network, the *link-state-id* can take one of two forms: • The network IP address (as in Type 3 summary link advertisements and in autonomous system external link advertisements). • A derived address obtained from the link-state ID. **Note** Masking the link-state ID of a network link advertisement with the subnet mask of the network yields the IP address of the network. When the LSA is describing a router, the link-state ID is always the OSPF router ID of the described router. When an autonomous system external advertisement (LS Type = 5) is describing a default route, its link-state ID is set to Default Destination (0.0.0.0).
internal	(Optional) Displays internal LSA information.
self-originate	(Optional) Displays only self-originated LSAs (from the local router).
database-summary	(Optional) Displays how many of each type of LSA for each area there are in the database and the total.
external	(Optional) Displays information only about the external LSAs.

network	(Optional) Displays information only about the network LSAs.
nssa-external	(Optional) Displays information only about the not-so-stubby area (NSSA) external LSAs.
opaque-area	(Optional) Displays information about the opaque Type 10 LSAs. Type 10 denotes an area-local scope. Refer to RFC 2370 for more information on the opaque LSA options.
opaque-as	(Optional) Displays information about the opaque Type 11 LSAs. Type 11 denotes that the LSA is flooded throughout the autonomous system.
opaque-link	(Optional) Displays information about the opaque Type 9 LSAs. Type 9 denotes a link-local scope.
router	(Optional) Displays information only about the router LSAs.
summary	(Optional) Displays information only about the summary LSAs.
detail	(Optional) With **database-summary** , the keyword displays information about the number of LSA counts per router sorted by total LSA count.

Command Default

IPv4 and unicast address prefixes

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.9.1	The **show ospf database database-summary** with the **adv-router** *router ID* keyword displays the router information and the LSAs received from a particular router.
Release 7.9.1	The **show ospf database database-summary** command with the **detail** keyword displays the number of LSA counts per router.
Release 6.0	This command was introduced.

Usage Guidelines

The various forms of the **show ospf database** command deliver information about different OSPF link-state advertisements. This command can be used to examine the link-state database (LSD) and its contents. Each router participating in an area having identical database entries pertaining to that area (with the exception of LSAs that are being flooded). Numerous options (such as **network** and **router**) are used to display portions of the database.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospf database** command when no arguments or keywords are used:

```
RP/0/RP0/CPU0:router# show ospf database

OSPF Router with ID (172.20.1.11) (Process ID 1)

                Router Link States (Area 0)

 Link ID         ADV Router      Age         Seq#       Checksum Link count
 172.20.1.8       172.20.1.8       1381      0x8000010D    0xEF60   2
 172.20.1.11      172.20.1.11      1460      0x800002FE    0xEB3D   4
 172.20.1.12      172.20.1.12      2027      0x80000090    0x875D   3
 172.20.1.27      172.20.1.27      1323      0x800001D6    0x12CC   3

                Net Link States (Area 0)

 Link ID         ADV Router      Age         Seq#       Checksum
 172.22.1.27      172.20.1.27      1323      0x8000005B    0xA8EE
 172.22.1.11      172.20.1.11      1461      0x8000005B    0x7AC

                Type-10 Opaque Link Area Link States (Area 0)

 Link ID         ADV Router      Age         Seq#       Checksum Opaque ID
 10.0.0.0         172.20.1.11      1461      0x800002C8    0x8483     0
 10.0.0.0         172.20.1.12      2027      0x80000080    0xF858     0
 10.0.0.0         172.20.1.27      1323      0x800001BC    0x919B     0
 10.0.0.1         172.20.1.11      1461      0x8000005E    0x5B43     1
```

This table describes the significant fields shown in the display.

Table 4: show ospf database Field Descriptions

Field	Description
Link ID	Router ID number.
ADV Router	ID of the advertising router.
Age	Link-state age.
Seq#	Link-state sequence number (detects old or duplicate LSAs).
Checksum	Fletcher checksum of the complete contents of the LSA.
Link count	Number of interfaces detected for the router.
Opaque ID	Opaque LSA ID number.

The following is sample output from the **show ospf database** command with the **asbr-summary** keyword:

```
RP/0/RSP0RP0/CPU0:router# show ospf database asbr-summary

OSPF Router with ID (192.168.0.1) (Process ID 300)

Summary ASB Link States (Area 0.0.0.0)

  LS age: 1463
```

```
  Options: (No TOS-capability)
  LS Type: Summary Links (AS Boundary Router)
  Link State ID: 172.17.245.1 (AS Boundary Router address)
  Advertising Router: 172.17.241.5
  LS Seq Number: 80000072
  Checksum: 0x3548
  Length: 28
  Network Mask: /0
TOS: 0  Metric: 1
```

This table describes the significant fields shown in the display.

Table 5: show ospf database asbr-summary Field Descriptions

Field	Description
OSPF Router with ID	Router ID number.
Process ID	OSPF process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only).
LS Type	Link-state type.
Link State ID	Link-state ID (ASBR).
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSAs.
Network Mask	Network mask implemented.
TOS	Type of service.
Metric	Link-state metric.

The following is sample output from the **show ospf database** command with the **external** keyword:

```
RP/0/RP0/CPU0:router# show ospf database external

OSPF Router with ID (192.168.0.1) (Process ID 300)

                Type-5 AS External Link States

  LS age: 280
  Options: (No TOS-capability)
  LS Type: AS External Link
  Link State ID: 172.17.0.0 (External Network Number)
  Advertising Router: 172.17.70.6
  LS Seq Number: 80000AFD
```

```
Checksum: 0xC3A
Length: 36
Network Mask: 255.255.0.0
      Metric Type: 2 (Larger than any link state path)
      TOS: 0
      Metric: 1
      Forward Address: 0.0.0.0
      External Route Tag: 0
```

This table describes the significant fields shown in the display.

Table 6: show ospf database external Field Descriptions

Field	Description
OSPF Router with Router ID	Router ID number.
Process ID	OSPF process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only).
LS Type	Link-state type.
Link State ID	Link-state ID (external network number).
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence number (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSA.
Network Mask	Network mask implemented.
Metric Type	External type.
TOS	Type of service.
Metric	Link-state metric.
Forward Address	Forwarding address. Data traffic for the advertised destination is forwarded to this address. If the forwarding address is set to 0.0.0.0, data traffic is forwarded instead to the originator of the advertisement.
External Route Tag	External route tag, a 32-bit field attached to each external route. This tag is not used by the OSPF protocol itself.

The following is sample output from the **show ospf database** command with the **network** keyword:

```
RP/0/RP0/CPU0:router# show ospf database network

 OSPF Router with ID (192.168.0.1) (Process ID 300)
```

```
Net Link States (Area 0.0.0.0)

  LS age: 1367
  Options: (No TOS-capability)
  LS Type: Network Links
  Link State ID: 172.23.1.3 (address of Designated Router)
  Advertising Router: 192.168.0.1
  LS Seq Number: 800000E7
  Checksum: 0x1229
  Length: 52
  Network Mask: /24
        Attached Router: 192.168.0.1
        Attached Router: 172.23.241.5
        Attached Router: 172.23.1.1
        Attached Router: 172.23.54.5
        Attached Router: 172.23.1.5
```

This table describes the significant fields shown in the display.

Table 7: show ospf database network Field Descriptions

Field	Description
OSPF Router with ID	Router ID number.
Process ID	OSPF process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only).
LS Type	Link-state type.
Link State ID	Link-state ID of the designated router.
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence number (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSA.
Network Mask	Network mask implemented.
Attached Router	List of routers attached to the network, by IP address.

The following is sample output, carrying Multiprotocol Label Switching traffic engineering (MPLS TE) specification information, from the **show ospf database** command with the **opaque-area** keyword and a *link-state-id* of adv-router:

```
RP/0/RP0/CPU0:router# show ospf database opaque-area adv-router 172.20.1.12

  OSPF Router with ID (172.20.1.11) (Process ID 1)

                Type-10 Opaque Link Area Link States (Area 0)
```

```
    LS age: 224
    Options: (No TOS-capability, DC)
    LS Type: Opaque Area Link
    Link State ID: 1.0.0.0
    Opaque Type: 1
    Opaque ID: 0
    Advertising Router: 172.20.1.12
    LS Seq Number: 80000081
    Checksum: 0xF659
    Length: 132
    Fragment number : 0

      MPLS TE router ID : 172.20.1.12

      Link connected to Point-to-Point network
        Link ID : 172.20.1.11
        Interface Address : 172.21.1.12
        Neighbor Address : 172.21.1.11
        Admin Metric : 10
        Maximum bandwidth : 193000
        Maximum reservable bandwidth : 125000
        Number of Priority : 8
        Priority 0 : 125000      Priority 1 : 125000
        Priority 2 : 125000      Priority 3 : 125000
        Priority 4 : 125000      Priority 5 : 125000
        Priority 6 : 125000      Priority 7 : 100000
        Affinity Bit : 0x0

      Number of Links : 1
```

The following is the sample output from the **show ospf database opaque-area** command displaying the extended link LSA information.

```
RP/0/RP0/CPU0:router# show ospf database opaque-area 4.0.0.0
LS age: 361
  Options: (No TOS-capability, DC)
  LS Type: Opaque Area Link
  Link State ID: 8.0.0.40
  Opaque Type: 8
  Opaque ID: 40
  Advertising Router: 100.0.0.3
  LS Seq Number: 8000012e
Checksum: 0xeab4
  Length: 92

    Extended Link TLV: Length: 68
      Link-type : 2
      Link ID   : 100.0.9.4
      Link Data : 100.0.9.3

     LAN Adj sub-TLV: Length: 16
       Flags      : 0x0
       MTID       : 0
       Weight     : 0
       Neighbor ID: 100.0.0.1

       SID/Label sub-TLV: Length: 3
          SID       : 24001

     LAN Adj sub-TLV: Length: 16
```

```
    Flags      : 0x0
    MTID       : 0
    Weight     : 0
    Neighbor ID: 100.0.0.2

    SID/Label sub-TLV: Length: 3
       SID       : 24000

  Adj sub-TLV: Length: 12
     Flags     : 0x0
     MTID      : 0
     Weight    : 0

    SID/Label sub-TLV: Length: 3
       SID       : 24002
```

The following is sample output from the **show ospf database** command that displays a Type 10, Router Information LSA:

```
RP/0/RP0/CPU0:router# show ospf database opaque-area 4.0.0.0

            OSPF Router with ID (3.3.3.3) (Process ID orange)

                Type-10 Opaque Link Area Link States (Area 0)

  LS age: 105
  Options: (No TOS-capability, DC)
  LS Type: Opaque Area Link
  Link State ID: 4.0.0.0
  Opaque Type: 4
  Opaque ID: 0
  Advertising Router: 3.3.3.3
  LS Seq Number: 80000052
  Checksum: 0x34e2
  Length: 52
  Fragment number: 0

    Router Information TLV: Length: 4
    Capabilities:
      Graceful Restart Helper Capable
      Traffic Engineering enabled area
      All capability bits: 0x50000000

    PCE Discovery TLV: Length: 20
      IPv4 Address: 3.3.3.3
      PCE Scope: 0x20000000
      Compute Capabilities:
      Inter-area default (Rd-bit)
      Compute Preferences:
      Intra-area: 0  Inter-area:  0
      Inter-AS:   0  Inter-layer: 0
```

This table describes the significant fields shown in the display.

Table 8: show ospf database opaque-area Field Descriptions

Field	Description
OSPF Router with ID	Router ID number.

Field	Description
Process ID	OSPF process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only).
LS Type	Link-state type.
Link State ID	Link-state ID.
Opaque Type	Opaque link-state type.
Opaque ID	Opaque ID number.
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).
Length	Length (in bytes) of the LSA.
Fragment number	Arbitrary value used to maintain multiple traffic engineering LSAs.
Link ID	Link ID number.
Interface Address	ID address of the interface.
Neighbor Address	IP address of the neighbor.
Admin Metric	Administrative metric value used by MPLS TE.
Maximum bandwidth	Specifies maximum bandwidth (in kbps).
Maximum reservable bandwidth	Specifies maximum reservable bandwidth (in kbps).
Number of Priority	Priority number.
Affinity Bit	Used by MPLS TE.
Router Information TLV	Router capabilities are advertised in this TLV.
Capabilities	Some router capabilities include stub router, traffic engineering, graceful restart, and graceful restart helper.
PCE Discovery TLV	PCE address and capability information is advertised in this TLV.
IPv4 Address	Configured PCE IPv4 address.
PCE Scope	Computation capabilities of the PCE.
Compute Capabilities	Compute capabilities and preferences of the PCE.

Field	Description
Inter-area default (RD-bit)	PCE compute capabilities such as intra-area, inter-area, inter-area default, inter-AS, inter-AS default and inter-layer.
Compute Preferences	Order or preference of path computation that includes intra-area, inter-area, inter-AS, and inter-layer preferences.

The following is sample output from the **show ospf database** command with the **router** keyword:

```
RP/0/RP0/CPU0:router# show ospf database router

OSPF Router with ID (192.168.0.1) (Process ID 300)

Router Link States (Area 0.0.0.0)

  LS age: 1176
  Options: (No TOS-capability)
  LS Type: Router Links
  Link State ID: 172.23.21.6
  Advertising Router: 172.23.21.6
  LS Seq Number: 80002CF6
  Checksum: 0x73B7
  Length: 120
  AS Boundary Router
  Number of Links: 8

 Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 172.23.21.5
(Link Data) Router Interface address: 172.23.21.6
 Number of TOS metrics: 0
  TOS 0 Metrics: 2
```

This table describes the significant fields shown in the display.

Table 9: show ospf database router Field Descriptions

Field	Description
OSPF Router with ID	Router ID number.
Process ID	OSPF process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only).
LS Type	Link-state type.
Link State ID	Link-state ID.
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).

Field	Description
Length	Length (in bytes) of the LSA.
AS Boundary Router	Definition of router type.
Number of Links	Number of active links.
Link ID	Link type.
Link Data	Router interface address.
TOS	Type of service metric (Type 0 only).

The following is sample output from **show ospf database** command with the **summary** keyword:

```
RP/0/RP0/CPU0:router# show ospf database summary

       OSPF Router with ID (192.168.0.1) (Process ID 300)

 Summary Net Link States (Area 0.0.0.0)

  LS age: 1401
  Options: (No TOS-capability)
  LS Type: Summary Links (Network)
  Link State ID: 172.23.240.0 (Summary Network Number)
  Advertising Router: 172.23.241.5
  LS Seq Number: 80000072
  Checksum: 0x84FF
  Length: 28
  Network Mask: /24
     TOS: 0  Metric: 1
```

This table describes the significant fields shown in the display.

Table 10: show ospf database summary Field Descriptions

Field	Description
OSPF Router with ID	Router ID number.
Process ID	OSPF process name.
LS age	Link-state age.
Options	Type of service options (Type 0 only).
LS Type	Link-state type.
Link State ID	Link-state ID (summary network number).
Advertising Router	ID of the advertising router.
LS Seq Number	Link-state sequence (detects old or duplicate LSAs).
Checksum	Link-state checksum (Fletcher checksum of the complete contents of the LSA).

Field	Description
Length	Length (in bytes) of the LSA.
Network Mask	Network mask implemented.
TOS	Type of service.
Metric	Link-state metric.

The following is sample output from **show ospf database** command with the **database-summary** keyword:

```
RP/0/RP0/CPU0:router# show ospf database database-summary

          OSPF Router with ID (172.19.65.21) (Process ID 1)

Area 0 database summary
  LSA Type      Count    Delete   Maxage
  Router        2        0        0
  Network       1        0        0
  Summary Net   2        0        0
  Summary ASBR  0        0        0
  Type-7 Ext    0        0        0
  Opaque Link   0        0        0
  Opaque Area   0        0        0
  Subtotal      5        0        0

Process 1 database summary
  LSA Type      Count    Delete   Maxage
  Router        2        0        0
  Network       1        0        0
  Summary Net   2        0        0
  Summary ASBR  0        0        0
  Type-7 Ext    0        0        0
  Opaque Link   0        0        0
  Opaque Area   0        0        0
  Type-5 Ext    2        0        0
  Opaque AS     0        0        0
  Total         7        0        0
```

The **show ospf database database-summary** command with the **detail** keyword displays the number of LSA counts per router:

```
Router#show ospf database database-summary detail
Tue Dec  6 19:20:34.090 IST


            OSPF Router with ID (192.168.0.1) (Process ID 1)

Router 192.168.0.4 LSA summary
  LSA Type      Count    Delete   Maxage
  Router        0        0        0
  Network       0        0        0
  Summary Net   0        0        0
  Summary ASBR  0        0        0
  Type-5 Ext    697      0        0
  Type-7 Ext    0        0        0
  Opaque Link   0        0        0
```

```
  Opaque Area   0         0         0
  Opaque AS     6         0         0
  Total         703       0         0

Router 192.168.0.1 LSA summary
  LSA Type      Count     Delete    Maxage
  Router        1         0         0
  Network       0         0         0
  Summary Net   0         0         0
  Summary ASBR  0         0         0
  Type-5 Ext    0         0         0
  Type-7 Ext    0         0         0
  Opaque Link   0         0         0
  Opaque Area   64        0         0
  Opaque AS     0         0         0
  Total         65        0         0

Router 192.168.0.2 LSA summary
  LSA Type      Count     Delete    Maxage
  Router        1         0         0
  Network       0         0         0
  Summary Net   21        0         0
  Summary ASBR  2         0         0
  Type-5 Ext    0         0         0
  Type-7 Ext    0         0         0
  Opaque Link   0         0         0
  Opaque Area   21        0         0
  Opaque AS     0         0         0
  Total         45        0         0

Router 192.168.0.6 LSA summary
  LSA Type      Count     Delete    Maxage
  Router        1         0         0
  Network       0         0         0
  Summary Net   21        0         0
  Summary ASBR  2         0         0
  Type-5 Ext    0         0         0
  Type-7 Ext    0         0         0
  Opaque Link   0         0         0
  Opaque Area   19        0         0
  Opaque AS     0         0         0
  Total         43        0         0

Router 192.168.0.3 LSA summary
  LSA Type      Count     Delete    Maxage
  Router        0         0         0
  Network       0         0         0
  Summary Net   0         0         0
  Summary ASBR  0         0         0
  Type-5 Ext    7         0         0
  Type-7 Ext    0         0         0
  Opaque Link   0         0         0
  Opaque Area   0         0         0
  Opaque AS     6         0         0
  Total         13        0         0
```

The **show ospf database database-summary** with the **adv-router** *router ID* keyword displays the router information and the LSAs received from a particular router:

```
Router#show ospf database database-summary adv-router 192.168.0.4
Tue Dec  6 19:21:04.549 IST

            OSPF Router with ID (192.168.0.1) (Process ID 1)
```

```
Router 192.168.0.4 LSA summary
  LSA Type      Count    Delete   Maxage
  Router        0        0        0
  Network       0        0        0
  Summary Net   0        0        0
  Summary ASBR  0        0        0
  Type-5 Ext    697      0        0
  Type-7 Ext    0        0        0
  Opaque Link   0        0        0
  Opaque Area   0        0        0
  Opaque AS     6        0        0
  Total         703      0        0
```

This table describes the significant fields shown in the display.

Table 11: show ospf database database-summary Field Descriptions

Field	Description
LSA Type	Link-state type.
Count	Number of advertisements in that area for each link-state type.
Delete	Number of LSAs that are marked “Deleted” in that area.
Maxage	Number of LSAs that are marked “Maxaged” in that area.

show ospf flood-list

To display a list of Open Shortest Path First (OSPF) link-state advertisements (LSAs) waiting to be flooded over an interface, use the **show ospf flood-list** command in XR EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **flood-list** [*type interface-path-id*]

Syntax Description

process-name	(Optional) OSPF process name that uniquely identifies an OSPF routing process. The process name is any alphanumeric string no longer than 40 characters. If this argument is included, only information for the specified routing process is included.
vrf	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.
all	(Optional) Specifies all OSPF VRF instances.
area-id	(Optional) Area number used to define the particular area.
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.

Command Default

All interfaces

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **show ospf flood-list** command to display LSAs in flood queue and queue length.

Flood list information is transient and normally the flood lists are empty.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospf flood-list** command for interface GigabitEthernet 3/0/0/0:

```
RP/0/RP0/CPU0:router# show ospf flood-list HundredGigE 3/0/0/0


 Interface GigabitEthernet3/0/0/0, Queue length 20
 Link state retransmission due in 12 msec
 Displaying 6 entries from flood list:

 Type  LS ID           ADV RTR         Seq NO      Age    Checksum
    5  10.2.195.0       200.0.0.163     0x80000009  0      0xFB61
    5  10.1.192.0       200.0.0.163     0x80000009  0      0x2938
    5  10.2.194.0       200.0.0.163     0x80000009  0      0x757
    5  10.1.193.0       200.0.0.163     0x80000009  0      0x1E42
    5  10.2.193.0       200.0.0.163     0x80000009  0      0x124D
    5  10.1.194.0       200.0.0.163     0x80000009  0      0x134C
```

This table describes the significant fields shown in the display.

Table 12: show ospf flood-list Field Descriptions

Field	Description
GigabitEthernet3/0/0/0	Interface for which information is displayed.
Queue length	Number of LSAs waiting to be flooded.
Link state retransmission due in	Length of time (in milliseconds) before next link-state transmission.
Type	Type of LSA.
LS ID	Link-state ID of the LSA.
ADV RTR	IP address of the advertising router.
Seq NO	Sequence number of the LSA.
Age	Age of the LSA (in seconds).
Checksum	Checksum of the LSA.

show ospf interface

To display strict-mode information use the **show ospf interface** command in XR EXEC mode.

show ospf interface [**interface** *type interface-path-id*]

Syntax Description

type	Interface type. For more information, use the question mark (**?**) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (**?**) online help function.

Command Default

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
bgp	read
ospf	read

show ospf interface

The following example shows the output from the **show ospf interface** command.

```
RP/0/0/CPU0:rt2(config-ospf-ar-if)# show ospf interface HundredGigE 0/2/0/0
Sun Feb 15 12:17:35.072 IST

HundredGigE 0/2/0/0 is up, line protocol is up
  Internet Address 10.1.1.2/24, Area 0
  Process ID 1, Router ID 2.2.2.2, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DR, Priority 1, MTU 1500, MaxPktSz 1500
  BFD enabled, BFD interval 150 msec, BFD multiplier 3, Mode: Strict
  Designated Router (ID) 2.2.2.2, Interface address 10.1.1.2
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:07:358
  Index 1/1, flood queue length 0
  Next 0(0)/0(0)
  Last flood scan length is 1, maximum is 1
```

```
Last flood scan time is 0 msec, maximum is 0 msec
LS Ack List: current length 0, high water mark 1
Neighbor Count is 1, Adjacent neighbor count is 0
Suppress hello for 0 neighbor(s)
Multi-area interface Count is 0
```

show ospf message-queue

To display the information about the queue dispatch values, peak lengths, and limits, use the **show ospf message-queue** command in XR EXEC mode.

show ospf message-queue

This command has no arguments or keywords.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospf message-queue** command:

```
RP/0/RP0/CPU0:router# show ospf 1 message-queue

OSPF 1
  Hello Input Queue:
    Current queue length: 0
    Event scheduled: 0
    Total queuing failures: 0
    Maximum length : 102
    Pkts pending processing: 0
    Limit: 5000

  Router Message Queue
    Current instance queue length: 0
    Current redistribution queue length: 0
    Current ex spf queue length: 0
    Current sum spf queue length: 0
    Current intra spf queue length: 0
    Event scheduled: 0
    Maximum length : 101
    Total low queuing failures: 0
    Total medium queuing failures: 0
    Total high queuing failures: 0
    Total instance events: 919
    Processing quantum : 300
    Low queuing limit: 8000
    Medium queuing limit: 9000
    High queuing limit: 9500
```

```
    Rate-limited LSA processing quantum: 150
    Current rate-limited LSA queue length: 0
    Rate-limited LSA queue peak len: 517

    Rate-limited LSAs processed: 4464
    Flush LSA processing quantum: 150
    Current flush LSA queue length: 0
    Flush LSA queue peak len: 274
    Rate-limited flush LSAs processed: 420

    SPF-LSA-limit processing quantum: 150
    Managed timers processing quantum: 50
    Instance message count: 0
    Instance pulse send count: 919
    Instance pulse received count: 919
    Global pulse count: 0
    Instance Pulse errors: 0

  TE Message Queue
    Current queue length: 0
    Total queuing failures: 0
    Maximum length : 0

Number of Dlink errors: 0
```

This table describes the significant fields shown in the display.

Table 13: show ospf message-queue Field Descriptions

Field	Description
Hello Input Queue	This section provides statistics on the number of events and incoming packets processed in the Hello (incoming packet) thread of the OSPF process.
Router Message Queue	This section provides statistics on the events and messages processed in the Router (primary) thread of the OSPF process.
TE Message Queue	This section provides statistics on traffic-engineering events and messages received by OSPF from TE (the te_control process). These events are processed in the Router thread of the OSPF process.
Number of Dlink errors	The number of enqueuing or dequeuing errors seen across all the linked-lists in the OSPF process.

show ospf neighbor

To display Open Shortest Path First (OSPF) neighbor information on an individual interface basis, use the **show ospf neighbor** command in XR EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **neighbor** [[*type* *interface-path-id*] [*neighbor-id*] [**detail**] | **area-sorted** | **capabilities**]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
vrf *vrf-name* **all**	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.
area-id	(Optional) Area ID. If you do not specify an area, all areas are displayed.
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.
neighbor-id	(Optional) Neighbor ID.
detail	(Optional) Displays all neighbors given in detail (lists all neighbors).
area-sorted	(Optional) Specifies that all neighbors are grouped by area.
capabilities	(Optional) Displays the capabilties of all neighbors advertised in the Router Information Link State Advertisement (RI LSA).

Command Default

All neighbors

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 24.2.1	The **capabilities** keyword was introduced.
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospf neighbor** command showing two lines of summary information for each neighbor:

```
RP/0/RP0/CPU0:router# show ospf neighbor

Neighbors for OSPF

Neighbor ID     Pri  State         Dead Time  Address           Interface
192.168.199.137 1    FULL/DR       0:00:31    172.31.80.37      HundredGigE 0/3/0/2
    Neighbor is up for 18:45:22
192.168.48.1     1    FULL/DROTHER  0:00:33    192.168.48.1       HundredGigE 0/3/0/3
    Neighbor is up for 18:45:30
192.168.48.200   1    FULL/DROTHER  0:00:33    192.168.48.200     HundredGigE 0/3/0/3
    Neighbor is up for 18:45:25
192.168.199.137 5    FULL/DR       0:00:33    192.168.48.189     HundredGigE 0/3/0/3
    Neighbor is up for 18:45:27
```

This table describes the significant fields shown in the display.

*Table 14: **show ospf neighbor** Field Descriptions*

Field	Description
Neighbor ID	Neighbor router ID.
Pri	Designated router priority.
State	OSPF state.
Dead time	Time (in hh:mm:ss) that must elapse before OSPF declares the neighbor dead.
Address	Address of next hop.
Interface	Interface name of next hop.
Neighbor is up	Amount of time (in hh:mm:ss) that the OSPF neighbor has been up.

The following is sample output showing summary information about the neighbor that matches the neighbor ID:

```
RP/0/RP0/CPU0:router# show ospf neighbor 192.168.199.137

 Neighbor 192.168.199.137, interface address 172.31.80.37
    In the area 0.0.0.0 via interface HundredGigE 0/3/0/2
    Neighbor priority is 1, State is FULL, 6 state changes
    DR is 0.0.0.0 BDR is 0.0.0.0
    Options is 0x2
    Dead timer due in 0:00:32
```

```
    Neighbor is up for 18:45:30
    Number of DBD retrans during last exhange 0
    Index 1/1, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum 0 msec
 Neighbor 192.168.199.137, interface address 192.168.48.189
    In the area 0.0.0.0 via interface HundredGigE 0/3/0/3
    Neighbor priority is 5, State is FULL, 6 state changes
    Options is 0x2
    Dead timer due in 0:00:32
    Neighbor is up for 18:45:30
    Number of DBD retrans during last exhange 0
    Index 1/1, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum 0 msec

Total neighbor count: 2
```

This table describes the significant fields shown in the display.

Table 15: show ospf neighbor 192.168.199.137 Field Descriptions

Field	Description
Neighbor	Neighbor router ID.
interface address	IP address of the interface.
In the area	Area and interface through which the OSPF neighbor is known.
Neighbor priority	Router priority of neighbor and neighbor state.
State	OSPF state.
state changes	Number of state changes for this neighbor.
DR is	Neighbor ID of the designated router.
BDR is	Neighbor ID of the backup designated router.
Options	Hello packet options field contents(E-bit only; possible values are 0 and 2; 2 indicates area is not a stub; 0 indicates area is a stub).)
Dead timer	Time (in hh:mm:ss) to elapse before OSPF declares the neighbor dead.
Neighbor is up	Amount of time (in hh:mm:ss) that the OSPF neighbor has been up.
Number of DBD retrans	Number of re-sent database description packets.
Index	Index and the remaining lines of this command give detailed information about flooding information received from the neighbor.

If you specify the interface along with the neighbor ID, the software displays the neighbors that match the neighbor ID on the interface, as in the following sample display:

```
RP/0/RP0/CPU0:router# show ospf neighbor HundredGigE 0/3/0/2 192.168.199.137

 Neighbor 192.168.199.137, interface address 172.31.80.37
    In the area 0.0.0.0 via interface HundredGigE 0/3/0/2
    Neighbor priority is 1, State is FULL, 6 state changes
    DR is 0.0.0.0 BDR is 0.0.0.0
    Options is 0x2
    Dead timer due in 0:00:32
    Neighbor is up for 18:45:30
    Number of DBD retrans during last exhange 0
    Index 1/1, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum 0 msec

Total neighbor count: 1
```

This table describes the significant fields shown in the display.

Table 16: show ospf neighbor HundredGigE 0/3/0/2 192.168.199.137 Field Descriptions

Field	Description
Neighbor	Neighbor router ID.
interface address	IP address of the interface.
In the area	Area and interface through which the OSPF neighbor is known.
Neighbor priority	Router priority of the neighbor.
State	OSPF state.
state changes	Number of state changes for this neighbor.
DR is	Neighbor ID of the designated router.
BDR is	Neighbor ID of the backup designated router.
Options	Hello packet options field contents (E-bit only; possible values are 0 and 2; 2 indicates area is not a stub; 0 indicates area is a stub)
Dead timer	Time (in hh:mm:ss) to elapse before OSPF declares the neighbor dead.
Neighbor is up	Amount of time (in hh:mm:ss) that the OSPF neighbor has been up.
Number of DBD retrans	Number of re-sent database description packets.
Index	Index and the remaining lines of this command give detailed information about flooding information received from the neighbor.

You can also specify the interface without the neighbor ID to show all neighbors on the specified interface, as in the following sample display:

```
RP/0/RP0/CPU0:router# show ospf neighbor HundredGigE POS 0/3/0/3
```

```
Neighbors for OSPF ospf1

   ID           Pri   State           Dead Time     Address         Interface
192.168.48.1     1    FULL/DROTHER  0:00:33    192.168.48.1       HundredGigE POS 0/3/0/3

    Neighbor is up for 18:50:52
192.168.48.200   1    FULL/DROTHER  0:00:32    192.168.48.200     HundredGigE POS 0/3/0/3

    Neighbor is up for 18:50:52
192.168.199.137 5    FULL/DR       0:00:32    192.168.48.189     HundredGigE POS 0/3/0/3
    Neighbor is up for 18:50:52

Total neighbor count: 3
```

This table describes the significant fields shown in the display.

Table 17: show ospf neighbor HundredGigE 0/3/0/3 Field Descriptions

Field	Description
ID	Neighbor router ID.
Pri	Route priority of the neighbor.
State	OSPF state.
Dead Time	Time (in hh:mm:ss) to elapse before OSPF declares the neighbor dead.
Address	Address of next hop.
Interface	Interface name of next hop.
Neighbor is up	Time (in hh:mm:ss) that the OSPF neighbor has been up.
Options	Hello packet options field contents (E-bit only; possible values are 0 and 2; 2 indicates area is not a stub; 0 indicates area is a stub)
Dead timer	Time (in hh:mm:ss) to elapse before OSPF declares the neighbor dead.
Neighbor is up	Amount of time (in hh:mm:ss) that the OSPF neighbor has been up.
Number of DBD retrans	Number of re-sent database description packets.
Index	Index and the remaining lines of this command give detailed information about flooding information received from the neighbor.

The following samples are from output from the **show ospf neighbor detail** command:

```
RP/0/RP0/CPU0:router# show ospf neighbor detail

 Neighbor 192.168.199.137, interface address 172.31.80.37
    In the area 0.0.0.0 via interface HundredGigE 0/3/0/2
    Neighbor priority is 1, State is FULL, 6 state changes
    DR is 0.0.0.0 BDR is 0.0.0.0
    Options is 0x2
    Dead timer due in 0:00:32
    Neighbor is up for 18:45:30
```

```
    Number of DBD retrans during last exhange 0
    Index 1/1, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum 0 msec

Total neighbor count: 1


 Neighbor 10.1.1.1, interface address 192.168.13.1
    In the area 0 via interface HundredGigE 0/3/0/1
    Neighbor priority is 1, State is FULL, 10 state changes
    DR is 0.0.0.0 BDR is 0.0.0.0
    Options is 0x52
    LLS Options is 0x1 (LR)
    Dead timer due in 00:00:36
    Neighbor is up for 1w2d
    Number of DBD retrans during last exchange 0
    Index 3/3, retransmission queue length 0, number of retransmission 5
    First 0(0)/0(0) Next 0(0)/0(0)
    Last retransmission scan length is 1, maximum is 1
    Last retransmission scan time is 0 msec, maximum is 0 msec

 Neighbor 10.4.4.4, interface address 192.168.34.4
    In the area 0 via interface HundredGigE 0/3/0/2
    Neighbor priority is 1, State is FULL, 48 state changes
    DR is 0.0.0.0 BDR is 0.0.0.0
    Options is 0x12
    LLS Options is 0x1 (LR)
    Dead timer due in 00:00:30
    Neighbor is up for 00:40:03
    Number of DBD retrans during last exchange 0
    Index 2/2, retransmission queue length 0, number of retransmission 6
    First 0(0)/0(0) Next 0(0)/0(0)
    Last retransmission scan length is 0, maximum is 1
    Last retransmission scan time is 0 msec, maximum is 0 msec
```

This table describes the significant fields shown in the display.

Table 18: show ospf neighbor detail Field Descriptions

Field	Description
Neighbor	Neighbor router ID.
interface address	IP address of the interface.
In the area	Area and interface through which the OSPF neighbor is known.
Neighbor priority	Router priority of neighbor and neighbor state.
State	OSPF state.
state changes	Number of state changes for this neighbor.
DR is	Neighbor ID of the designated router.
BDR is	Neighbor ID of the backup designated router.

Field	Description
Options	Hello packet options field contents. (E-bit only; possible values are 0 and 2; 2 indicates that the area is not a stub; 0 indicates that the area is a stub).)
LLS Options is 0x1 (LR)	Neighbor is NFS Cisco capable.
Dead timer	Time (in hh:mm:ss) to elapse before OSPF declares the neighbor dead.
Neighbor is up	Amount of time (in hh:mm:ss) that the OSPF neighbor has been up.
Number of DBD retrans	Number of re-sent database description packets.
Index	Index and the remaining lines of this command give detailed information about flooding information received from the neighbor.

The following example displays the traffic engineering support, hostname, and graceful restart and helper capabilities of neighboring connected routers, which are enabled for Segment Routing.

```
Router#show ospf neighbor capabilities
Mon Apr 29 08:27:07.289 UTC

Capabilities advertised in RI LSA
  TE - Traffic Engineering
  GR - Graceful Restart
  GH - Graceful Restart Helper

  NA - Capabilities TLV Not Advertised

Neighbors for OSPF 1

Area 1
Neighbor ID            Address              Interface       SR-Enabled    Capabilities
192.168.199.137        209.165.200.225/27       Gi100/0/0/0     Yes           TE, GR, GH
    HostName: Device_3
192.168.48.1           209.165.200.226/27        Gi100/0/0/4     No            NA
192.168.48.200         209.165.200.227/27       Gi10/0/0/0      Yes           TE, GR, GH
    HostName: Device_0

Total neighbor count: 3
```

This table describes the significant fields shown in the display.

Table 19: show ospf neighbor capabilties Field Descriptions

Field	Description
SR-Enabled	Indicates whether a neighbor router is enabled for Segment Routing or not. The valid values are Yes and No.

Field	Description
Capabilities	Includes the capabilities advertised in the RI LSA of a neighbor router. In general, Traffic Engineering (TE), Graceful Restart (GR), and Graceful Restart Helper (GH) capabilities are advertised. By default, routers advertise the GH capability in the RI LSA. To enable the TE and GR capabilities on neighboring routers, you require configuration at either the area or OSPF level. When Segment Routing is not enabled for a neighbor router, it does not advertise any capabilities in the RI LSA, which is indicated by NA.
Hostname	Indicates the hostname of the neighbor router. If the hostname is not configured, the RI LSA uses the default hostname of the neighbor router.
Neighbor ID	Router ID of the neighbor router.
Address	IP address of the neighbor router.
Interface	Interface IP address of the neighbor router.

The following example shows the capabilities of neighbor 192.168.199.130 listed in the summary information:

```
Router#show ospf neighbor 192.168.199.130
Mon Nov 6 19:13:08.861 IST
* Indicates MADJ interface
# Indicates Neighbor awaiting BFD session up

Neighbors for OSPF 1

 Neighbor 192.168.199.130, interface address 172.31.80.37
    In the area 0 via interface GigabitEthernet0/2/0/4
    Neighbor priority is 1, State is FULL, 6 state changes
    DR is 0.0.0.0 BDR is 0.0.0.0
    Options is 0x52
    LLS Options is 0x1 (LR)
    Dead timer due in 00:00:32
    Neighbor is up for 00:08:50
    Number of DBD retrans during last exchange 0
    Index 3/6, retransmission queue length 0, number of retransmission o
    First 0(0) /0(0) Next 0(0) /0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 1 msec, maximum is 8 msec
    LS Ack list: NSR-sync pending 0, high water mark o
    Adjacency SIDs:
        Label: 24010,    Dynamic, Unprotected
        Label: 24011,    Dynamic,   Protected (Has backup)
    Neighbor Interface ID: 4
    Segment-Routing Enabled
      Strict SPF capable
      Start SID: 16000 Range Size: 8000
      Flex algo Enabled
    Capabilities advertised in RI LSA
      Graceful Restart capable
```

```
      Graceful Restart Helper capable
HostName: r5

Total neighbor count: 1
```

show ospf request-list

To display the first ten link-state requests pending that the local router is making to the specified Open Shortest Path First (OSPF) neighbor and interface, use the **show ospf request-list** command in XR EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **request-list** [*type interface-path-id*] [*neighbor-id*]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
vrf	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.
all	(Optional) Specifies all OSPF VRF instances.
area-id	(Optional) Area ID. If you do not specify an area, all areas are displayed.
type	Interface type. For more information, use the question mark (?) online help function.
i *nterface-path-id*	Physical interface or virtual interface. Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.
neighbor-id	(Optional) IP address of the OSPF neighbor.

Command Default

All neighbors

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

You might use this command when the databases of two neighboring routers are out of synchronization or if the adjacency does not form between them. Adjacency means that the routers synchronize their databases when they discover each other.

You can look at the list to determine if one router is trying to request a particular database update. Entries that are suspended in the list usually indicate that updates are not being delivered. One possible reason for this behavior is a maximum transmission unit (MTU) mismatch between the routers.

You might also look at this list to make sure it is not corrupted. The list should refer to database entries that actually exist.

Request list information is transient and normally the lists are empty.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospf request-list** command:

```
RP/0/RP0/CPU0:router# show ospf request-list 10.0.124.4 HundredGigE 3/0/0/0

Request Lists for OSPF pagent

 Neighbor 10.0.124.4, interface HundredGigE 3/0/0/0 address 10.3.1.2

 Type   LS ID            ADV RTR        Seq NO        Age   Checksum
    1   192.168.58.17     192.168.58.17   0x80000012     12   0x0036f3
    2   192.168.58.68     192.168.58.17   0x80000012     12   0x00083f
```

This table describes the significant fields shown in the display.

Table 20: show ospf request-list 10.0.124.4 HundredGigE 3/0/0/0 Field Descriptions

Field	Description
Neighbor	Specific neighbor receiving the request list from the local router.
Interface	Specific interface over which the request list is being sent.
Address	Address of the interface over which the request list is being sent.
Type	Type of link-state advertisement (LSA).
LS ID	Link-state ID of the LSA.
ADV RTR	IP address of the advertising router.
Seq NO	Sequence number of the LSA.
Age	Age of the LSA (in seconds).
Checksum	Checksum of the LSA.

show ospf retransmission-list

To display the first ten link-state entries in the Open Shortest Path First (OSPF) retransmission list that the local router sends to the specified neighbor over the specified interface, use the **show ospf retransmission-list** command in XR EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] [*area-id*] **retransmission-list** [*type interface-path-id*] [*neighbor-id*]

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
vrf *vrf-name* **all**	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.
area-id	(Optional) Area ID. If you do not specify an area, all areas are displayed.
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.
neighbor-id	(Optional) IP address of the OSPF neighbor.

Command Default

All neighbors

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

You might use this command when the databases of two neighboring routers are out of synchronization or if the adjacency is not forming between them. Adjacency means that the routers synchronize their databases when they discover each other.

You can look at the list to determine if one router is trying to request a particular database update. Entries that appear to be suspended in the list usually indicate that updates are not being delivered. One possible reason for this behavior is a maximum transmission unit (MTU) mismatch between the routers.

You might also look at this list to make sure it is not corrupted. The list should refer to database entries that actually exist.

Retransmission list information is transient, and normally the lists are empty.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospf retransmission-list** command:

```
RP/0/RP0/CPU0:router# show ospf retransmission-list 10.0.124.4 HundredGigE 3/0/0/0

 Neighbor 10.0.124.4, interface HundredGigE 3/0/0/0 address 10.3.1.2
```

This table describes the significant fields shown in the display.

Table 21: show ospf retransmission-list 10.0.124.4 GigabitEthernet3/0/0/0 Field Descriptions

Field	Description
Neighbor	Specified neighbor receiving the retransmission list from the local router.
Interface	Specified interface over which the retransmission list is being sent.
Address	Address of the interface.

show ospf routes

To display the Open Shortest Path First (OSPF) topology table, use the **show ospf routes** command in XR EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] **routes** [**connected** | **external** | **local**] [*prefix mask*] [*prefix/length*] **[multicast-intact] [backup-path]**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
vrf *vrf-name* **all**	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.
connected	(Optional) Displays connected routes.
external	(Optional) Displays routes redistributed from other protocols.
local	(Optional) Displays the local routes redistributed from the Routing Information Base (RIB).
prefix	(Optional) IP prefix, which limits output to a specific route. If the *prefix* argument is specified, either the *length* or *mask* argument is required.
mask	(Optional) IP address mask.
/ *length*	(Optional) Prefix length, which can be indicated as a slash (/) and number. For example, /8 indicates that the first eight bits in the IP prefix are network bits. If *length* is used, the slash is required.
multicast-intact	(Optional) Displays multicast intact paths.
backup-path	(Optional) Displays fast-reroute backup path information.

Command Default

All route types

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **show ospf routes** command to display the OSPF private routing table (which contains only routes calculated by OSPF). If there is something wrong with a route in the RIB, then it is useful to check the OSPF copy of the route to determine if it matches the RIB contents. If it does not match, there is a synchronization problem between OSPF and the RIB. If the routes match and the route is incorrect, OSPF has made an error in its routing calculation.

Task ID

Task ID	Operations
ospf	read

show ospf routes command output with TI-LFA information

This is sample output from the **show ospf routes** command with **backup-path** keyword that displays backup-path information, including TI-LFA:

```
RP/0/RP0/CPU0:routersh ospf 1 routes 2.2.2.2/32 backup-path
Fri Apr  4 02:08:04.210 PDT

Topology Table for ospf 1 with ID 1.1.1.1

Codes: O - Intra area, O IA - Inter area
     O E1 - External type 1, O E2 - External type 2
     O N1 - NSSA external type 1, O N2 - NSSA external type 2

O    2.2.2.2/32, metric 3
     10.1.0.2, from 2.2.2.2, via HundredGigE 0/0/0/7, path-id 1
         Backup path: TI-LFA, P node: 4.4.4.4, Labels: 16004, 123
            10.0.3.2, from 2.2.2.2, via HundredGigE 0/0/0/3, protected bitmap 0x1
            Attribues: Metric: 104, SRLG Disjoint
```

This table describes the significant fields shown in the display.

Table 22: show ospf route Field Descriptions

Field	Description
O	OSPF route.
E	External Type 1 or 2 route.
N	NSSA Type 1 or 2
2.2.2.2/32	Network and subnet mask to which the local router has a route.
metric	Cost to reach network 10.3.1.0.
10.1.0.2	Next-hop router on the path to network 10.3.1.0.
from 2.2.2.2	Router ID 172.16.10.1 is the router that advertised this route.
via GigabitEthernet0/0/0/7	Packets destined for the given prefix (10.3.1.0/24) are sent over GigabitEthernet interface 0/0/0/7.
Backup path	Indicates the topology independent loop-free alternate backup path. Here, the backup path uses the P node 4.4.4.4.

Examples

The following is sample output from the **show ospf routes** command:

```
RP/0/RP0/CPU0:router# show ospf routes
```

```
Topology Table for ospf 1 with ID 10.3.4.2

Codes:O - Intra area, O IA - Inter area
       O E1 - External type 1, O E2 - External type 2
       O N1 - NSSA external type 1, O N2 - NSSA external type 2

O E2 10.3.1.0/24, metric 1
       10.3.4.1, from 172.16.10.1, via HundredGigE 0/1/0/1
O    10.3.4.0/24, metric 1562
       10.3.4.2, directly connected, via HundredGigE 0/1/0/1
O E2 10.1.0.0/16, metric 1
       10.3.4.1, from 172.16.10.1, via HundredGigE 0/1/0/1
O IA 10.10.10.0/24, metric 1572
       10.3.4.1, from 172.16.10.1, via HundredGigE 0/1/0/1
O E2 130.10.10.0/24, metric 20
       10.3.4.1, from 172.16.10.1, via HundredGigE 0/1/0/1
```

This table describes the significant fields shown in the display.

Table 23: show ospf route Field Descriptions

Field	Description
O	OSPF route.
E	External Type 1 or 2 route.
N	NSSA Type 1 or 2
10.3.1.0/24	Network and subnet mask to which the local router has a route.
metric	Cost to reach network 10.3.1.0.
10.3.4.1	Next-hop router on the path to network 10.3.1.0.
from 172.16.10.1	Router ID 172.16.10.1 is the router that advertised this route.
via GigabitEthernet 0/1/0/1	Packets destined for the given prefix (10.3.1.0/24) are sent over GigabitEthernet interface 0/1/0/1.

This table describes the significant fields shown in the display.

Table 24: show ospf route Field Descriptions

Field	Description
O	OSPF route.
E2	External Type 2 route.
10.3.1.0/24	Network and subnet mask to which the local router has a route.
metric 1	Cost to reach network 10.3.1.0.
10.3.4.1	Next-hop router on the path to network 10.3.1.0.

Field	Description
from 172.16.10.1	Router ID 172.16.10.1 is the router that advertised this route.
via POS 0/1/0/1	Packets destined for the given prefix (10.3.1.0/24) are sent over POS interface 0/1/0/1.

The following is sample output from the **show ospf routes** command with a process name of 100:

```
RP/0/RP0/CPU0:router# show ospf 100 routes

Topology Table for ospf 100 with ID 172.23.54.14

Codes:O - Intra area, O IA - Inter area
       O E1 - External type 1, O E2 - External type 2
       O N1 - NSSA external type 1, O N2 - NSSA external type 2

O    10.1.5.0/24, metric 1562
       10.1.5.14, directly connected, via HundredGigE 0/3/0/3
O IA 21.0.0.0/24, metric 1572
       10.1.5.12, from 172.23.54.12, via HundredGigE 0/3/0/3
O    10.0.0.0/24, metric 10
       10.0.0.12, directly connected, via HundredGigE 0/2/0/3
```

This table describes the significant fields shown in the display.

Table 25: show ospf 100 route Field Descriptions

Field	Description
O	OSPF route.
IA	Interarea route.
10.1.5.0/24	Network and subnet mask to which the local router has a route.
metric 1562	Cost to reach network 10.1.5.0.
10.1.5.14	Next-hop router on the path to network 10.1.5.0.
from 172.23.54.12	Router ID 172.23.54.12 is the router that advertised this route.
via GigabitEthernet 0/3/0/3	Packets destined for the given prefix (10.3.1.0/24) are sent over GigabitEthernet interface 0/3/0/3.

The following is sample output from the **show ospf routes** command with a prefix of 10.0.0.0 and a length of 24:

```
RP/0/RP0/CPU0:router# show ospf routes 10.0.0.0/24

Topology Table for ospf 100 with ID 172.23.54.14

Codes:O - Intra area, O IA - Inter area
       O E1 - External type 1, O E2 - External type 2
       O N1 - NSSA external type 1, O N2 - NSSA external type 2
```

```
O IA 10.0.0.0/24, metric 1572
       10.1.5.12, from 172.23.54.12, via GigabitEthernet 0/3/0/3
```

This table describes the significant fields shown in the display.

Table 26: show ospf route 10.0.0.0/24 Field Descriptions

Field	Description
O	Route is an OSPF route.
IA	Route to network 10.0.0.0 is an interarea route.
10.0.0.0/24	Network and subnet mask to which the local router has a route.
metric 1572	Cost to reach network 10.0.0.0.
10.1.5.12	IP address of next-hop router on the path to network 10.0.0.0.
from 172.23.54.12	Router ID 172.23.54.12 is the router that advertised this route.
via GigabitEthernet 0/3/0/3	Packets destined for the given prefix (10.0.0.0/24) are sent over GigabitEthernet interface 0/3/0/3.

show ospf sham-links

To display Open Shortest Path First (OSPF) sham-link information, use the **show ospf sham-links** command in XR EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] **sham-links**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
vrf	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance.
vrf-name	(Optional) Name of the OSPF VRF. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.
all	(Optional) Specifies all OSPF VRF instances.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **show ospf sham-links** command to display OSPF sham-link information.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospf sham-links** command:

```
RP/0/RP0/CPU0:router# show ospf 1 vrf vrf_1 sham-links

Sham Links for OSPF 1, VRF vrf_1

Sham Link OSPF_SL0 to address 10.0.0.3 is up
Area 0, source address 10.0.0.1
IfIndex = 185
Run as demand circuit
DoNotAge LSA allowed., Cost of using 1
Transmit Delay is 1 sec, State POINT_TO_POINT,
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:04
Adjacency State FULL (Hello suppressed)
Number of DBD retrans during last exchange 0
Index 2/2, retransmission queue length 0, number of retransmission 0
```

```
First 0(0)/0(0) Next 0(0)/0(0)
Last retransmission scan length is 0, maximum is 0
Last retransmission scan time is 0 msec, maximum is 0 msec
Keychain-based authentication enabled
Key id used is 2
```

This table describes the significant fields shown in the display.

Table 27: show ospf sham-links Field Descriptions

Field	Description
Sham Link OSPF_SL0 to address	Address of the destination endpoint of the sham link.
IfIndex	ifindex associated with the sham link.
Run as demand circuit	Sham link is treated as a demand circuit.
DoNotAge LSA allowed	DoNotAge LSAs are allowed to be flooded over the sham link.
Cost of using	Sham-link cost.
Transmit Delay	Sham-link transmit delay.
State	Sham-link interface state.
Timer intervals configured	Various sham-link interface-related timers.
Hello due in	Time before the next Hello is sent over the sham link.
Adjacency State	State of the adjacency with the neighbor over the sham link.
Number of DBD retrans during last exchange	Number of DBD retransmissions during the last exchange over the sham link.
Index	Area flood index.
retransmission queue length	Retransmission queue length on the sham link.
number of retransmission	Number of retransmissions over the sham-link interface.
First	First flood information.
Next	Next flood information.
Last retransmission scan length is	Last retransmission scan length on the sham-link interface.
maximum is	Maximum retransmission scan length on the sham-link interface.
Last retransmission scan time is	Last retransmission scan time on the sham-link interface.
maximum is 0 msec	Maximum retransmission scan time on the sham-link interface.
Keychain-based authentication enabled	Keychain-based authentication is enabled.
Key id used is	Key ID used.

show ospf summary-prefix

To display Open Shortest Path First (OSPF) aggregated summary address information, use the **show ospf summary-prefix** command in XR EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] **summary-prefix**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
vrf *vrf-name* **all**	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.

Command Default

All summary prefixes

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **show ospf summary-prefix** command if you configured summarization of external routes with the **summary-prefix** command and you want to display configured summary addresses.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospf summary-prefix** command:

```
RP/0/RP0/CPU0:router# show ospf summary-prefix

OSPF Process 1, summary-prefix

10.1.0.0/255.255.0.0 Metric 20, Type 2, Tag 0
```

This table describes the significant fields shown in the display.

Table 28: show ospf summary-prefix Field Descriptions

Field	Description
10.1.0.0/255.255.0.0	Summary address designated for a range of addresses. The IP subnet mask used for the summary route.

Field	Description
Metric	Metric used to advertise the summary routes.
Type	External link-state advertisements (LSA) metric type.
Tag	Tag value that can be used as a “match” value for controlling redistribution through route maps.

show ospf virtual-links

To display parameters and the current state of Open Shortest Path First (OSPF) virtual links, use the **show ospf virtual-links** command in XR EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name* | **all**}] **virtual-links**

Syntax Description

process-name	(Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the **router ospf** command. If this argument is included, only information for the specified routing process is displayed.
vrf *vrf-name* **all**	(Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.

Command Default

All virtual links

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **show ospf virtual-links** command to display useful information for debugging OSPF routing operations.

Task ID

Task ID	Operations
ospf	read

Examples

The following is sample output from the **show ospf virtual-links** command:

```
RP/0/RP0/CPU0:router# show ospf virtual-links

Virtual Link to router 172.31.101.2 is up
Transit area 0.0.0.1, via interface GigabitEthernet 0/3/0/0, Cost of using 10
Transmit Delay is 1 sec, State POINT_TO_POINT
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 0:00:08
Adjacency State FULL
```

This table describes the significant fields shown in the display.

Table 29: show ospf virtual-links Field Descriptions

Field	Description
Virtual Link to router 172.31.101.2 is up	OSPF neighbor and whether the link to that neighbor is up or down.
Transit area 0.0.0.1	Transit area through which the virtual link is formed.
via interface GigabitEthernet 0/3/0/0	Interface through which the virtual link is formed.
Cost of usingusing 10	Cost of reaching the OSPF neighbor through the virtual link.
Transmit Delay is 1 sec	Transmit delay (in seconds) on the virtual link.
State POINT_TO_POINT	State of the OSPF neighbor.
Timer intervals	Various timer intervals (in seconds) configured for the link.
Hello due in 0:00:08	When the next hello message is expected from the neighbor (in hh:mm:ss).
Adjacency State FULL	Adjacency state between the neighbors.

show protocols (OSPF)

To display information about the OSPFv2 processes running on the router, use the **show protocols** command in XR EXEC mode.

show protocols [**afi-all** | **ipv4** | **ipv6**] [**all***protocol*]

Syntax Description

afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies an IPv4 address family.
ipv6	(Optional) Specifies an IPv6 address family.
all	(Optional) Specifies all protocols for a given address family.
protocol	(Optional) Specifies a routing protocol. For the IPv4 address family, the options are: • **bgp** • **isis** • **ospf** • **rip** For the IPv6 address family, the options are: • **bgp** • **isis** • **ospfv3**

Command Default

No default behavior or value

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
ospf	read
rib	read

Examples

The following is an OSPF configuration and the resulting **show protocols ospf** display:

```
RP/0/RP0/CPU0:router#show running router ospf 1
```

```
router ospf 1
 router-id Loopback0
 nsf
 redistribute connected
 redistribute isis 3
 area 0
  mpls traffic-eng
  interface Loopback0
  !
  interface Loopback1
  !
  interface Loopback2
  !
  interface HundredGigE  0/3/0/0
  !
  interface HundredGigE 0/3/0/1
  !
  interface HundredGigE 0/3/0/2
  !
  interface HundredGigE 0/3/0/3
  !
 !
 mpls traffic-eng router-id Loopback0

!

RP/0/RP0/CPU0:router# show protocols ospf
Routing Protocol OSPF 1
  Router Id: 55.55.55.55
  Distance: 110
  Non-Stop Forwarding: Enabled
  Redistribution:
    connected
    isis 3
  Area 0
    MPLS/TE enabled
    HundredGigE 0/3/0/3
    HundredGigE 0/3/0/2
    HundredGigE 0/3/0/1
    HundredGigE 0/3/0/0
    Loopback2
    Loopback0
```

This table describes the significant fields shown in the display.

Table 30: show protocols ospf Field Descriptions

Field	Description
Router Id	ID of the router for this configuration.
Distance	Administrative distance of OSPF routes relative to routes from other protocols.
Non-Stop Forwarding	Status of nonstop forwarding.
Redistribution	Lists the protocols that are being redistributed.
Area	Information about the current area including list of interfaces and the status of Multiprotocol Label Switching traffic engineering (MPLS TE).

snmp context (OSPF)

To specify an SNMP context for an OSPF instance, use the **snmp context** command in XR Config mode or in VRF configuration mode. To remove the SNMP context, use the **no** form of this command.

snmp context *context_name*
no snmp context *context_name*

Syntax Description

context_name	Specifies name of the SNMP context for OSPF instance.

Command Default

SNMP context is not specified.

Command Modes

XR Config mode

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The snmp-server commands need to be configured to perform SNMP request for the OSPF instance. Refer *SNMP Server Commands* module in *System Management Command Reference* for information on using the **snmp-server** commands.

Note To map an SNMP context with a protocol instance, topology or VRF entity, use the **snmp-server context mapping** command. However, the **feature** option of this command does not work with OSPF protocol.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to configure an SNMP context *foo* for OSPF instance *100*:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router ospf 100
RP/0/RP0/CPU0:router(config-ospf)#snmp context foo
```

This example shows how to configure **snmp-server** commands to be used with the **snmp context** command:

```
RP/0/RP0/CPU0:router(config)#snmp-server host 10.0.0.2 traps version 2c public udp-port
1620
RP/0/RP0/CPU0:router(config)#snmp-server community public RW
RP/0/RP0/CPU0:router(config)#snmp-server contact foo
```

```
RP/0/RP0/CPU0:router(config)#snmp-server community-map public context foo
```

This is a sample SNMP context configuration for OSPF instance *100*:

```
snmp-server host 10.0.0.2 traps version 2c public udp-port 1620
snmp-server community public RW
snmp-server contact foo

snmp-server community-map public context foo


router ospf 100
 router-id 2.2.2.2
 bfd fast-detect
 nsf cisco
snmp context foo
 area 0
  interface Loopback1
  !
 !
 area 1
  interface HundredGigE 0/2/0/1
   demand-circuit enable
  !
  interface HundredGigE 0/3/0/0
  !
  interface HundredGigE 0/3/0/1
  !
 !
!
```

snmp trap (OSPF)

To enable SNMP trap for an OSPF instance, use the **snmp trap** command in VRF configuration mode. To disable SNMP trap for the OSPF instance, use the **no** form of this command.

snmp trap
no snmp trap

Syntax Description

This command has no keywords or arguments.

Command Default

Disabled.

Command Modes

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
ospf	read, write

This example shows how to enable SNMP trap for OSPF instance *100* under VRF *vrf-1*:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router ospf 100
RP/0/RP0/CPU0:router(config-ospf)#vrf vrf-1
RP/0/RP0/CPU0:router(config-ospf-vrf)#snmp trap
```

snmp trap rate-limit (OSPF)

To control the number of traps that OSPF sends by configuring window size and the maximum number of traps during that window, use the **snmp trap rate-limit** command in XR Config mode. To disable configuring the window size and maximum number of traps during the window, use the **no** form of this command.

snmp trap rate-limit *window-size max-num-traps*
no snmp trap rate-limit *window-size max-num-traps*

Syntax Description

window-size	Specifies the trap rate limit sliding window size.
max-num-traps	Specifies the maximum number of traps sent in window time.

Command Default

The default window size is 10 seconds and the maximum number of traps sent in the window time is 7.

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
ospf	read,write

Examples

The following example shows how to set the trap rate limit sliding window size to 30 and the maximum number of traps sent to 100:

```
RP/0/RP0/CPU0:router(config)#router ospf 100
RP/0/RP0/CPU0:router(config-ospf)#snmp trap rate-limit 30 100
```

spf prefix-priority (OSPFv2)

To prioritize OSPFv2 prefix installation into the global Routing Information Base (RIB) during Shortest Path First (SPF) run, use the **spf prefix-priority** command in router configuration mode. To return to the system default value, use the **no** form of this command.

spf prefix-priority route-policy *policy-name*
no spf prefix-priority route-policy *policy-name*

Syntax Description

route-policy *policy-name*	Specifies the route policy to apply to OSPFv2 prefix prioritization.

Note
If SPF prefix prioritization is configured, /32 prefixes are no longer preferred by default. To retain the /32 prefixes in higher-priority queues, define the route-policy accordingly.

Command Default

SPF prefix prioritization is disabled.

Command Modes

OSPF router configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

SPF prefix prioritization is disabled, by default. In disabled mode, the /32 prefixes are installed into the global RIB before other prefixes.

If SPF prefix prioritization is enabled, routes are matched against the route-policy criteria and are assigned to the appropriate priority queue based on the spf-priority set. Unmatched prefixes, including the /32 prefixes, are placed in the low-priority queue.

If all /32 prefixes are desired in the high-priority queue or medium-priority queue, configure the following single route map:

```
prefix-set ospf-medium-prefixes
   0.0.0.0/0 ge 32
end-set
```

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure OSPFv2 SPF prefix prioritization:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# prefix-set ospf-critical-prefixes
```

```
  RP/0/RP0/CPU0:router(config-pfx)# 66.0.0.0/16
  RP/0/RP0/CPU0:router(config-pfx)# end-set
  RP/0/RP0/CPU0:router(config)# route-policy ospf-spf-priority
  RP/0/RP0/CPU0:router(config-rpl)# if destination in ospf-critical-prefixes then set
spf-priority critical
  endif
  RP/0/RP0/CPU0:router(config-rpl)# end-policy
  RP/0/RP0/CPU0:router(config)# router ospf 1
  RP/0/RP0/CPU0:router(config-ospf)# router-id 66.0.0.1
  RP/0/RP0/CPU0:router(config-ospf)# spf prefix-priority route-policy ospf-spf-priority
```

stub (OSPF)

To define an area as a stub area, use the **stub** command in area configuration mode. To disable this function, use the **no** form of this command.

stub [**no-summary**]
no stub

Syntax Description

no-summary	(Optional) Prevents an Area Border Router (ABR) from sending summary link advertisements into the stub area.

Command Default

No stub area is defined.

Command Modes

Area configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

You must configure the **stub** command on all routers in the stub area.

Use the **default-cost** command on the ABR of a stub area to specify the cost of the default route advertised into the stub area by the ABR.

To further reduce the number of link-state advertisements (LSAs) sent into a stub area, you can configure the **no-summary** keyword on the ABR to prevent it from sending summary LSAs (LSA Type 3) into the stub area.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to assign a default cost of 20 to stub network 10.0.0.0:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 201
RP/0/RP0/CPU0:router(config-ospf)# area 10.0.0.0
RP/0/RP0/CPU0:router(config-ospf-ar)# stub
RP/0/RP0/CPU0:router(config-ospf-ar)# default-cost 20
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/3/0/3
```

summary-prefix (OSPF)

To create aggregate addresses for routes being redistributed from another routing protocol into the Open Shortest Path First (OSPF) protocol, use the **summary-prefix** command in the appropriate mode. To stop summarizing redistributed routes, use the **no** form of the command.

summary-prefix *address mask* [**not-advertise** | **tag** *tag*]
no summary-prefix *address mask*

Syntax Description

address	Summary address designated for a range of addresses.
mask	IP subnet mask used for the summary route.
not-advertise	(Optional) Suppresses summary routes that match the address and mask pair from being advertised.
tag *tag*	(Optional) Tag value that can be used as a "match" value for controlling redistribution through route policies.

Command Default

When this command is not used, specific addresses are created for each route from another route source being distributed into the OSPF protocol.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **summary-prefix** command to cause an OSPF Autonomous System Boundary Router (ASBR) to advertise one external route as an aggregate for all redistributed routes that are covered by the address. This command summarizes only routes from other routing protocols that are being redistributed into OSPF.

You can use this command multiple times to summarize multiple groups of addresses. The metric used to advertise the summary is the lowest metric of all the more specific routes. This command helps reduce the size of the routing table.

If you want to summarize routes between OSPF areas, use the **range** command.

Task ID

Task ID	Operations
ospf	read, write

Examples

In the following example, summary address 10.1.0.0 includes address 10.1.1.0, 10.1.2.0, 10.1.3.0, and so on. Only the address 10.1.0.0 is advertised in an external link-state advertisement.

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 201
RP/0/RP0/CPU0:router(config-ospf)# summary-prefix 10.1.0.0 255.255.0.0
```

timers lsa group-pacing

To change the interval at which Open Shortest Path First (OSPF) link-state advertisements (LSAs) are collected into a group and refreshed, checksummed, or aged, use the **timers lsa group-pacing** command in the appropriate mode. To restore the default value, use the **no** form of this command.

timers lsa group-pacing *seconds*
no timers lsa group-pacing

Syntax Description

seconds	Interval (in seconds) at which LSAs are grouped and refreshed, checksummed, or aged. Range is 10 seconds to 1800 seconds.

Command Default

seconds : 240 seconds

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

OSPF LSA group pacing is enabled by default. For typical customers, the default group pacing interval for refreshing, checksumming, and aging is appropriate and you need not configure this feature.

The duration of the LSA group pacing is inversely proportional to the number of LSAs the router is handling. For example, if you have approximately 10,000 LSAs, decreasing the pacing interval would benefit you. If you have a very small database (40 to 100 LSAs), increasing the pacing interval to 10 to 20 minutes might benefit you slightly.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to change the OSPF pacing between LSA groups to 60 seconds:

```
RP/0/RP0/CPU0:router# configure
 RP/0/RP0/CPU0:router(config)# router ospf 1
 RP/0/RP0/CPU0:router(config-ospf)# timers lsa group-pacing 60
```

timers lsa min-arrival

To limit the frequency that new instances of any particular Open Shortest Path First (OSPF) link-state advertisements (LSAs) can be accepted during flooding, use the **timers lsa min-arrival** command in the appropriate mode. To restore the default value, use the **no** form of this command.

timers lsa min-arrival *milliseconds*

no timers lsa min-arrival

Syntax Description

milliseconds	Minimum interval (in milliseconds) between accepting same LSA. Range is 0 to 600000 milliseconds.

Command Default

milliseconds : 100 milliseconds

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to change the minimum interval between accepting the same LSA to 2 seconds:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# timers lsa min-arrival 2
```

timers throttle lsa all (OSPF)

To modify the Open Shortest Path First (OSPF) link-state advertisement (LSA) throttling, use the **timers throttle lsa all** command in the appropriate mode. To revert LSA throttling to default settings, use the **no** form of this command

timers throttle lsa all *start-interval hold-interval max-interval*
no timers throttle lsa all

Syntax Description

start-interval	Delay to generate first occurance of LSA in milliseconds. Range is 0 to 600000 milliseconds.
hold-interval	Minimum delay between originating the same LSA in milliseconds. Range is 1 to 600000 milliseconds.
max-interval	Maximum delay between originating the same LSA in milliseconds. Range is 1 to 600000 milliseconds.

Command Default

start-interval : 50 milliseconds

hold-interval : 200 milliseconds

max-interval : 5000 milliseconds

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The *lsa-start* time is the delay before flooding the first instance of an LSA. The *lsa-hold* interval is the minimum time to elapse before flooding an updated instance of an LSA. The *lsa-max-wait* time is the maximum time that can elapse before flooding an updated instance of an LSA.

For quick convergence, use smaller times for the *lsa-start* time and *lsa-hold* interval. However, in relatively large networks, this may result in a large number of LSAs being flooded in a relatively short time. A balance with the *lsa-start* time and *lsa-hold* interval can be iteratively arrived at for the size of your network. The *lsa-max-wait* time can be used to ensure that OSPF reconverges within a reasonable amount of time.

Note LSA throttling is always enabled. You can change the timer values with the **timers throttle lsa all** command or specify the **no** keyword to revert back to the default settings.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to change the start, hold, and maximum wait interval values to 500, 1000, and 90,000 milliseconds, respectively:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# timers throttle lsa all 500 1000 90000
```

The following example is output from the show ospf command that displays the modified LSA throttle settings:

```
RP/0/RP0/CPU0:router# show ospf

Routing Process "ospf 1" with ID 1.1.1.1
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 It is an area border router
 Initial SPF schedule delay 5000 msecs
 Minimum hold time between two consecutive SPFs 10000 msecs
 Maximum wait time between two consecutive SPFs 10000 msecs
 Initial LSA throttle delay 500 msecs
 Minimum hold time for LSA throttle 1000 msecs
 Maximum wait time for LSA throttle 90000 msecs
 Minimum LSA interval 1000 msecs. Minimum LSA arrival 1 secs
 Maximum number of configured interfaces 255
 Number of external LSA 0. Checksum Sum 00000000
 Number of opaque AS LSA 0. Checksum Sum 00000000
 Number of DCbitless external and opaque AS LSA 0
 Number of DoNotAge external and opaque AS LSA 0
 Number of areas in this router is 2. 2 normal 0 stub 0 nssa
 External flood list length 0
 Non-Stop Forwarding enabled
    Area BACKBONE(0) (Inactive)
        Number of interfaces in this area is 2
        SPF algorithm executed 8 times
        Number of LSA 2. Checksum Sum 0x01ba83
        Number of opaque link LSA 0. Checksum Sum 00000000
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0
    Area 1
        Number of interfaces in this area is 1
        SPF algorithm executed 9 times
        Number of LSA 2. Checksum Sum 0x0153ea
        Number of opaque link LSA 0. Checksum Sum 00000000
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0
```

timers throttle spf (OSPF)

To modify the Open Shortest Path First (OSPF) shortest path first (SPF) throttling, use the **timers throttle spf** command in the appropriate mode. To revert SPF throttling to default settings, use the **no** form of this command.

timers throttle spf *spf-start spf-hold spf-max-wait*
no timers throttle spf

Syntax Description

spf-start	Initial SPF schedule delay (in milliseconds). Range is 1 to 600000 milliseconds.
spf-hold	Minimum hold time (in milliseconds) between two consecutive SPF calculations. Range is 1 to 600000 milliseconds.
spf-max-wait	Maximum wait time (in milliseconds) between two consecutive SPF calculations. Range is 1 to 600000 milliseconds.

Command Default

spf-start:50 milliseconds

spf-hold: 200 milliseconds

spf-max-wait: 5000 milliseconds

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The *spf-start* time is the delay before running SPF for the first time. The *spf-hold* interval is the minimum time to elapse between subsequent SPF runs. The *spf-max-wait* time is the maximum time that can elapse before running SPF again.

Tip Setting a low *spf-start* time and *spf-hold* time causes routing to switch to the alternate path more quickly if there is a failure; however, it consumes more CPU processing time.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to change the start, hold, and maximum wait interval values to 5, 1000, and 90000 milliseconds, respectively:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# timers throttle spf 5 1000 90000
```

transmit-delay (OSPF)

To set the estimated time required to send a link-state update packet on the interface, use the **transmit-delay** command in the appropriate mode. To return to the default value, use the **no** form of this command.

transmit-delay *seconds*
no transmit-delay *seconds*

Syntax Description

seconds	Time (in seconds) required to send a link-state update. Range is 1 to 65535 seconds.

Command Default

seconds: 1 second

Command Modes

XR Config mode

Area configuration

Interface configuration

Virtual-link configuration

VRF configuration

Multi-area configuration

Sham-link configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Link-state advertisements (LSAs) in the update packet must have their ages incremented by the amount specified in the *seconds* argument before transmission. The value assigned should take into account the transmission and propagation delays for the interface.

If the delay is not added before transmission over a link,the time in which the LSA propagates over the link is not considered. This setting has significance only on very low-speed networks not supported in Cisco IOS XR software or on networks such as satellite circuits that incur a very long (greater than one second) delay time.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure a transmit delay for interface HundredGigE 0/3/0/0:

```
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# area 0
RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE 0/3/0/0
```

```
RP/0/RP0/CPU0:router(config-ospf-ar-if)# transmit-delay 3
```

virtual-link (OSPF)

To define an Open Shortest Path First (OSPF) virtual link, use the **virtual-link** command in area configuration mode. To remove a virtual link, use the **no** form of this command.

virtual-link *router-id*
no virtual-link *router-id*

Syntax Description

router-id	Router ID associated with the virtual link neighbor. The router ID appears in the **show ospf** command display. The router ID can be any 32-bit router ID value specified in four-part, dotted-decimal notation.

Command Default

No virtual links are defined.

Command Modes

Area configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

All areas in an OSPF autonomous system must be physically connected to the backbone area (area 0). In some cases in which this physical connection is not possible, you can use a virtual link to connect to the backbone through a nonbackbone area. You can also use virtual links to connect two parts of a partitioned backbone through a nonbackbone area. The area through which you configure the virtual link, known as a transit area, must have full routing information. The transit area cannot be a stub or not-so-stubby area.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to establish a virtual link with default values for all optional parameters:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 201
RP/0/RP0/CPU0:router(config-ospf)# area 10.0.0.0
RP/0/RP0/CPU0:router(config-ospf-ar)# virtual-link 10.3.4.5
RP/0/RP0/CPU0:router(config-ospf-ar-vl)#
```

The following example shows how to establish a virtual link with clear text authentication called mykey:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 201
RP/0/RP0/CPU0:router(config-ospf)# area 10.0.0.0
```

```
RP/0/RP0/CPU0:router(config-ospf-ar)# virtual-link 10.3.4.5
RP/0/RP0/CPU0:router(config-ospf-ar-vl)# authentication-key 0 mykey
```

vrf (OSPF)

To configure an Open Shortest Path First (OSPF) VPN routing and forwarding (VRF) instance, use the **vrf** command in XR Config mode. To terminate an OSPF VRF, use the **no** form of this command.

vrf *vrf-name*
no vrf *vrf-name*

Syntax Description

vrf-name	Identifier of an OSPF VRF. The *vrf-name* argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names.

Command Default

No OSPF VRF is defined.

Command Modes

Router configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

Use the **vrf** command to explicitly configure a VRF. Commands configured under the VRF configuration mode (such as the **interface** [OSPF] and **authentication** commands) are automatically bound to that VRF.

To modify or remove the VRF, the *vrf-id* argument format must be the same as the format used when creating the area.

Note To remove the specified VRF from the router configuration, use the **no vrf** *vrf-id* command. The **no vrf** *vrf-id* command removes the VRF and all VRF options, such as **authentication** , **default-cost** , **nssa** , **range** , **stub** , **virtual-link** , and **interface.**

To avoid possibly having the router ID change under a VRF, explicitly configure the router ID using the **router-id** command.

Task ID

Task ID	Operations
ospf	read, write

Examples

The following example shows how to configure VRF vrf1 and HundredGigE interface 0/2/0/0. HundredGigE interface 0/2/0/0 is bound to VRF vrf1 automatically.

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# vrf vrf1
RP/0/RP0/CPU0:router(config-ospf-vrf)# area area1
```

```
RP/0/RP0/CPU0:router(config-ospf-vrf-ar)# interface HundredGigE 0/2/0/0
```