

Chapter 29

Open Shortest Path First (OSPF)

Introduction	29-3
OSPF Features	29-3
Adjacency and Designated Routers	29-5
Link State Advertisements	29-5
OSPF Packet Types	29-6
OSPF States	29-7
Automatic Cost Calculation	29-8
Routing with OSPF	29-9
Filtering OSPF Routes	29-9
Addressless Interfaces	29-10
Network Types	29-10
Passive Interfaces	29-12
Authenticating OSPF	29-13
Password Authentication	29-13
Cryptographic Authentication	29-13
Exchanging Information Between OSPF and RIP	29-14
Importing BGP routes into OSPF	29-15
Redistributing External Routes	29-15
Summarising Routes for Redistribution	29-17
OSPF On-Demand Circuits	29-18
Configuration Examples	29-20
Basic OSPF Network	29-20
OSPF Network with Addressless PPP Links	29-22
OSPF Network with Virtual Links	29-23
Command Reference	29-27
add ospf area	29-28
add ospf host	29-30
add ospf interface	29-31
add ospf md5key	29-35
add ospf neighbour	29-36
add ospf range	29-37
add ospf redistribute	29-38
add ospf stub	29-40
add ospf summaryaddress	29-41
delete ospf area	29-42
delete ospf host	29-42
delete ospf interface	29-43
delete ospf md5key	29-44
delete ospf neighbour	29-45
delete ospf range	29-45
delete ospf redistribute	29-46

delete ospf stub	29-47
delete ospf summaryaddress	29-48
disable ospf	29-48
disable ospf debug	29-49
disable ospf interface	29-50
disable ospf log	29-50
enable ospf	29-51
enable ospf debug	29-51
enable ospf interface	29-53
enable ospf log	29-54
purge ospf	29-55
reset ospf	29-55
reset ospf counter	29-56
reset ospf interface	29-57
reset ospf spf	29-58
set ospf	29-59
set ospf area	29-63
set ospf host	29-65
set ospf interface	29-66
set ospf neighbour	29-70
set ospf range	29-71
set ospf redistribute	29-72
set ospf stub	29-74
set ospf summaryaddress	29-75
show ospf	29-76
show ospf area	29-79
show ospf debug	29-82
show ospf host	29-83
show ospf interface	29-85
show ospf lsa	29-90
show ospf md5key	29-95
show ospf neighbour	29-96
show ospf range	29-98
show ospf redistribute	29-100
show ospf route	29-101
show ospf stub	29-103
show ospf summaryaddress	29-105

Introduction

This chapter describes the Open Shortest Path First (OSPF) protocol, support for OSPF on the switch, and how to configure the switch to use OSPF as a routing protocol.

Some interface and port types mentioned in this chapter may not be supported on your switch. The interface and port types that are available vary depending on your product's model, and whether an expansion unit (PIC, NSM) is installed. For more information, see the Hardware Reference for the switch.

OSPF Features

Open Shortest Path First (OSPF) is an Interior Gateway Routing Protocol, based on Shortest Path First (SPF) or link-state technology. OSPF is defined in RFCs 1245–1247, 1253 and 1583. OSPF was designed specifically for the TCP/IP Internet environment, and supports the following features:

- Authentication of routing updates.
- Tagging of externally-derived routes.
- Fast response to topology changes with low overhead.
- Load sharing over meshed links.

In SPF-based routing protocols, each router maintains a database describing the Autonomous System's (AS) topology. Each router has an identical database. Each piece of this database describes a particular router and its current state, which includes the state of the interfaces, reachable neighbours, and other information. The router distributes this information about the Autonomous System by "flooding".

Each router runs the algorithm in parallel with other SPF routers, and from the internal database, constructs a tree of shortest paths with itself as the root. The tree contains a route to each destination in the Autonomous System. External routes are added to the tree as "leaves".

OSPF allows the grouping of networks into a set, called an *area*. The topology of an area is hidden from the rest of the Autonomous System. This technique minimizes the routing traffic required for the protocol. When multiple areas are used, each area has its own copy of the topological database.

Another feature of OSPF is that it allows IP subnets to be configured in a very flexible way. Each route distributed by OSPF has a destination and a mask. During the routing process, routes with the longest mask to a destination are used in preference to shorter masks. Host routes are also supported by OSPF; these are considered to be subnets with masks of all ones.

All OSPF protocol exchanges can be authenticated so that only trusted routers participate in the creation of the topology database, and hence the Autonomous System's routing. Authentication is disabled by default.

Externally derived routing data can be passed into the Autonomous System transparently. The externally derived routing information is kept separate from the OSPF protocol's link state data.

Physical networks The following table describes the physical networks OSPF supports.

Network type	Description
Point-to-point	A network of two routers, one at either end of a single connection. The point-to-point interfaces can be set up as numbered or unnumbered interfaces.
Broadcast	A network with potentially more than two routers, and capable of sending a single physical message to all the routers. An example of this type of network is an Ethernet network.
Non-broadcast	A network with potentially more than two routers, but without a mechanism to send a single physical message to all the routers. Examples of these types of network are frame relay or X.25 networks.

The Autonomous System (AS) can be split into multiple areas. This means that routing within the AS takes place on two levels, depending on whether the route to the destination lies entirely within an area (intra-area routing) or in another area (inter-area routing). To link together multiple areas, OSPF uses the concept of a *backbone* that consists of networks and routers linking the other areas. These routers typically have interfaces to the backbone and to other areas. The backbone must be contiguous. *Virtual links* can be used to make the backbone contiguous. Virtual links are links configured between any two backbone routers through a non-backbone area. The backbone itself has all the properties of an area.

When a packet must be routed between two areas, the backbone is used. The packet is first routed to the router that is connected both to the originating area and to the backbone (such a router is called an *Area Border Router*). The packet is then routed through the backbone to another area border router acting for the destination area. The packet is finally routed through the destination area to the specific destination.

The following table summarises the types of OSPF routers.

Router type	Description
internal	Internal Routers route packets within a single area. The internal router can also be a backbone router if that router has no interfaces to other areas.
area border	Area border routers have interfaces in multiple areas, and route packets between these areas. Area border routers condense topological information before passing it to the backbone. This reduces the amount of routing information passed across the backbone.
backbone	A backbone router has an interface on the backbone area.
AS boundary	AS boundary routers exchange routing information with other autonomous systems.

OSPF supports the concept of *stub* areas. External advertisements (external routing information) are not flooded into stub areas. Instead, a single default route is advertised by the Area Border Router. This feature reduces the number of routes the switch needs to store, which may become critical if many external routes are known.

Adjacency and Designated Routers

OSPF creates adjacencies between neighbouring routers. The reason for forming adjacencies is to exchange topological information. Not every router needs to become adjacent to every other router. Adjacencies are established and maintained with *hello* packets. These packets are sent periodically on all router interfaces. Bidirectional communication is determined by a router seeing itself listed in hello packets from its neighbours. On broadcast and non-broadcast multi-access networks, one of the routers becomes a *designated router*. This router performs the following additional tasks:

- **Network Links Advertisement**

The designated router originates the network link state advertisement for the network.

- **Adjacency**

The designated router becomes adjacent to all other routers on the network. Since the topological database is spread over adjacencies, the designated router coordinates the synchronization of the topological database on all the routers attached to the network.

The designated router for a broadcast network is determined dynamically via hello packets. On non-broadcast multi-access networks, static configuration information is used to initiate the search for a designated router. To help in dynamic fallover, OSPF also determines a backup designated router for a network via hello packets. The backup designated router, like the designated router maintains an adjacency to all other routers on the network. If the designated router fails for any reason, the backup designated router takes over.

Link State Advertisements

Link state advertisements are records in the topological database. Routers may generate five different types of link state advertisements ([Table 29-1 on page 29-6](#)). Each type of link state advertisement describes a different set of features of the Autonomous System.

Link state advertisements *age* to a maximum age called *MaxAge* (3600 seconds) while stored in the topological database. When a link state advertisement reaches *MaxAge*, the router tries to flush it from the routing domain by reflooding the advertisement. A link state advertisement that has reached *MaxAge* is not used in further routing table calculations. The *MaxAge* link state advertisement is removed totally from the topological database when it is no longer contained on a neighbour link state retransmission list or none of the neighbours are in exchange or loading state. It is relatively rare for a link state advertisement to reach *MaxAge* because advertisements are usually replaced by more recent instances by normal refresh processes.

Table 29-1: OSPF link state advertisement types

Type	Meaning
Router Links	The router originates a router links advertisement for each area to which it belongs. The advertisement describes the collected states of the router's links to the area. This advertisement also indicates if the router is an area border router or an AS boundary router.
Network Links	A network link advertisement is originated for every transit multi-access network. This advertisement is originated by the designated router for the transit network, and describes all the OSPF routers fully adjacent to the designated router.
Summary Links	Summary Link advertisements describe a single route to a destination. The destinations described are external to the area but internal to the Autonomous System. Some condensing of routing information occurs when creating these summary link state advertisements.
AS Summary Links	These are like summary link advertisements but they describe routes to Autonomous System boundary routers.
AS External Links	AS external advertisements describe routes external to the Autonomous System.

OSPF Packet Types

The OSPF protocol runs directly over IP, using the assigned number 89. The following table describes OSPF packet types.

Type	Purpose
Hello	Used to discover and maintain neighbours.
Database Description	Used to form adjacencies. The router summarises all its link state advertisements and passes this information, via database description packets to the router it is forming an adjacency with.
Link State Request	After the database description packets have been exchanged with a neighbour, the router may detect link state advertisements it requires to update or complete the topological database. Link state request packets are sent to the neighbour requesting these link state advertisements.
Link State Update	Used for transmission of link state advertisements between routers. This could be in response to a link state request packet or to flood a new or more recent link state advertisement.
Link State Acknowledgment	Used to make the flooding of link state advertisements reliable. Each link state advertisement received is explicitly acknowledged.

OSPF States

The following table describes the states that neighbours can be in.

State	Meaning
Down	This is the initial state. No hello packets have been received from the neighbour recently or at all.
Attempt	This state applies to non-broadcast multi-access networks. The router is making a determined attempt to contact a statically configured neighbour. Hello packets are sent every hello interval.
Init	A hello packet has been seen from the neighbour, however the hello packet does not list the router as known.
Two-Way	This state is entered when the communication between two neighbours is bidirectional (the hello packet from the neighbour lists this router as a neighbour).
ExStart	This is the first step in creating an adjacency between two routers. The two routers decide which is going to control the exchange between them.
Exchange	In this state, the neighbours exchange database description packets. Each packet summarises the link state advertisements held by that router.
Loading	After all the database description information has been exchanged, the routers exchange link state advertisements required to update or complete each router's topological database thereby synchronising the two router's databases.
Full	This is the final state and the adjacency is complete. Reaching this state in itself may cause new instances of some link state advertisements, such as the network and router advertisements related to the two routers.

The following table describes interface states the router can be in.

State	Meaning
Down	The initial state. No traffic can be routed with the interface in this state.
Loopback	The router's interface to the network is looped back.
Waiting	Interfaces to broadcast and non-broadcast multi-access networks enter this state when they are started. In this state the router tries to determine the backup designated router. The router is not allowed to elect a backup or designated router while in the waiting state. This stops unnecessary changes in the designated router.
Point-to-point	The interface is operational and is connected to a point-to-point network or virtual link.
DROther	The interface to a broadcast or a non-broadcast multi-access network has not been selected as either the designated router or backup designated router for the network.
Backup	The interface to a broadcast or a non-broadcast multi-access network has been selected as the backup designated router for the network.
DR	The interface to a broadcast or a non-broadcast multi-access network has been selected as the designated router for the network.

The metrics used by OSPF are not simple distance metrics, as used by RIP, but are measures of the bandwidth. Interface metrics should be set using the formula $10^8 / \text{InterfaceSpeed}$. This gives metrics such as 10 for Ethernet and 1562 for a 64 kbps serial line.

Automatic Cost Calculation

OSPF interfaces can automatically set the OSPF metric of an IP interface on the basis of the bandwidth of the interface, instead of the system administrator manually setting the OSPF metric. Automatic setting takes into account that the speed of an interface can change over time, when ports change link state or change speed via auto negotiation or manual setting. If metrics are manually set, some interfaces are preferred when they should not be because the network configuration dynamically changes.

Note that the interface speed used in the cost calculation is the average interface speed. For example, if the interface is a VLAN with two ports up, and one port has a speed of 10 and the other a speed of 100, then the metric will be 18.

To configure auto cost calculation:

1. Do not set the OSPF metric manually with the **add ip interface** command. If you have, remove the manual setting by using the command:

```
set ip interface=int ospfmetric=default
```

The **ospfmetric** parameter specifies the cost of crossing the logical interface, for OSPF. If you specify **default**, the switch restores the interface to the default metric value. If you set the OSPF metric to a value other than **default**, the switch uses the value you specify instead of the OSPF automatic metric setting, even if the automatic setting is enabled via **set ospf autocost=on**.

2. Set **autocost** to **on** and change the reference bandwidth if necessary, by using the command:

```
set ospf autocost=on [refbandwidth=10..10000]
```

The **autocost** parameter specifies whether or not the switch will assign OSPF interface metrics based on the available interface bandwidth. The default is **off**.

The **refbandwidth** parameter specifies the reference bandwidth that the switch uses to calculate the metric, in megabits per second. The cost is calculated as:

$$\text{refbandwidth} / \text{Interface bandwidth}$$

Using the default reference bandwidth of 1000, the automatic cost calculation results in an OSPF metric of 10 for a fast Ethernet (100Mbps) interface.

3. Check the settings, by using the command:

```
show ospf
```


Routing with OSPF

To route an IP packet, the switch looks up the routing table entry that best matches the destination of the packet. This entry contains the interface and nexthop router to forward the IP packet to its destination. The entry that best matches the destination is determined first by the path type, then the longest (most specific) network mask. At this point there may still be multiple routing entries to the destination; if so, then equi-cost multi-path routes exist to the destination. Such equi-cost routes are appropriately used to share the load to the destination.

Path types The following table describes path types.

Path Type	Description
INTRA	Route to the destination is within a single OSPF area.
INTER	Route to the destination is within the Autonomous System but spans more than one OSPF area.
EXT1	Route to the destination is via an AS router within the Autonomous system. This is an OSPF external route of Type 1. Type 1 external routes add the external metric as received by the AS router and the internal OSPF metric to the AS router to determine the final metric to the destination.
EXT2	Route to the destination is via an AS router within the Autonomous system. This is an OSPF external route of Type 2. Type 2 external routes use two metrics to determine how to route traffic to the destination. The first metric is the internal OSPF metric to the AS router, the second metric is the EGP-derived external metric to the destination.

Filtering OSPF Routes

Instead of adding all OSPF routes to the IP route table (the RIB), you can filter to select a subset. This uses route maps, which are described in the following sections of [Chapter 31, Filtering IP Routes](#).

For this information...	See...
what they are and when to use them	"About Route Maps" on page 31-5
how to create an appropriate route map	"Creating Route Maps for OSPF Routes" on page 31-16
how to apply the route map	"Applying Filters When Writing to the RIB" on page 31-22

You can also use route maps to set the metric or route type of some or all routes, and to tag routes with a number.

Addressless Interfaces

OSPF supports addressless interfaces on PPP links. There are a few restrictions, however, to the use of such addressless interfaces. If access to the switch is required via ping or Telnet, then there must be at least one other interface with an assigned Internet address.

A switch generally has at least one LAN interface, and this would usually have an assigned IP address.

A virtual link cannot be configured through an area where the last interface to the area border router is an addressless interface. The reason for this is that to create a virtual link, each end must have a real Internet address to send and receive OSPF packets that are sent as unicast packets. Except for these few restrictions it is prudent to use addressless interfaces on PPP links to save needless use of IP address space.

RFC 1583 provides additional support for addressless interfaces. Addressless interfaces are distinguished from each other by the *ifIndex* of the interface. This data is placed into router LSAs, and is flooded to all OSPF routers. This data is also used to provide next hop information when routes are calculated. The addressless interface index is displayed in the output of the [show ospf lsa command on page 29-90](#), and the output of the [show ip route command on page 23-192 of Chapter 23, Internet Protocol \(IP\)](#). In the **show ip route** command, the next hop for OSPF routes that are down an addressless PPP interface is displayed as the ifIndex in dotted decimal notation. For example, if the addressless interface index for a PPP interface is 5, the next hop appears as 0.0.0.5.

Network Types

OSPF treats the networks attached to OSPF interfaces as one of the following network types, depending on the physical media:

- broadcast
- non-broadcast multi-access (NBMA)
- point-to-point
- point-to-multipoint
- virtual

By default, VLAN and Ethernet networks are treated as broadcast networks. You can configure a VLAN interface as either a broadcast or an NBMA network. Configure a VLAN interface as an NBMA interface when:

- some devices on the network do not support multicast addressing
- you want to select which devices on the network are to become OSPF neighbours, rather than allow all the devices on the network to become OSPF neighbours

Configuring the network type

To add a VLAN interface to OSPF as a broadcast network (the default), use the command:

```
add ospf interface=interface [other-options...]
```

To add a VLAN interface to OSPF and set the network type to NBMA, use the command:

```
add ospf interface=interface network=non-broadcast
[other-options...]
```

To change the network type of an existing a VLAN interface, use the command:

```
set ospf interface=interface
network={broadcast|non-broadcast} [other-options...]
```

To display the network type of an OSPF interface, use the command:

```
show ospf interface=interface full
```

To display the network types of all OSPF interfaces, use the command:

```
show ospf interface full
```

Neighbours on non-broadcast networks

When you change the network type of a VLAN interface from broadcast to non-broadcast:

- All OSPF packets are sent as unicast messages, not broadcast messages, so neighbours need to be statically configured.
- Any existing dynamically learned neighbours are automatically converted to static neighbours, and will appear in any configuration script created by using the [create config command on page 5-22 of Chapter 5, Managing Configuration Files and Software Versions](#).
- Hello messages are not transmitted until at least one static neighbour exists.

You can add, delete or modify static neighbours by using the commands:

```
add ospf neighbour=ipadd priority=0..255
delete ospf neighbour=ipadd
set ospf neighbour=ipadd
```

You can display the list of currently configured static neighbours using the command:

```
show ospf neighbour
```

You can configure the time interval between hello messages sent to neighbours that are deemed to be inactive. To do this, use the **pollinterval** parameter on the [add ospf interface](#) and [set ospf interface](#) commands.

Neighbours on broadcast networks

When you change the network type of a VLAN interface from non-broadcast to broadcast:

- Any existing statically defined neighbours are cleared.
- Hello messages are sent as broadcast messages, so neighbours are dynamically learned.

You can display the list of current neighbours using the command:

```
show ospf neighbour
```

Passive Interfaces

General A passive interface does not take part in normal OSPF interface operations:

- OSPF does not transmit or receive Hello messages via the interface.
- The interface does not experience interface state transitions.
- OSPF does not associate neighbours with the interface.

If the interface is up, OSPF adds the network attached to the interface as a stub network to the router LSA of the area in which the interface resides.

When to use Use a passive interface when you want OSPF neighbours to know about the networks attached to the interface without involving the interface in OSPF protocol exchanges. Using passive interfaces reduces the amount of OSPF protocol traffic the switch transmits and receives.

Per-interface setting To explicitly add an interface to OSPF as a passive interface, use the command:

```
add ospf interface=interface passive={on|yes|true}
[other-options...]
```

To make an existing OSPF interface into a passive interface, use the command:

```
set ospf interface=interface passive={on|yes|true}
[other-options...]
```

To explicitly add an interface to OSPF as a normal interface, use the command:

```
add ospf interface=interface passive={off|no|false}
[other-options...]
```

To make an existing passive interface into a normal OSPF interface, use the command:

```
set ospf interface=interface passive={off|no|false}
[other-options...]
```

To delete a passive interface, use the command:

```
delete ospf interface=interface
```

To display a list of passive interfaces, use the command:

```
show ospf interface full
```

Global setting To set all OSPF interfaces to passive interfaces, use the command:

```
set ospf passiveinterfacedefault={on|yes|true}
[other-options...]
```

To reset all passive interfaces back to normal OSPF interfaces, use the command:

```
set ospf passiveinterfacedefault={off|no|false}
[other-options...]
```

The **passiveinterfacedefault** parameter also sets the default mode for all interfaces subsequently added to OSPF. You can override this for individual interfaces using the **passive** parameter of the **add ospf interface** and **set ospf interface** commands.

Authenticating OSPF

You can authenticate OSPF packets as described in Appendix D of RFC 2328. See this RFC for a detailed description of how OSPF packets are authenticated.

An authentication type can be chosen for each interface. RFC 1583 states that authentication is configured per area, but RFC 2328 states that authentication is configured per interface. This implementation of OSPF authentication represents a compromise for these two solutions. An authentication type is set up per area to act as a default for all interfaces in the area. The default setting for interfaces is to use the area default, but each interface can be individually set to any authentication method.

There are two ways to authenticate an OSPF packet:

- **Password Authentication**
- **Cryptographic Authentication**

Password Authentication

Password authentication can be configured for OSPF areas to authenticate incoming packets. The password can be up to 8 characters long, and is configured for each interface.

To configure an OSPF area with password authentication, use the command:

```
add ospf area={backbone | area-number} authentication=password
```

The password itself is configured on a per-interface basis with the **add ospf interface** and **set ospf interface** commands.

To configure an OSPF interface with password authentication, use the command:

```
add ospf interface=interface authentication=password  
password={none | password}
```

Valid characters for the password are any printable character. If the password contains spaces it must be enclosed in double quotes. If **none** is specified, no password is set on the interface, and any previously set password is removed.

For password authentication to succeed, you need to configure all interfaces in the same physical network with the same password.

Cryptographic Authentication

An MD5 digest can be appended to the OSPF packet for authentication. The digest is based on the contents of the packet and a shared secret key. The key that you configure must be the same for all interfaces sharing the same physical network. MD5 keys are defined for each interface, and can be up to 16 characters long. The key is case-sensitive, and valid characters are letters and digits only. If authentication is set to MD5 and no key is configured, a default key is created that has an ID of 0 and no key.

For MD5 authentication to succeed, first configure the OSPF area, then the interface, and then add the MD5 key to the interface. You can use the **set ospf area** and **set ospf interface** commands to modify the authentication type.

To configure an OSPF area with MD5 authentication, use the command:

```
add ospf area={backbone|area-number} authentication=md5
```

To configure an OSPF interface with MD5 authentication, use the command:

```
add ospf interface=interface authentication=md5
```

To add an MD5 key that will be used for interface authentication, use the command:

```
add ospf md5key=key id=1..255 interface=interface
```

Only one MD5 key is usually added at a time for any given OSPF interface. We recommend changing the MD5 key every month, and periodically deleting old and inactive keys. When changing MD5 keys, add the new key on all routers in the network. While the keys are being added, routers will send duplicate messages using both keys. The old key becomes inactive when all interfaces transmit packets using the new key. The packets will not be duplicated using the inactive key. The output of the **show ospf md5key** command shows whether a key is active or inactive.

Deleting MD5 keys

You can delete an MD5 key when it has become inactive or when it is being used. You may want to delete a key that is currently active if an illicit router is using the key. To delete an active key immediately, specify the **force** parameter in the **add ospf md5key** command.



Caution Forcibly deleting an active MD5 key may lead to partial network failure if succeeding keys are not configured on all interfaces in the physical network.

Before deleting a key, configure a succeeding key on all interfaces in the physical network. You can delete the previously active key without using the **force** parameter. The new key can take over authentication duties when the active key is deleted.

Exchanging Information Between OSPF and RIP

OSPF allows the bidirectional exchange of routing information with RIP.

RIP metrics are based on a simple distance scheme (i.e. the number of routers in the path to the destination), while OSPF metrics are based on bandwidth. To overcome this problem, OSPF treats routes derived from RIP as exterior routing information, and converts them into Type-2 routes. The RIP metric is simply added to the internal OSPF metric. If the RIP metric is infinity the OSPF metric is similarly set to infinity.

To import RIP routes into OSPF, you can use the command:

```
set ospf rip={import|both}
```

However, we recommend that you use route redistribution definitions to import and redistribute RIP routes into OSPF, as they provides more control over how the routes are imported. For more information, see [“Redistributing External Routes” on page 29-15](#).

Exporting OSPF-derived routes out to RIP is more difficult. There is no easy way to convert bandwidth metrics to distance metrics. OSPF simply uses any

OSPF metric below 8 as is, and converts metrics above 8 to 8. OSPF metrics of infinity are converted to RIP infinity 16.

To export OSPF routes to RIP, use the command:

```
set ospf rip={export|both}
```

Then configure RIP as detailed in [Chapter 23, Internet Protocol \(IP\)](#).

Importing BGP routes into OSPF

You can import routes from BGP into OSPF, using the command:

```
set ospf bgpimport=on [bgpfilter={0..999|NONE}]  
[bgplimit=1..4000]
```



Caution BGP can learn thousands of routes. Do not enable importing of BGP routes into OSPF without considering the network impact of importing these routes. Routing devices in the OSPF domain may become overloaded if they store too many routes. You can prevent this by limiting the number of routes to be imported.

We recommend, however, that you use route redistribution definitions to import and redistribute BGP routes into OSPF, as they provides more control over how the routes are imported. For more information, see [“Redistributing External Routes” on page 29-15](#).

Redistributing External Routes

OSPF can import and redistribute BGP, RIP, non-OSPF interface, and statically configured routes. It can also optionally assign any of the following settings to all routes it imports:

- a route metric
- the External metric type (see [“Path types” on page 29-9](#))
- a tag—a number to label the route

Alternatively, you can assign a route map to select particular routes and set their route parameters (see [“Creating Route Maps for OSPF Routes” on page 31-16 of Chapter 31, Filtering IP Routes](#)). The route map can also filter out a subset of routes, so you do not have to import all routes.

The import settings also allow you to select whether to redistribute subnets (classless network routes), or only classfull network routes.

To import and redistribute external routes into OSPF, create a route redistribution definition for the source routing protocol, using the command:

```
add ospf redistribute protocol={bgp|interface|rip|static}  
[other-options...]
```

To delete a route redistribution definition and stop importing routes, use the command:

```
delete ospf redistribute protocol={bgp|interface|rip|static}
```

To change a route redistribution definition, use the command:

```
set ospf redistribute protocol={bgp|interface|rip|static}
[other-options]
```

To display the currently configured route redistribution definitions, use the command:

```
show ospf redistribute
```

Interaction with global OSPF parameters

You can also use the **bgpfilter**, **bgpimport**, **bgplimit**, **rip** and **staticexport** parameters of the [set ospf command on page 29-59](#) to configure OSPF to import BGP, RIP and static routes. However, we recommend that you use route redistribution definitions to import and redistribute routes into OSPF, as they provides more control over how the routes are imported.

For compatibility, the **bgpimport**, **rip** and **staticexport** parameters are synchronised with the equivalent redistribution definition. Changing the setting of these parameters will add or delete the corresponding route redistribution definition, as summarised in the following table.

When you change this set ospf parameter...	From...	To...	Then OSPF...
rip	off or export	import or both	adds a RIP route redistribution definition
	import or both	off or export	deletes the RIP route redistribution definition
bgpimport	off	on	adds a BGP route redistribution definition
	on	off	deletes the BGP route redistribution definition
staticexport	off	on	adds a static route redistribution definition, if asexternal is set to on or nssa
	on	off	deletes the static route redistribution definition, if asexternal is set to on or nssa

Similarly, adding or deleting a route redistribution definition changes the setting of the corresponding **bgpimport**, **rip** or **staticexport** parameter, as summarised in the following table.

When you do this...	Then this parameter...	Changes from...	To...
add a BGP route redistribution definition	bgpimport	off	on
delete a BGP route redistribution definition	bgpimport	on	off
add a RIP route redistribution definition	rip	off or export	import or both
delete a RIP route redistribution definition	rip	import or both	off or export
add a static route redistribution definition	staticexport	off	on
delete a static route redistribution definition	staticexport	on or nssa	off

These changes are also reflected in the output of the [show config](#) and [create config](#) commands:

- If **bgpimport** is set to **on** in the **set ospf** command, then **bgpimport** will be set to **off** (default) in the output, and the corresponding BGP redistribution definition will be added to the output.
- If **rip** is set to **import** in the **set ospf** command, then **rip** will not be written to the output (default is **off**). Instead, the corresponding RIP redistribution definition will be written to the output.

- If **rip** is set to **both** in the **set ospf** command, then **rip** will be set to **export** in the output, and the corresponding RIP redistribution definition will be added to the output.
- If **staticexport** is set to **on** in the **set ospf** command, then **staticexport** will be set to **off** (default) in the output, and the corresponding static redistribution definition will be added to the output.

Summarising Routes for Redistribution

OSPF can summarise external routes to be redistributed using a list of administratively defined summary addresses specified as network/mask pairs. The summary addresses replace the original routes in AS external LSAs. Use summary addresses to reduce the number of AS external routes advertised by the switch.

You can set the following attributes for summary addresses:

- Whether the summary address is advertised.
- The tag to be inserted in the AS external LSA. The tag overrides tags set by the original route and by the route map used to select the original routes for redistribution.

To create summary addresses for route redistribution, use the command:

```
add ospf summaryaddress=ipadd mask=ipadd [other-options...]
```

To delete a summary address, use the command:

```
delete ospf summaryaddress=ipadd/prefix-length
```

To modify a summary address, use the command:

```
set ospf summaryaddress=ipadd mask=ipadd [other-options...]
```

To display the list of configured summary addresses, use the command:

```
show ospf summaryaddress
```

OSPF On-Demand Circuits

OSPF on demand circuits allow data link connections to be closed when not carrying application traffic. It is used on cost-conscious networks, such as ISDN, X.25 and dial-up networks. If there is no traffic crossing the network, (either routing protocol traffic or application traffic), the data link connection is closed. When there is traffic to send, the data link connection is established, the data is sent, and the connection stays open until the link has been idle for a specified time. The data link connection is then closed to conserve cost and resources. [Figure 29-1 on page 29-18](#) and [Figure 29-2 on page 29-19](#) illustrate before and after OSPF on demand scenarios.

Some interface and port types mentioned in this example may not be supported on your switch. The interface and port types that are available vary depending on your product's model, and whether an expansion unit (PIC, NSM) is installed. For more information, see the Hardware Reference for the switch.

Figure 29-1: Example of dial-on-demand ISDN before configuring OSPF on demand





Figure 29-2: Example of dial-on-demand ISDN after configuring OSPF on demand



OSPF on demand is defined in RFC 1793, *Extending OSPF to Support Demand Circuits*. All routers in the network must support at least Part II of the RFC. Routers attached to on-demand links must support Part III of the RFC.

To enable OSPF to be used for routing from Switch B to Switch A and beyond, OSPF on-demand is required. Turning OSPF on-demand on at either end of the ISDN link disables static and RIP routes.

OSPF uses LSAs (Link State Advertisements) to advertise routes. An LSA includes a *Route Age* field. When the *Route Age* field of a route in the OSPF routing database reaches the value **maxage**, it is deleted. OSPF incorporates a process for refreshing active routes so they don't age out. Therefore, only dead, inactive, and obsolete routes are aged out. The refresh process involves hello messages, and these bring up the dial-on-demand ISDN link.

OSPF on demand overcomes this by setting the *DO_NOT_AGE* bit of the *Route Age* field in LSAs received via the dial-on-demand link. Therefore, these LSAs are not aged and do not require the hello process. When Switch B transmits a packet over the dial-on-demand link using such an LSA, a timer is set to MAXAGE. If there is no response before the timer expires, the route is assumed to be dead and is immediately deleted. Switch A can also remove a route it knows is dead from Switch B's LSA database by sending an LSA with the *DO_NOT_AGE* bit unset and the *Route Age* field set to **maxage**.

The *demand circuit* option (DC) is added to the OSPF *Options* field to support the demand circuit extensions. The DC option indicates whether the switch supports OSPF on demand as per the RFC. If Switch C does not support OSPF on demand, it transmits LSAs without the DC bit set. When Switch B unsets the DC bit, it transmits a special LSA to all other routers in the connected areas with the DC bit unset. This causes all conforming routers to stop sending LSAs with the *DO_NOT_AGE* bit set.

Configuration Examples

This section contains examples of how to configure the following common OSPF networks:


- **Basic OSPF Network**
- **OSPF Network with Addressless PPP Links**
- **OSPF Network with Virtual Links**

Basic OSPF Network

This example is a simple network of two switches connected together, each with its own local area network. The switches belong to a single class B network 172.31.0.0, which has been subnetted using the subnet mask 255.255.255.0 (Figure 29-3).

Some interface and port types mentioned in this example may not be supported on your switch. The interface and port types that are available vary depending on the model, and whether an expansion unit (PIC, NSM) is installed. For more information, see the Hardware Reference for the switch.

Figure 29-3: Basic OSPF network



To configure a basic OSPF network

I. Configure interfaces on switch 1.

Create IP interfaces to use the interfaces, and assign an OSPF metric to each IP interface:

```
create ppp=0 over=syn0
enable ip
add ip interface=ppp0 ip=172.31.2.1 mask=255.255.255.0
    ospfmetric=1
add ip interface=eth0 ip=172.31.1.1 mask=255.255.255.0
    ospfmetric=1
```

Assign IP addresses and assign an OSPF metric to each IP interface.

```
enable ip
add ip interface=ppp0 ip=172.31.2.1 mask=255.255.255.0
    ospfmetric=1
add ip interface=eth0 ip=172.31.1.1 mask=255.255.255.0
    ospfmetric=1
```

2. Configure switch 1 as an OSPF router.

Create an OSPF area, assign the IP interfaces to the area, and configure OSPF routing parameters:

```
enable ospf
add ospf area=0.0.0.1 authentication=password
add ospf range=172.31.0.0 area=0.0.0.1 mask=255.255.0.0
add ospf interface=ppp0 area=0.0.0.1 password=asecret
add ospf interface=eth0 area=0.0.0.1 password=bsecret
```

3. Configure interfaces on switch 2.

Create IP interfaces to use the interfaces, and assign an OSPF metric to each IP interface:

```
create ppp=0 over=syn0
enable ip
add ip interface=ppp0 ip=172.31.2.2 mask=255.255.255.0
    ospfmetric=1
add ip interface=eth0 ip=172.31.108.10 mask=255.255.255.0
    ospfmetric=1
```

Assign IP addresses and assign an OSPF metric to each IP interface.

```
enable ip
add ip interface=ppp0 ip=172.31.2.2 mask=255.255.255.0
    ospfmetric=1
add ip interface=eth0 ip=172.31.108.10 mask=255.255.255.0
    ospfmetric=1
```

4. Configure switch 2 as an OSPF router.

Create an OSPF area, assign the IP interfaces to the area, and configure OSPF routing parameters:

```
enable ospf
add ospf area=0.0.0.1 authentication=password
add ospf range=172.31.0.0 area=0.0.0.1 mask=255.255.0.0
add ospf interface=ppp0 area=0.0.0.1 password=csecret
add ospf interface=eth0 area=0.0.0.1 password=bsecret
```


See the [add ospf interface command on page 29-31](#) for a more information about interfaces.

OSPF Network with Addressless PPP Links

This example is similar to the basic network example except that the PPP link has been replaced with an addressless PPP link (Figure 29-4 on page 29-22).

Some interface and port types mentioned in this example may not be supported on your switch. The interface and port types that are available vary depending on your product's model, and whether an expansion unit (PIC, NSM) is installed. For more information, see the Hardware Reference for the switch.

Figure 29-4: OSPF network with an addressless PPP link



To configure an OSPF network with addressless PPP interfaces

1. Configure the PPP and Ethernet interfaces on switch 1.

Create IP interfaces to use the PPP and Ethernet interfaces, and assign an OSPF metric to each IP interface. Note that the PPP interface has no IP address:

```
create ppp=0 over=syn0
enable ip
add ip interface=ppp0 ip=0.0.0.0 mask=0.0.0.0 ospfmetric=1
add ip interface=vlan1 ip=172.31.1.1 mask=255.255.255.0
    ospfmetric=1
```

2. Configure switch 1 as an OSPF router.

Create an OSPF area, assign the IP interfaces to the area, and configure OSPF routing parameters:

```
enable ospf
add ospf area=0.0.0.1 authentication=password
add ospf range=172.31.0.0 area=0.0.0.1 mask=255.255.0.0
add ospf interface=vlan1 area=0.0.0.1 password=asecret
add ospf interface=PPP0 area=0.0.0.1 password=bsecret
```

3. Configure the PPP and Ethernet interfaces on switch 2.

Create IP interfaces to use the PPP and Ethernet interfaces, and assign an OSPF metric to each IP interface. Note that the PPP interface has no IP address:

```
create ppp=0 over=syn0
enable ip
add ip interface=ppp0 ip=0.0.0.0 mask=0.0.0.0 ospfmetric=1
add ip interface=vlan1 ip=172.31.108.10 mask=255.255.255.0
    ospfmetric=1
```

4. Configure switch 2 as an OSPF router.

Create an OSPF area, assign the IP interfaces to the area, and configure OSPF routing parameters:

```
enable ospf
add ospf area=0.0.0.1 authentication=password
add ospf range=172.31.0.0 area=0.0.0.1 mask=255.255.0.0
add ospf interface=vlan1 area=0.0.0.1 password=csecret
add ospf interface=ppp0 area=0.0.0.1 password=bsecret
```


OSPF Network with Virtual Links

This example shows three switches forming a segmented backbone, and connected via two switches in a transit area, using a virtual link. This configuration would join the two backbone areas together via the virtual link using area 1 as the transit area ([Figure 29-5 on page 29-24](#)).

Some interface and port types mentioned in this example may not be supported on your switch. The interface and port types that are available vary depending on your product's model, and whether an expansion unit (PIC, NSM) is installed. For more information, see the Hardware Reference for the switch.

Note that for areas that are not backbone areas, the default setting for the **stubarea** parameter is **on**. However, if an area is the transit area of a virtual link, **stubarea** must be **off**. Therefore in this example, **stubarea** is explicitly set to **off** for all definitions of area 0.0.0.1.

Figure 29-5: OSPF network with a virtual link



To configure an OSPF network with virtual links

1. Configure interfaces on switch 1.

Create IP interfaces to use the interfaces, and assign an OSPF metric to each one:

```
create ppp0 over=syn0
enable ip
add ip int=ppp0 IP=172.31.2.2 mask=255.255.255.0 ospf=1
add ip int=eth0 IP=172.31.1.1 mask=255.255.255.0 ospf=1
```

2. Configure switch 1 as an OSPF router.

Create an OSPF area, assign the IP interfaces to the area, and configure OSPF routing parameters:

```
enable ospf
add ospf area=0.0.0.0 authentication=password
add ospf range=172.31.0.0 area=0.0.0.0 mask=255.255.0.0
add ospf range=172.29.0.0 area=0.0.0.0 mask=255.255.255.0
add ospf interface=ppp0 area=0.0.0.0 password=asecret
add ospf interface=eth0 area=0.0.0.0 password=secret
```

3. Configure interfaces on switch 2.

Create IP interfaces to use the interfaces, and assign an OSPF metric to each IP interface:

```
create ppp0 over=syn0
enable ip
add ip int=ppp0 ip=172.31.2.254 mask=255.255.255.0 ospf=1
add ip int=eth0 ip=172.30.100.4 mask=255.255.255.0 ospf=1
```


4. Configure switch 2 as an OSPF router.

Create two OSPF areas, assign the IP interfaces to the areas, establish the virtual link and configure OSPF routing parameters. Explicitly assign the local router identification number and set the IP address of the virtual link to the IP address at the remote end of the link:

```
enable ospf
set ospf routerid=2.2.2.2
add ospf area=0.0.0.0 authentication=password
add ospf area=0.0.0.1 authentication=password stubarea=off
add ospf range=172.31.0.0 area=0.0.0.0 mask=255.255.0.0
add ospf range=172.29.0.0 area=0.0.0.0 mask=255.255.0.0
add ospf range=172.30.0.0 area=0.0.0.1 mask=255.255.0.0
add ospf interface=ppp0 area=0.0.0.0 password=asecret
add ospf interface=eth0 area=0.0.0.1 password=bsecret
add ospf interface=virt0 area=0.0.0.1 password=esecret
virt=4.4.4.4
```

5. Configure interfaces on switch 3.

Create IP interfaces to use the interfaces, and assign an OSPF metric to each IP interface:

```
create ppp=0 over=syn0
enable ip
add ip int=ppp0 IP=172.30.12.1 mask=255.255.255.0 ospf=1
add ip int=eth0 IP=172.30.100.2 mask=255.255.255.0 ospf=1
```

6. Configure switch 3 as an OSPF router.

Create an OSPF area, assign the IP interfaces to the area, and configure OSPF routing parameters:

```
enable ospf
add ospf area=0.0.0.1 authentication=password stubarea=off
add ospf range=172.30.0.0 area=0.0.0.1 mask=255.255.0.0
add ospf interface=ppp0 area=0.0.0.1 password=csecret
add ospf interface=eth0 area=0.0.0.1 password=bsecret
```

7. Configure interfaces on switch 4.

Create IP interfaces to use the interfaces, and assign an OSPF metric to each IP interface:

```
create ppp=0 over=syn0
enable ip
add ip int=ppp0 IP=172.30.12.2 mask=255.255.255.0 ospf=1
add ip int=eth0 IP=172.29.5.1 mask=255.255.255.0 ospf=1
```

8. Configure switch 4 as an OSPF router.

Create two OSPF areas, assign the IP interfaces to the areas, establish the virtual link, and configure OSPF routing parameters:

```
enable ospf
set ospf routerid=4.4.4.4
add ospf area=0.0.0.0 authentication=password
add ospf area=0.0.0.1 authentication=password stubarea=off
```

```
add ospf range=172.31.0.0 area=0.0.0.0 mask=255.255.0.0
add ospf range=172.30.0.0 area=0.0.0.1 mask=255.255.0.0
add ospf range=172.29.0.0 area=0.0.0.0 mask=255.255.0.0
add ospf interface=virt0 area=0.0.0.1 password=esecret
    virt=2.2.2.2
add ospf interface=ppp0 area=0.0.0.1 password=csecret
add ospf interface=eth area=0.0.0.0 password=dsecret
```

Command Reference

This section describes the commands available to configure and monitor OSPF routing on the switch. OSPF requires the IP module to be enabled and configured correctly. See [Chapter 23, Internet Protocol \(IP\)](#) for detailed descriptions of the commands required to enable and configure IP.

Some interface and port types mentioned in this example may not be supported on your switch. The interface and port types that are available vary depending on your product's model, and whether an expansion unit (PIC, NSM) is installed. For more information, see the Hardware Reference for the switch.

The shortest valid command is denoted by capital letters in the Syntax section. See [“Conventions” on page lxvi of About this Software Reference](#) in the front of this manual for details of the conventions used to describe command syntax. See [Appendix A, Messages](#) for a complete list of messages and their meanings.

add ospf area

Syntax ADD OSPF AREa={Backbone|*area-number*}
[Authentication={NONE|PASSWORD|MD5}]
[NSSAStability=1..3600] [NSSATranslator={CANDidate|
ALways}] [STUBArea={ON|OFF|YES|NO|NSSA|True|False}]
[STUBMetric=0..16777215] [SUMmary={SEND|NONE|OFF|NO|
False}]

where *area-number* is a 4-byte OSPF area number in dotted decimal notation

Description This command adds an OSPF area to the OSPF area table.

The **area** parameter specifies the area number, and is required; other parameters are optional. The area number 0.0.0.0 is reserved for the backbone area, and can be used interchangeably with **backbone**. The specified area must not already exist in the area table.

The **authentication** parameter specifies the type of authentication required for the area. If **none** is specified then no authentication is used on incoming OSPF packets. If **password** is specified, a simple password up to eight characters is used to authenticate each incoming OSPF packet. The password is configured on a per-interface basis. If **md5** is specified, MD5 message digests are used to authenticate OSPF packets. Define the MD5 keys used in the process for every interface using the [add ospf md5key command on page 29-35](#). The default is **none**.

RFC 1583 specifies that authentication be configured per area, but RFC 2328 specifies that authentication be configured per interface. This command represents a hybrid of these solutions, to allow backward compatibility. An authentication type is set up per area to act as a default for all interfaces in the area. The default setting for interfaces is to use the area default, but each interface can be individually set to any authentication method.

The **nssatranslator** parameter sets the NSSA translator role when the switch is acting as an NSSA border router. If you specify **always**, the switch will always translate Type-7 LSAs to Type-5 LSAs, regardless of the translator state of other border routers in the NSSA, as long as it retains border router status. If it loses border router status it will stop translating Type-7 LSAs until it regains border router status. If you specify **candidate**, the switch will participate in the NSSA translator election process. The NSSA border router with the highest router identifier is elected as the translator. Once elected, the switch will translate Type-7 LSAs until it loses border router status or another NSSA border router with a higher router identifier is elected as the translator. The default is **candidate**. If the switch is acting as a translator it will set the Nt bit in router LSAs it originates into the NSSA. The **nssatranslator** parameter is only valid when **stubarea** is set to **nssa**.

The **nssastability** parameter specifies the additional time, in seconds, that the switch will continue to translate Type-7 LSAs after losing the translator role. An elected translator loses its translator role when another NSSA border router with a higher router identifier is elected as translator, or an NSSA router configured to always translate gains border router status. The time interval allows for a more stable transition to the newly elected translator and minimises excessive flushing of translated Type-7 LSAs. The default is 40. The **nssastability** parameter is only valid when **stubarea** is set to **nssa** and **nssatranslator** is set to **candidate**.

The **stubarea** parameter specifies whether the switch treats the area as a stub area. AS external advertisements are not flooded into or out of stub areas. The backbone cannot be configured as a stub area, nor can a virtual link be configured through a stub area. If **area** is set to 0.0.0.0 or **backbone**, then **stubarea** must be **off**. If **area** specifies the transit area of a virtual link, **stubarea** must be **off**. All switches within a particular area must have the same setting for **stubarea**. The **nssa** value specifies that the area is a not-so-stubby-area (NSSA). External routes can be imported as type 7 advertisements in a NSSA. The default is **off** when **area** is set to 0.0.0.0 or **backbone**; or **on** when **area** is set to any other (non-backbone) value.

If the area has been configured as a stub area, and the switch is to act as the area border router for the stub area, then the **stubmetric** parameter specifies the metric (cost) of the default route as advertised by the switch in the default summary link. The default is 1.

The **summary** parameter controls the generation of summary LSA's into stub areas. By default, the default (0.0.0.0) summary LSA is emitted into a stub area by an area border router. If **send**, summary LSA's from other areas are also emitted into the stub area. If **none**, the default (0.0.0.0) summary LSA is emitted into the stub area. If **stubarea** is **nssa**, then the default is **send**; otherwise, the default is **none**.

Examples To set up two areas on a switch, with the first one being the backbone, use the commands:

```
add ospf are=0.0.0.0
add ospf are=0.0.0.1
```

To subsequently change area 1 into a stub area, use the command:

```
set ospf are=0.0.0.1 stuba=on stubm=10
```

Because this switch has multiple areas defined, it acts as an area border router and as such, the **stubmetric** parameter must be defined when setting the area to a stub area. All other switches within area 1 should also have the **stubarea** parameter set to **on**. They do not require that a **stubmetric** be assigned.

Related Commands

- [add ospf range](#)
- [delete ospf area](#)
- [delete ospf range](#)
- [set ospf area](#)
- [set ospf range](#)
- [show ospf area](#)
- [show ospf range](#)

add ospf host

Syntax ADD OSPF HOSt=*ipadd* [METric=0..65535]

where *ipadd* is an IP address in dotted decimal notation

Description This command adds a static OSPF host route to the routing table.

This command is a synonym for the [add ospf stub](#) command. The switch treats OSPF host routes and stub networks identically. If you add a host route with this command, it will appear in the output of the [create config](#) and [show config](#) commands as the equivalent stub network.

OSPF allows host routes to be advertised to the local OSPF area. Such routes are advertised with a route mask of 255.255.255.255. Host routes are normally used for point-to-point networks where it is undesirable to run OSPF with IP hosts directly connected to the switch.

The **host** parameter specifies the IP address of the host or point-to-point network. The IP address must fall within one of the OSPF ranges defined on the switch. The host must not already exist in the routing table.

The **metric** parameter specifies the metric for the route to the host. The default is 1.

Examples Ethernet interface vlan1 has been configured, and an OSPF interface has been defined with an IP address of 172.30.1.1 and network mask 255.255.255.0 to be in area 0.0.0.1. To define a host 172.30.1.2 on a PPP link, use the command:

```
add ospf ho=172.30.1.2 met=1
```

Traffic for address 172.30.1.2 is directed down the PPP link instead of the Ethernet interface because the mask for the host LSA entry is 255.255.255.255 and is more specific than the Ethernet network mask of 255.255.255.0.

Related Commands

- [add ospf interface](#)
- [delete ospf host](#)
- [set ospf host](#)
- [show ospf host](#)
- [show ospf interface](#)

add ospf interface

Syntax ADD OSPF INTerface=*interface* AREa={BACkbone|*area-number*}
 [AUTHentication={AREadefault|NONE|PASSword|MD5}]
 [BOOST1=0..1023] [DEadinterval=2..2147483647]
 [DEMAND={ON|OFF|YES|NO|True|False}]
 [HELlointerval=1..65535] [NETwork={BROadcast|
 NON-broadcast}] [PASSive={ON|OFF|YES|NO|True|False}]
 [PASSword={NONE|*password*}] [POLLIinterval=1..2147483647]
 [PRIOrity=0..255] [RXmtinterval=1..3600]
 [TRansitdelay=1..3600] [VIRTuallink=*router-id*]

where:

- *interface* is a valid interface name.
- *area-number* is a 4-byte OSPF area number in dotted decimal notation.
- *password* is a string 1 to 8 characters long. Valid characters are any printable character. If *password* contains spaces, then it must be in double quotes.

Description This command adds an OSPF interface attached to the specified IP interface, or a unique OSPF virtual interface. Virtual interfaces are used for OSPF virtual links. To add an interface the associated area and its range or ranges must be defined first.

The **interface** parameter specifies the name of the OSPF interface to add, and must be a defined IP interface, or a valid virtual interface instance. Valid interfaces are:

- PPP (such as ppp0)
- VLAN (such as vlan1)
- FR (such as fr0)
- virtual interface (such as virt9)

The interface must already exist. To see a list of all currently available interfaces, use the [show interface command on page 11-87 of Chapter 11, Interfaces](#).

Disabled ghost OSPF interfaces exist for each IP interface that is attached to OSPF, for use with SNMP. These interfaces must be added to OSPF by using the **add ospf interface** command before they can be used by OSPF. Deleting the interface by using the [delete ospf interface command on page 29-43](#) turns it into a disabled ghost OSPF interface again.

Before the interface can be added, both the area and ranges for the area must be defined. A check is made to see if the interface is within one of the defined ranges for the area.

Important Addressless PPP links are the only exception to the range checking mechanism since there is no range to check them against.

The **area** parameter specifies the area to which the interface belongs. When defining a virtual link, the area number is the area number of the transit area used for the virtual link. The area number 0.0.0.0 is reserved for the backbone area, and can be used interchangeably with **backbone**. The specified area must already exist in the area table.

The **authentication** parameter specifies the type of authentication required for the interface. If **areadefault** is specified, the authentication type is taken from the area setting. If **none** is specified, then no authentication is used on incoming OSPF packets. If **password** is specified, a simple password up to eight characters long is used to authenticate each incoming OSPF packet. The password is configured for each interface. If **md5** is specified, MD5 message digests are used to authenticate OSPF packets. Use the [add ospf md5key command on page 29-35](#) to define the MD5 keys for every interface in the process. The default is **areadefault**.

RFC 1583 specifies that authentication be configured per area, but RFC 2328 specifies that authentication be configured per interface. This command represents a hybrid of these solutions, to allow backward compatibility. An authentication type is set up per area to act as a default for all interfaces in the area. The default setting for interfaces is to use the area default, but each interface can be individually set to any authentication method.

The **boost1** parameter increases the type 1 metrics of all OSPF LSAs flooded out the interface by the value specified. Setting the value to anything other than 0 increases the LSA metrics by that amount. For example, setting the parameter in the switch to **boost1=2** boosts the metrics by 2, so an LSA in this switch with a metric of 3 has a value of 5 when it is flooded out the interface. The default is 0.



Caution The **boost1** parameter is a departure from the OSPF specification so should be used with extreme caution.

The **deadinterval** parameter specifies the interval in seconds after which no hello packets are received by a neighbour, the neighbour declares the switch to be down. The timer is advertised in the hello packets. All switches on the same network where the interface is attached must have the same **deadinterval** timer. The value must be at least twice the value of the **hellointerval** timer. The recommended multiplier is four. For example, if the **hellointerval** timer is set to 10 seconds, then the **deadinterval** timer should be set to 40 seconds. The default is four times the value of the **hellointerval** timer.

The **demand** parameter specifies whether the interface connects to a demand circuit. Two switches connecting to the same common network segment need not agree on that segment's demand circuit status. This means that configuring one switch does not require configuring other switches that connect to the same common network segment. If one switch has been configured and the common network is a broadcast or non-broadcast multi-access (NBMA) network, the behaviour (e.g. sending, receiving hello packets) of the network remains the same, just as if the interface has not been configured as a demand circuit. If one switch has been configured and the common network segment is a point-to-point link, the switch on the other end may agree to treat the link as a demand circuit and the point-to-point network receives the full benefit. When broadcast and non-broadcast multi-access (NBMA) networks are declared as demand circuits (i.e. more than one switch has the network configured as a demand circuit), routing update traffic is reduced but the periodic sending of hellos is not, which requires that the data link connection remain constantly open. The default is **off**.

Virtual links created using the **virtuallink** parameter are always treated as demand circuits, and ignore the setting of the **demand** parameter.

The **hellointerval** parameter specifies the interval in seconds between hello packets transmitted over the interface. All switches on the network where the

interface is attached must have the same **hellointerval** timer. The value must be less than the value of the **pollinterval** timer. The default is 10.

The **network** parameter specifies the OSPF network type of the interface, and is only valid for VLAN interfaces. Specify **broadcast** if you want OSPF to treat the network as a broadcast network. Hello messages are transmitted as broadcast messages, and neighbours are learned dynamically. You can not configure static neighbours or use the **pollinterval** parameter to set the time interval between hello messages to inactive neighbours. Specify **non-broadcast** if you want OSPF to treat the network as an NBMA network. All OSPF packets are transmitted as unicast messages, so neighbours must be statically defined. You can use the **pollinterval** parameter to set the time interval between hello messages to inactive neighbours. The default is **broadcast**.

The **passive** parameter specifies whether the interface acts as a passive interface. Specify **on**, **yes** or **true** if you want the interface to act as a passive OSPF interface. The interface does not take part in the OSPF protocol and OSPF packets are not sent or received on the interface. However, the network that the interface belongs to is added to the router LSA as a stub network, so that routing information about the network will be carried through OSPF to the rest of the routing domain. Specify **off**, **no** or **false** if you want the interface to act as a normal OSPF interface. The default is the value of the **passiveinterfacedefault** parameter of the [set ospf command on page 29-59](#), or **off** if the **passiveinterfacedefault** parameter has not been set.

The **password** parameter specifies the password used for authentication. A password is required if the authentication scheme for the area has been set to **password** by using the [add ospf area command on page 29-28](#) or the [set ospf area command on page 29-63](#). If **none** is specified, no password is configured on the interface. The default is **none**.

The **pollinterval** timer defines the time in seconds that hello packets are sent to neighbouring switches that are deemed to be inactive. The **pollinterval** timer is used on NBMA networks (for example, Frame Relay X.25) and point-to-point networks configured as OSPF on demand circuits. The value must be greater than the value of the **hellointerval** timer. When the neighbour is in the "Down" state, hellos are sent to the neighbour at the interval set for **pollinterval**. The default is four times the value of the **hellointerval** timer.

The **priority** parameter is used on multi-access networks to set the router priority. When two switches attached to a network attempt to become the designated router, the one with the highest priority takes precedence. If the priorities are the same then the switch with the highest router identification number takes precedence. A switch with a priority of zero is ineligible to become the designated router. Router priority can be configured only for switches attached to multi-access networks. The default is 1.

The **rxmtinterval** parameter specifies the time interval, in seconds, between link state retransmissions, for adjacencies on the interface. The timer is also used when retransmitting database description and link state request packets. The value should be set well above the round trip time between the two switches. It should be set higher for slow serial and virtual links. A typical value is 5 seconds on a local area network. The default is 5.

The **transitdelay** parameter specifies an estimate in seconds required to transmit a link state update packet over this interface. This time is added to link state advertisements sent over the interface. The value of this parameter should take into account the transmission and propagation delays of the interface. This mechanism helps keep the link state advertisement timers

synchronised on different switches. A typical value is 1 second on a local area network. The default is 1.

The **virtuallink** parameter specifies the router identification number for another area border router to be included in the backbone using a virtual link. Each end of the virtual link must be configured. This parameter is required when a virtual interface is added to OSPF. Virtual links are always treated as demand circuits. This is because when a virtual link's underlying physical path contains one or more demand circuits, periodic OSPF protocol exchanges over the virtual link would unnecessarily keep the underlying demand circuits open. When configuring virtual links, the **area** parameter specifies the area number of the transit area used for the virtual link.

Examples To associate a PPP interface that connects to the demand circuit with the area 0.0.0.1, use the command:

```
add ospf int=ppp0 are=0.0.0.1 dem=on
```

To associate a VLAN interface with the backbone area, use the command:

```
add ospf int=vlan1 are=ba
```

Related Commands

- [add ospf area](#)
- [add ospf range](#)
- [delete ospf interface](#)
- [disable ospf interface](#)
- [enable ospf interface](#)
- [reset ospf interface](#)
- [set ospf area](#)
- [set ospf interface](#)
- [set ospf range](#)
- [show ospf area](#)
- [show ospf interface](#)
- [show ospf range](#)

add ospf md5key

Syntax `ADD OSPF MD5key=key ID=1..255 INTerface=interface`

where:

- *key* is a string 1 to 16 characters long and is case sensitive. Valid characters are uppercase and lowercase letters and digits.
- *interface* is a valid interface name.

Description This command adds an MD5 key for use when authenticating OSPF packets on a particular interface. For MD5 authentication to succeed, all key and ID values must be identical on all interfaces in the same physical network.

The **md5key** parameter specifies key used for MD5 authentication. The more characters in the key, the greater the security it offers. We recommend that the MD5 key be at least 9 characters long and be changed every month.

The **id** parameter specifies the identification number for this key. The ID is used in the authentication of packets to identify to the remote device which key is being used in this packet.

The **interface** parameter specifies the OSPF interface with which this key is associated. Each interface has its own set of keys, and keys must be identified by interface as well as key ID. The interface can be any OSPF interface for which MD5 authentication is required.

Example To add an MD5 authentication key called “mj48dhw05” with an ID of 3 to interface vlan1, use the command:

```
add ospf md5key=mj48dhw05 id=3 interface=vlan1
```

Related Commands

- [add ospf interface](#)
- [delete ospf md5key](#)
- [show ospf md5key](#)

add ospf neighbour

Syntax ADD OSPF NEighbour=*ipadd* PRIOrity=0..255

where *ipadd* is an IP address in dotted decimal notation

Description This command adds a non-broadcast multi-access neighbour and sets the neighbour parameters. On non-broadcast multi-access networks, such as Frame Relay, the neighbours cannot be discovered by dynamic means. To overcome this problem, neighbours are configured statically. All parameters are required to add a neighbour.

Static OSPF neighbours should be defined on switches that are attached to an OSPF network and are eligible to become the *Designated Router* (DR) or *Backup Designated Router* (BDR) for that network. To be eligible to become the DR or BDR for an OSPF network, a switch must have at least one interface on the OSPF network with a non-zero priority. The priority of an OSPF interface is set with the **priority** parameter in the **add ospf interface** or the **set ospf interface** command.

Any OSPF router connected to an OSPF network via an OSPF interface with a non-zero priority should be configured with static OSPF neighbours corresponding to all other OSPF routers attached to the OSPF network. OSPF routers that do not have an OSPF interface with a non-zero priority should not be configured with static OSPF neighbours.

The **neighbour** parameter specifies the IP address of the neighbour. The IP address must fall within an IP address range associated with an OSPF interface. A neighbour with the specified IP address must not already exist in the neighbour table.

The **priority** parameter specifies the priority of the neighbour router. If the priority for the neighbour is set to zero then it is not initially considered eligible to become the designated router. The priority for the switch itself is set in the interface priority.

Only a few OSPF routers in an NBMA network should be eligible to become the designated router.

Examples To define a neighbour with IP address 172.30.1.2, use the command:

```
add ospf nei=172.30.1.2 prio=1
```

Related Commands

- [add ospf interface](#)
- [delete ospf interface](#)
- [delete ospf neighbour](#)
- [set ospf interface](#)
- [set ospf neighbour](#)
- [show ospf interface](#)
- [show ospf neighbour](#)

add ospf range

Syntax ADD OSPF RANGE=*ipadd* AREA={BACKbone|*area-number*}
[MASK=*ipadd*] [EFFECT={ADVERTISE|DONOTADVERTISE}]

where:

- *ipadd* specifies an IP address in dotted decimal notation.
- *area-number* is a 4-byte OSPF area number in dotted decimal notation.

Description This command adds an OSPF range to an OSPF area. An OSPF area is defined as a list of IP address ranges. A IP address range is defined by an IP address and network mask pair. After an area has been created, the ranges that are to be active in the area must be defined. All the OSPF routers within an area must use the same set of ranges for the area.

Configuring a range on an OSPF ABR enables the ABR to aggregate the network advertisements from one area into another area in the form of summary LSAs. However, networks advertised into a transit area should not be aggregated into summary LSAs, so the switch ignores any configured range when it is advertising into a transit area.

The **area** parameter specifies the area number of the area to which the range is added. The area number 0.0.0.0 is reserved for the backbone area, and can be used interchangeably with **backbone**. The specified area must already exist in the area table.

The **range** parameter specifies the base IP address of the range.

The **mask** parameter specifies a mask used in combination with the base address to determine addresses that belong to the range. The value must be consistent with the value specified for the **range** parameter. The default is the network mask for the address class of the IP address (e.g. 255.255.0.0 for a Class B address and 255.255.255.0 for a Class C address).

The range defined by the combination of **range** and **mask** must not overlap any other range defined on the switch.

The **effect** parameter controls the exporting of ranges into other areas by an area border router. This parameter is valid when the switch is configured as an area border router, and can be used to control where the traffic into and out of an area for a particular network travels. If **advertise**, summary LSAs are created for the range; if **donotadvertise**, they are not created. The default is **advertise**.

Examples To add an IP address range of 172.30.0.0–172.30.255.255 to area 1, use the command:

```
add ospf ran=172.30.0.0 mask=255.255.0.0 are=0.0.0.1
```

Related Commands

- [delete ospf range](#)
- [set ospf range](#)
- [show ospf range](#)

add ospf redistribute

Syntax ADD OSPF REDistribute PROTOcol={BGP|INTERface|RIP|STATIC}
[LIMIT=1..4000] [METric={0..16777214|ORIGINAL}]
[ROUTEMap=*routermap*] [SUBNET={ON|OFF|YES|NO|True|False}]
[TAG={1..65535|ORIGINAL}] [TYPE={1|2|ORIGINAL}]

where *routermap* is the name of an IP route map

Description This command creates a redistribution definition for OSPF to use when it imports BGP, RIP, non-OSPF interface, and statically configured routes into AS external or NSSA AS external LSAs that it advertises. Only one redistribution definition can be created for each route type.

Adding a BGP, RIP, or static route redistribution definition will change the setting of the **bgpimport**, **rip** and **staticexport** parameters of the [set ospf command on page 29-59](#). For more information, see “[Redistributing External Routes](#)” on page 29-15. If you configure a BGP route filter using the **bgpfilter** parameter of the [set ospf](#) command, the filter will be applied before any BGP route redistribution definition.

The **protocol** parameter specifies the type of route to redistribute. Specify **bgp** or **rip** to redistribute routes derived from BGP or RIP. Specify **interface** to redistribute non-OSPF interface routes. Specify **static** to redistribute statically configured routes.

The **limit** parameter specifies the maximum number of routes that can be redistributed into OSPF for the specified protocol. The default is 1000. If you add a BGP redistribution definition, the **limit** parameter overwrites the setting of the **bgplimit** parameter in the [set ospf command on page 29-59](#). Note that the limit sets the maximum number of routes that OSPF can include after route map filtering. For example, if 2000 routes are included after route map filtering but the limit is 1000 routes, only the first 1000 of those routes will be imported. This means you should do one of the following:

- make sure that the limit is set higher than the maximum possible number of routes that your route map would include, or
- assign low entry numbers to the route map entries that match the most preferred routes. That way, if the number of routes reaches your limit, OSPF will have imported the most important routes.

The **metric** parameter specifies the route metric that OSPF assigns to routes that it redistributes. If you specify **original**, the original route metric is preserved in the redistributed route—metric1 for Type-1 routes or metric2 for Type-2 routes. If you assign a route map that sets the metric, the route map overrides the setting in this parameter. The default is 20.

The **routermap** parameter specifies a route map to filter and/or modify routes before redistributing them through OSPF. This route map can match on interface, prefix list, next hop, metric, or tag. The route map must already exist. To create a route map, use the [add ip routermap command on page 31-49 of Chapter 31, Filtering IP Routes](#). If you do not specify a route map, OSPF will import all routes for the protocol, depending on the setting of the **subnet** parameter.

The **subnet** parameter specifies whether OSPF redistributes subnets. Specify **off** if you want OSPF to only redistribute classfull network routes. Specify **on** if

you want OSPF to redistribute classless network routes as well. The default is **on**.

The **tag** parameter specifies a number OSPF uses to label routes that it redistributes. If you specify **original**, the original route tag is preserved in the redistributed route. If you assign a route map that sets the tag, the route map overrides the setting in this parameter. The default is **original**.

The **type** parameter specifies the OSPF external route type that OSPF assigns to routes that it redistributes. Use the **type** parameter to ensure that all externally-sourced OSPF routes are the same type and therefore use the same method to calculate route metrics. Specify **1** if you require the routes to have a Type-1 external metric, or **2** if you require the routes to have a Type-2 external metric. If you assign a route map that sets the type, the route map overrides the setting in this parameter. See [“Path types” on page 29-9](#) for more information about the types. The default is **2**.

Example To import static routes into OSPF and use the default settings, use the command:

```
add ospf red prot=sta
```

To import interface routes into OSPF and set their metric to 30, use the command:

```
add ospf red prot=int met=30
```

Related Commands [add ip routemap in Chapter 31, Filtering IP Routes](#)
[delete ospf redistribute](#)
[set ospf redistribute](#)
[show ospf redistribute](#)

add ospf stub

Syntax `ADD OSPF STUB=ipadd MASK=ipadd [METric=0..65535]`

where *ipadd* is an IP address in dotted decimal notation

Description This command adds a static non-OSPF network as a stub network to the OSPF routing table. This enables networks, such as an LAN with only one OSPF router, to be added as stub networks to OSPF without running OSPF on the interface to which the network is attached. The stub links are added to the switch's LSA.

The **stub** parameter specifies the base IP address of the stub network. The IP address must fall within one of the OSPF ranges defined on the switch. The stub must not already exist in the routing table.

The **mask** parameter specifies a mask used in combination with the base address to determine addresses that belong to the stub network. The value must be consistent with the value specified for the **stub** parameter.

The **metric** parameter specifies the metric for the route to the stub network. The default is 1.

Examples Ethernet interface, vlan1 is attached to an Ethernet LAN with an IP subnet address of 172.30.1.0 and network mask 255.255.255.0. To add the LAN attached to the Ethernet interface as a non-OSPF stub network, use the command:

```
add ospf stub=172.30.1.0 mask=255.255.255.0
```

Related Commands

- [add ospf host](#)
- [add ospf interface](#)
- [delete ospf stub](#)
- [set ospf stub](#)
- [show ospf stub](#)

add ospf summaryaddress

Syntax ADD OSPF SUMM^{ary}address=*ipadd* MASK=*ipadd* [ADV^{er}tise={ON|OFF|YES|NO|True|False}] [TAG=0..65535]

where *ipadd* is an IP address in dotted decimal notation

Description This command adds a summary address to OSPF, configuring a single route to advertise (or not) in an OSPF AS external LSA. When OSPF redistributes routes, networks that fall within the range of this summary address are advertised by a single AS external LSA containing this summary address.

The AS external LSAs for static routes use the metric and metric type set by the parameters in the OSPF redistribution definition for the protocol. See [“Redistributing External Routes” on page 29-15](#).

The AS external LSAs for routes imported from BGP and RIP use the metric and metric type of the original route.

The **summaryaddress** parameter specifies the base IP address of the summary address.

The **mask** parameter specifies a mask used in combination with the base IP address to define a summary address range.

The **advertise** parameter specifies whether OSPF advertises the summary address in an AS external LSA for this network range. Specify **on**, **yes** or **true** if you want OSPF to advertise the summary address in AS external LSAs. Individual networks that fall within the range of the summary address are not advertised. Specify **off**, **no** or **false** if you do not want OSPF to advertise either the summary address or the individual networks within the range of the summary address. The default is **on**.

The **tag** parameter specifies the tag value that OSPF places in the AS external LSAs created as a result of redistributing the summary route. The tag overrides tags set by the original route and by the route map used to select the original routes for redistribution. If you do not specify a value, OSPF places a tag value of 0 in the AS external LSAs.

Examples To create a summary address for the class C network 198.168.1.*, and advertise the address with a tag of 13, use the command:

```
add ospf summ=192.168.1.0 mask=255.255.255.0. tag=13
```

Related Commands

- [delete ospf summaryaddress](#)
- [set ospf summaryaddress](#)
- [show ospf summaryaddress](#)

delete ospf area

Syntax DELEte OSPF AREa={BAckbone | *area-number*}

where *area-number* is a 4-byte OSPF area number in dotted decimal notation

Description This command deletes an OSPF area from the OSPF area table.

The **area** parameter specifies the area number of the area to delete. The area must already exist in the area table. The area number 0.0.0.0 is reserved for the backbone area, and can be used interchangeably with **backbone**.

Examples To delete an OSPF area with the area number 0.0.0.1, use the command:

```
del ospf are=0.0.0.1
```

Related Commands

- [add ospf area](#)
- [add ospf range](#)
- [delete ospf range](#)
- [set ospf area](#)
- [set ospf range](#)
- [show ospf area](#)
- [show ospf range](#)

delete ospf host

Syntax DELEte OSPF HOSt=*ipadd*

where *ipadd* is an IP address in dotted decimal notation

Description This command deletes a static OSPF host route or all static OSPF host routes from the routing table.

OSPF allows host routes to be advertised to the local OSPF area. Such routes are advertised with a route mask of 255.255.255.255. Host routes are normally used for point-to-point networks where it is undesirable to run OSPF with IP hosts directly connected to the switch.

The **host** parameter specifies the IP address of the host or point-to-point network to be deleted. The host must exist in the routing table.

Examples To delete the OSPF host with IP address 172.30.1.2, use the command:

```
del ospf ho=172.30.1.2
```

Related Commands

- [add ospf host](#)
- [set ospf host](#)
- [show ospf host](#)

delete ospf interface

Syntax `DELEte OSPF INTerface=interface`

where *interface* is a valid interface name

Description This command deletes an OSPF interface.

The **interface** parameter specifies the name of the OSPF interface to delete, and must be a defined IP interface, or a valid virtual interface instance. Valid interfaces are:

- PPP (such as ppp0)
- VLAN (such as vlan1)
- FR (such as fr0)
- virtual interface (such as virt9)

To see a list of current valid interfaces, use the **show ospf interface** command or the [show interface command on page 11-87 of Chapter 11, Interfaces](#). An interface cannot be deleted if static neighbours exist that depend on the interface. Delete the static neighbours first, then delete the interface.

Disabled ghost OSPF interfaces exist for each IP interface that is attached to OSPF, for use with SNMP. These interfaces must be added to OSPF with the **add ospf interface** command before they can be used by OSPF. Deleting the interface with the **delete ospf interface** command turns it back into a disabled ghost OSPF interface.

Examples To delete OSPF interface ppp0, use the command:

```
del ospf int=ppp0
```

Related Commands

- [add ospf interface](#)
- [delete ospf interface](#)
- [disable ospf interface](#)
- [enable ospf interface](#)
- [reset ospf interface](#)
- [set ospf interface](#)
- [show ospf interface](#)

delete ospf md5key

Syntax DELEte OSPF MD5key ID=1..255 INTErface=*interface* [FORce]

where *interface* is a valid interface name

Description This command deletes an MD5 key used for authenticating OSPF packets on a particular interface. If the key is being used when the command is executed, the command will fail unless the **force** parameter is also specified. Before deleting a key, configure a succeeding key on all interfaces in the physical network. You can delete the previously active key without using the **force** parameter. The new key can take over authentication duties when the active key is deleted. The output of the **show ospf md5key** command shows whether a key is active or inactive.

The **id** parameter specifies the identification number of the key that will be deleted.

The **interface** parameter specifies the OSPF interface with which this key is associated. Each interface has its own set of keys, and keys must be identified by interface as well as key ID.

The **force** parameter specifies that the MD5 key should be deleted even when it is being used. If an illicit router is using the key, using the **force** parameter ensures that the key is deleted.

Important Forcibly deleting an active MD5 key may lead to partial network failure if succeeding keys are not configured on all interfaces in the physical network.

Example To delete a key with the ID 35 that was used last week and should be replaced, use the command:

```
delete ospf md5key id=35 interface=vlan1
```

Related Commands [add ospf md5key](#)
[show ospf md5key](#)

delete ospf neighbour

Syntax `DELEte OSPF NEIghbour=ipadd`

where *ipadd* is an IP address in dotted decimal notation

Description This command deletes a non-broadcast multi-access neighbour.

The **neighbour** parameter specifies the IP address of the neighbour. The neighbour with the specified IP address must exist in the neighbour table.

Examples To delete the neighbour with IP address 172.30.1.2, use the command:

```
del ospf nei=172.30.1.2
```

Related Commands

- [add ospf neighbour](#)
- [set ospf neighbour](#)
- [show ospf neighbour](#)

delete ospf range

Syntax `DELEte OSPF RANge=ipadd`

where *ipadd* specifies an IP address in dotted decimal notation

Description This command deletes an OSPF range. An OSPF area is defined as a list of IP address ranges. A IP address range is defined by an IP address and network mask pair. All OSPF routers in an area must use the same set of ranges for the area.

The **range** parameter specifies the base IP address of the range to be deleted. The specified range must already exist.

Examples To delete the range with a base IP address of 172.30.0.0, use the command:

```
del ospf ran=172.30.0.0
```

Related Commands

- [add ospf area](#)
- [add ospf range](#)
- [set ospf range](#)
- [show ospf range](#)

delete ospf redistribute

Syntax DELEte OSPF REDistribute PROTOcol={BGP|INTERface|RIP|STATIC}

Description This command deletes a route redistribution definition that OSPF uses when it imports routes from other protocols into AS external or NSSA AS external LSAs that it advertises. Only one redistribution definition can exist for each route type.

Deleting a BGP, RIP, or static interface route redistribution definition will change the setting of the **bgpimport**, **rip** and **staticexport** parameters of the [set ospf command on page 29-59](#). For more information, see “[Redistributing External Routes](#)” on page 29-15.

The **protocol** parameter specifies the route redistribution definition to delete. OSPF no longer imports and redistributes routes from the protocol. Specify **bgp** or **rip** to delete the redistribution definition for BGP or RIP routes. Specify **interface** to delete the redistribution definition for non-OSPF interface routes. Specify **static** to delete the redistribution definition for statically configured routes.

Example To stop OSPF importing and redistributing static routes, use the command:

```
del ospf red prot=sta
```

Related Commands [add ospf redistribute](#)
[set ospf redistribute](#)
[show ospf redistribute](#)

delete ospf stub

Syntax `DELEte OSPF STUB=ipadd MASK=ipadd`

where *ipadd* is an IP address in dotted decimal notation

Description This command deletes a static non-OSPF network as a stub network from the OSPF routing table. Networks, such as an LAN with only one OSPF router, can be added as stub networks to OSPF without running OSPF on the interface to which the network is attached. The stub links are added to the switch's LSA.

The **stub** parameter specifies the base IP address of the stub network. The IP address must fall within one of the OSPF ranges defined on the switch. The stub must already exist in the routing table.

The **mask** parameter specifies a mask used in combination with the base address to determine addresses that belong to the stub network. The value must be consistent with the value specified for the **stub** parameter.

Examples To delete the stub network attached to Ethernet interface vlan1 with an IP subnet address of 172.30.1.0 and network mask 255.255.255.0, use the command:

```
del ospf stub=172.30.1.0 mask=255.255.255.0
```

Related Commands

- [add ospf stub](#)
- [delete ospf host](#)
- [delete ospf interface](#)
- [set ospf stub](#)
- [show ospf stub](#)

delete ospf summaryaddress

Syntax DELEte OSPF SUMMaryaddress=*ipadd*

where *ipadd* is an IP address in dotted decimal notation

Description This command deletes a summary address from OSPF. Networks falling within the summary address range are no longer replaced by the summary address in AS external LSAs, and are redistributed individually.

The **summaryaddress** parameter specifies the base IP address of the summary address to delete. To see a list of existing summary addresses, see the [show ospf summaryaddress command on page 29-105](#).

Examples To delete the summary address 192.168.1.0, use the command:

```
del ospf summ=192.168.1.0
```

Related Commands [add ospf summaryaddress](#)
[set ospf summaryaddress](#)
[show ospf summaryaddress](#)

disable ospf

Syntax DISable OSPF

Description This command disables OSPF routing. OSPF must currently be enabled. Networks and hosts that were reachable via OSPF are no longer accessible.

OSPF is disabled by default.

Examples To disable OSPF, use the command:

```
dis ospf
```

Related Commands [enable ospf](#)
[show ospf](#)

disable ospf debug

Syntax `DISable OSPF DEBug={ALL|AUTOcost|IFStAte|LSU|NBRStAte|NRL|NSSA|PACKet|REDistribute|SPF|StAte}`

Description This command disables the generation of OSPF debugging messages.

The **debug** parameter specifies the debugging options to disable. If **all** is specified, all debug options are disabled. If **autocost** is specified, debugging of the calculation of interface metrics using the autocost method is disabled. If **ifstate** is specified, interface state debugging is disabled. If **lsu** is specified, debugging of the reception and processing of Link State Update packets is disabled. If **nbrstate** is specified, neighbour state debugging is disabled. If **nrl** is specified, debugging of changes to the neighbour retransmission list is disabled. If **nssa** is specified, debugging of the operation of NSSAs (Not So Stubby Areas) is disabled. If **packet** is specified, OSPF packet debugging is disabled. If **redistribute** is specified, route redistribution debugging is disabled. If **spf** is specified, debugging for the Shortest Path First routing calculations is disabled. If **state** is specified, both interface and neighbour state debugging are disabled.

Examples To disable OSPF packet debugging, use the command:

```
dis ospf deb=pac
```

Related Commands

- [disable debug active](#) in Chapter 4, Configuring and Monitoring the System
- [disable ospf log](#)
- [enable ospf debug](#)
- [enable ospf log](#)
- [show debug active](#) in Chapter 4, Configuring and Monitoring the System
- [show ospf](#)

disable ospf interface

Syntax DISable OSPF INTerface=*interface*

where *interface* is a valid interface name

Description This command disables an OSPF interface. The interface along with areas assigned to it no longer participate in the OSPF routing protocol.

The **interface** parameter specifies the name of the OSPF interface to disable, and must be a defined IP interface, or a valid virtual interface instance. Valid interfaces are:

- PPP (such as ppp0)
- VLAN (such as vlan1)
- FR (such as fr0)
- virtual interface (such as virt9)

To see a list of current valid interfaces, use the [show ospf interface command on page 29-85](#) or the [show interface command on page 11-87 of Chapter 11, Interfaces](#).

Examples To disable OSPF interface ppp0, use the command:

```
dis ospf int=ppp0
```

Related Commands

- [add ospf interface](#)
- [delete ospf interface](#)
- [enable ospf interface](#)
- [reset ospf interface](#)
- [set ospf interface](#)
- [show ospf interface](#)

disable ospf log

Syntax DISable OSPF LOG

Description This command disables the logging of significant OSPF events to the switch's logging facility. OSPF logging must currently be enabled. See the **enable ospf log** command for a list of events that can be logged. RFC 1247 has a detailed description of each message.

Examples To disable OSPF logging, use the command:

```
dis ospf log
```

Related Commands

- [enable ospf log](#)
- [show ospf](#)

enable ospf

Syntax ENABle OSPF

Description This command enables OSPF routing. OSPF must currently be disabled. OSPF interfaces form adjacencies, link state advertisements (LSAs) are exchanged, and the routing table is built by OSPF by using the Dijkstra algorithm.

OSPF is disabled by default.

Examples To enable OSPF, use the command:

```
ena ospf
```

Related Commands [disable ospf](#)
[show ospf](#)

enable ospf debug

Syntax ENABle OSPF DEBUg={ALL|AUTOCost|IFSTate|LSU|NBRSTate|NRL|NSSA|PACKet|REDistribute|SPF|STate} [DETAil={BRIef|HEADer|LSAFull|LSASummary}] [TIMEOut={NONE|1..2400}]

Description This command enables the generation of OSPF debugging messages.

The **debug** parameter specifies the debugging options to enable. If **all** is specified, all debug options are enabled. If **autocost** is specified, debugging of the calculation of interface metrics using the autocost method is enabled. If **ifstate** is specified, interface state debugging is enabled. If **lsu** is specified, debugging of the reception and processing of Link State Update packets is enabled. If **nbrstate** is specified, neighbour state debugging is enabled. Output from **ifstate** and **nbrstate** includes the interface or neighbour the state change relates to, the event that caused the state change, and the previous and current states of the interface or neighbour. If **nrl** is specified, changes to the neighbour retransmission list are displayed. Note that this option may generate large amounts of debugging output on a large OSPF network. Use it with care. If **nssa** is specified, debugging of the operation of NSSAs (Not So Stubby Areas) is enabled. If **packet** is specified, OSPF packet debugging is enabled. The level of detail shown in packet debugging is set with the **detail** parameter, but the output always contains the direction of the packet, the type of packet, the version of OSPF, the packet's source and destination, the router ID, area, length, checksum and authentication type. If **redistribute** is specified, route redistribution debugging is enabled. If **spf** is specified, debugging for the Shortest Path First routing calculations is enabled. If **state** is specified, both interface and neighbour state debugging are enabled.

The **detail** parameter specifies the level of packet debugging. If **brief** is specified, a short description about the contents of each packet is displayed as well as the OSPF header. If **header** is specified, the OSPF header for each packet is displayed. If **lsafull** is specified, the details of the LSA are displayed as well as the OSPF header. If **lsasummary** is specified, Link State Advertisements (LSA) header information is displayed as well as the OSPF header. The default is **header**.

The **timeout** parameter specifies the length of time that OSPF debugging output is generated. This is limited to a maximum of 40 minutes. For any longer than that, debugging must be enabled without a timeout, then can be disabled using time triggers. The **timeout** parameter overrides any previously entered timeout. If **none** is specified, debugging continues indefinitely. The default is the previously entered **timeout** value, or **none** if no previous timeout value exists.

Examples To enable all OSPF packet debugging and show details of the LSA and the OSPF header, use the command:

```
ena ospf deb=pac det=lsaf
```

Related Commands

- [disable ospf debug](#)
- [disable ospf log](#)
- [enable ospf log](#)
- [show ospf](#)
- [disable debug active](#) in Chapter 4, Configuring and Monitoring the System
- [show debug active](#) in Chapter 4, Configuring and Monitoring the System

enable ospf interface

Syntax `ENABle OSPF INTerface=interface`

where *interface* is a valid interface name

Description This command enables an OSPF interface. The interface along with areas assigned to it now participate in the OSPF routing protocol.

The **interface** parameter specifies the name of the OSPF interface to enable, and must be a defined IP interface, or a valid virtual interface instance. Valid interfaces are:

- PPP (such as ppp0)
- VLAN (such as vlan1)
- FR (such as fr0)
- virtual interface (such as virt9)

To see a list of current valid interfaces, use the [show ospf interface](#) command on page 29-85 or the [show interface](#) command on page 11-87 of Chapter 11, [Interfaces](#).

Examples To enable OSPF interface ppp0, use the command:

```
ena ospf int=ppp0
```

Related Commands

- [add ospf interface](#)
- [delete ospf interface](#)
- [disable ospf interface](#)
- [reset ospf interface](#)
- [set ospf interface](#)
- [show ospf interface](#)

enable ospf log

Syntax ENAbLe OSPF LOG

Description This command enables the logging of significant OSPF events to the switch's logging facility. The following table shows events that can be logged and associated codes. See RFC 1247 for a detailed description of the meaning of each message.

Code	Event
T1	Interface state change
T2	Neighbour state change
T3	Designated router change
T4	New LSA generated by switch
T5	New LSA received
T6	Routing table change
C1	OSPF packet rejected due to header errors
C2	Hello packet rejected
C3	Other packet rejected due to Neighbour state being incorrect
C4	Database Description packet retransmitted
E1	Received LSA checksum incorrect
E2	Database LSA checksum incorrect
R1	Different instances of same LSA
R2	Different ages for 2 copies of LSA
R3	Received more recent self-originated LSA
R4	Received LSA Ack for unknown LSA
R5	Received older LSA
N1	Link state refresh timer gone off
N2	LSA aged to MaxAge
N3	MaxAge aged LSA has been flushed

Examples To enable OSPF logging, use the command:

```
ena ospf log
```

Related Commands [disable ospf log](#)
[show ospf](#)

purge ospf

Syntax PURge OSPF

Description This command purges all OSPF configuration information and resets global OSPF parameters to their defaults. This includes disabling OSPF routing.

Examples To purge OSPF, use the command:

```
pur ospf
```

Related Commands [disable ospf](#)
[enable ospf](#)
[reset ospf](#)
[show ospf](#)

reset ospf

Syntax RESET OSPF

Description This command resets the OSPF routing software and re-initialises dynamic data structures. It does not make OSPF operational when it is incompletely or incorrectly configured. Nor does it reset OSPF counters; use the **reset ospf counter** command for this. A message is sent to the switch's logging facility when one has been defined.

This command disrupts OSPF routing on the switch. Networks and hosts that were reachable through the switch are inaccessible while OSPF restarts.

Examples To reset OSPF, use the command:

```
reset ospf
```

Related Commands [disable ospf](#)
[enable ospf](#)
[purge ospf](#)
[reset ospf counter](#)
[reset ospf interface](#)
[reset ospf spf](#)
[show ospf](#)

reset ospf counter

Syntax RESET OSPF COUnTer

Description This command resets all OSPF counters to their initial values.

Examples To reset all OSPF counters, use the command:

```
reset ospf cou
```

Related Commands

- [purge ospf](#)
- [reset ospf](#)
- [reset ospf interface](#)
- [reset ospf spf](#)
- [show ospf](#)

reset ospf interface

Syntax RESET OSPF INTerface=*interface*

where *interface* is a valid interface name

Description This command resets an enabled OSPF interface. The OSPF interface is shut down, destroying all current routing information for networks attached to the interface. The interface is then restarted and the routing information is re-learned from the interface's network, effectively refreshing the routing database. The OSPF interface state machine is also restarted.

The **interface** parameter specifies the name of the OSPF interface to reset, and must be a defined IP interface, or a valid virtual interface instance. Valid interfaces are:

- PPP (such as ppp0)
- VLAN (such as vlan1)
- FR (such as fr0)
- virtual interface (such as virt9)

To see a list of current valid interfaces, use the **show ospf interface** command or the [show interface command on page 11-87 of Chapter 11, Interfaces](#).

Examples To reset OSPF interface ppp0 that is currently enabled, use the command:

```
reset ospf int=ppp0
```

Related Commands

- [add ospf interface](#)
- [delete ospf interface](#)
- [disable ospf interface](#)
- [enable ospf interface](#)
- [reset ospf](#)
- [reset ospf counter](#)
- [reset ospf spf](#)
- [set ospf interface](#)
- [show ospf interface](#)

reset ospf spf

Syntax RESET OSPF SPF [DEBUg]

Description This command forces a recalculation of the OSPF route table, re-running the Shortest Path First calculation. This command can be used, in conjunction with the **show ospf lsa** and **show ospf debug** commands, to diagnose network problems.

The **debug** parameter specifies that debugging information for the route table recalculation will be output to the port from which this command was entered, rather than the port specified by a previous **enable ospf debug=spf** command, for this one recalculation.

Examples To recalculate the OSPF route table, use the command:

```
reset ospf spf
```

To recalculate the OSPF route table and send debugging information to the local port, use the command:

```
reset ospf spf deb
```

Related Commands

- [purge ospf](#)
- [reset ospf](#)
- [reset ospf counter](#)
- [reset ospf interface](#)
- [show ospf](#)
- [show ospf debug](#)
- [show ospf lsa](#)

set ospf

Syntax SET OSPF [ASExternal={ON|OFF|NSSA}] [BGPFILTER={0..999|NONE}] [BGPIMPORT={ON|OFF|True|False|YES|NO}] [BGPLIMIT=1..4000] [AUTOCOST={ON|OFF}] [DEFROUTE={ON|OFF|True|False|YES|NO}] [DYNINTERFACE={STUB|ASExternal|NONE|NO|OFF|False}] [INRoutemap={*routemap*|NONE}] [METRIC=0..16777215] [PASSiveinterfacedefault={ON|OFF|True|False|YES|NO}] [REFBANDWIDTH=10..10000] [RIP={OFF|EXPORT|IMPORT|BOTH}] [ROUTERID=*ipadd*] [PTPStub={ON|OFF|YES|NO|True|False}] [STATICEXPORT={YES|NO}] [TYPE={1|2}]

where:

- *ipadd* is an IP address in dotted decimal notation
- *routemap* is the name of an IP route map

Description This command sets general OSPF routing configuration parameters.

You can use the **bgpfilter**, **bgpimport**, **bgplimit**, **rip** and **staticexport** parameters to configure OSPF to import BGP, RIP and static routes. However, we recommend that you use route redistribution definitions to import and redistribute routes into OSPF, as they provide more control over how the routes are imported. For compatibility, the **bgpimport**, **bgplimit**, **rip** and **staticexport** parameters are synchronised with the equivalent redistribution definition. Changing the setting of these parameters will add or delete the corresponding route redistribution definition. Similarly, adding or deleting a route redistribution definition changes the setting of the corresponding **bgpimport**, **rip** or **staticexport** parameter. For more information, see [“Redistributing External Routes” on page 29-15](#) and the [add ospf redistribute command on page 29-38](#).

The **asexternal** parameter specifies whether the switch acts as an Autonomous System boundary router. A switch is said to be an *Autonomous System* (AS) boundary router when it has some interfaces in the OSPF AS and some interfaces that are not in the AS. The switch typically has “external” routes in its routing table associated with the interfaces that are not in the AS. If **asexternal** is **on**, these external routes are advertised into the AS as type 5 LSAs for non-NSSAs and as type 7 LSAs for NSSAs. If **asexternal** is **nssa**, external routes will only be added to NSSAs as type 7 LSAs, which is translated if appropriate to a type 5 LSA at an NSSA ABR. The **asexternal** parameter should be set to **nssa** if the switch uses only NSSAs. If **asexternal** is **off**, these external routes are not be advertised. The default is **off**.

Important Changing the **rip** and **asexternal** parameters can temporarily disrupt the network’s integrity.

The **autocost** parameter specifies whether or not the switch assigns OSPF interface metrics based on the available interface bandwidth. If an OSPF metric has been manually assigned with the **add ip int=interface ospfmetric=value** command, the manual metric setting takes priority over an automatic metric setting. If the **ospfmetric** parameter has been specified for an interface, it must be restored to its defaults using the command **add ip int=interface ospfmetric=default** before an automatic metric can be applied. The default is **off**.

The **bgpfilter** parameter specifies the route filter to use when importing BGP routes into OSPF. Route filters are created with the [add ip filter](#) command. If a route filter is defined, the entries for the filter include or exclude routes for importation. If routes have not been included by a previous entry, they are excluded from the import. If **none** is specified, no filter is assigned and any existing filter is removed, so all routes from BGP are imported into OSPF. If you change the **bgpfilter**, the device re-imports the BGP routes according to the new filter. The default is **none**. We recommend, however, that you use a BGP route redistribution definition with a routemap instead of this parameter. For more information about creating redistribution definitions, see the [add ospf redistribute command on page 29-38](#). If you use both a BGP route filter and a redistribution definition, the BGP route filter is applied first.

The **bgpimport** parameter specifies whether importing BGP routes into OSPF is enabled. If **yes**, this feature is enabled; if **no**, it is disabled. The default is **no**. For compatibility, **bgpimport** is synchronised with the BGP redistribution definition, which we recommend you use instead. For more information, see [“Redistributing External Routes” on page 29-15](#) and the [add ospf redistribute command on page 29-38](#).



Caution BGP can learn thousands of routes. Do not enable importing of BGP routes into OSPF without considering the network impact of importing these routes. Routing devices in the OSPF domain may become overloaded if they store too many routes. You can prevent this by limiting the number of routes to be imported.

The **bgplimit** parameter specifies the maximum number of BGP routes that can be imported into OSPF at a time. Once this limit is reached, importing stops until existing routes are removed. If you change the **bgplimit**, the device can re-import BGP routes. The default is 1000. For compatibility, **bgplimit** is synchronised with the **limit** parameter of the BGP redistribution definition, which we recommend you use instead. For more information, see [“Redistributing External Routes” on page 29-15](#) and the [add ospf redistribute command on page 29-38](#).

The **defroute** parameter specifies whether a default destination (0.0.0.0) AS external LSA should be generated by this switch, and is only valid when **asexternal** is set to **on** or **nssa**. Specify **on**, **yes** or **true** if you want this switch to generate a default destination (0.0.0.0) AS external LSA. The switch must have routes in its routing table corresponding to the external routing information. The default is **off**.

The **dyninterface** parameter controls the importation of dynamic interface route information. If **stub**, dynamic interface routes are imported into the OSPF routing table as stub routes. If **asexternal**, dynamic interface routes are imported into the OSPF routing table as AS external LSAs. If **none**, dynamic interface routes are not imported. The default is **none**.

The **inroutemap** parameter specifies a route map to filter OSPF routes before adding them to the Routing Information Base (RIB) for IP to use. This route map can match on any of the route map options that are valid for OSPF: interface, prefix list, next hop, route source, metric, route type or tag. See [“Creating Route Maps for OSPF Routes” on page 31-16 of Chapter 31, Filtering IP Routes](#) for more information. Specify **none** to remove an existing routemap.

The **metric** parameter specifies the route metric of the default destination AS external LSA to be generated, and is valid when **defroute** is also present. The default is 1.

The **passiveinterfacedefault** parameter specifies whether all existing OSPF interfaces, and all OSPF interfaces added in the future, are treated as passive interfaces. Specify **on**, **yes** or **true** if you want all existing OSPF interfaces configured as passive OSPF interfaces, and all OSPF interfaces subsequently added using the [add ospf interface command on page 29-31](#) to default to passive. OSPF adds a stub network link to the router LSA if the OSPF routing process can identify the area to which the interface belongs. It does this by finding an area's range that includes the address of the interface. If it finds such a range, that range's area becomes the area for the passive interface. Specify **off**, **no** or **false** if you want all existing OSPF interfaces configured as normal OSPF interfaces, and all OSPF interfaces subsequently added using the [add ospf interface command on page 29-31](#) to default to normal OSPF interfaces. The default is **off**.

You can override the setting of the **passiveinterfacedefault** parameter on individual interfaces using the **passive** parameter of the [add ospf interface](#) and [set ospf interface](#) commands.

The **refbandwidth** parameter specifies the reference bandwidth in Mbits per second used for calculating the OSPF metric. The cost is calculated as **refbandwidth** / Interface Bandwidth (Mbps). Using default settings, the automatic cost calculation results in a OSPF metric of 1 for a fast Ethernet (100M) interface. The **autocost** parameter must be **on** for the **refbandwidth** parameter to take effect. The range is 10 to 10000; the default is 1000.

The **rip** parameter controls the import and export of RIP information. If **export** is specified, OSPF routes are exported as RIP routes. If **import** is specified, RIP information is imported into the OSPF LSA database. If **both** is specified, routing information is both imported from and exported to RIP. If **rip** is **off**, there is no exchange of routing information between OSPF and RIP. The default is **off**. For compatibility, **rip** is synchronised with the RIP redistribution definition, which we recommend you use instead. For more information, see ["Redistributing External Routes" on page 29-15](#) and the [add ospf redistribute command on page 29-38](#).

Important Changing the **rip** and **asexternal** parameters can temporarily disrupt the network's integrity.

The **routerid** parameter specifies a 4-byte number that uniquely identifies the switch in the autonomous system. One scheme for assigning router identification numbers is to choose the largest or smallest IP address assigned to the switch. If **routerid** is never set, the default is to use the highest interface IP address on the switch.

The **ptpstub** parameter allows control over the formation of stub network links. The OSPF RFC states that whenever a numbered point to point link comes up, a stub network to the other end of the link should be added. Each stub network adds an extra link to the switch's LSA. The extra link has no useful purpose, but does increase the LSA size. To limit the LSA size in cases where there are a large number of numbered point to point links, it may be desirable to stop generating stub networks. The default is **yes**, which means stub network links are created. Note that if **ptpstub** is set to **no**, the switch is not strictly compliant with the OSPF RFC but the non-compliance is minor and does not cause problems.

The **staticexport** parameter specifies whether static routing information is exported by this switch. If **yes**, static routes are included in routing exports. If **no**, static routes are omitted. The default is **yes**. For compatibility, **staticexport** is synchronised with the static route redistribution definition, which we

recommend you use instead. For more information, see [“Redistributing External Routes” on page 29-15](#) and the [add ospf redistribute](#) command on [page 29-38](#).

The **type** parameter specifies the LSA type (Type-1 or Type-2) of the default destination AS external LSA to be generated, and is valid only when **defroute** is also present. The default is 2 (Type-2).

Examples To assign a switch the IP address 172.31.1.2 and disable the importing and exporting of RIP information, use the command:

```
set ospf ro=172.31.1.2 rip=off
```

Related Commands

- [disable ospf debug](#)
- [disable ospf log](#)
- [enable ospf debug](#)
- [show ospf](#)

set ospf area

Syntax SET OSPF AREa={BACKbone | *area-number*}
[AUTHentication={NONE | PASSword | MD5}]
[NSSAStability=1..3600] [NSSATranslator={CANDidate |
ALWays}] [STUBArea={ON | OFF | YES | NO | NSSA | True | False}]
[STUBMetric=0..16777215] [SUMmary={SENd | NONE | OFF | NO |
FALSE}]

where *area-number* is a four-byte OSPF area number in dotted decimal notation

Description This command modifies the parameters of an existing OSPF area in the OSPF area table. The **area** parameter is required because it specifies the area number. Other parameters are optional. The area number 0.0.0.0 is reserved for the backbone area, and can be used interchangeably with **backbone**. The specified area must not already exist in the area table.

The **authentication** parameter specifies the type of authentication required for the area. If **none** is specified, no authentication is used on incoming OSPF packets. If **password** is specified, a simple password up to eight characters is used to authenticate each incoming OSPF packet. The password is configured on a per-interface basis. If **md5** is specified, MD5 message digests are used to authenticate OSPF packets. Define the MD5 keys used in the process for every interface using the [add ospf md5key command on page 29-35](#). The default is **none**.

RFC 1583 specifies that authentication be configured per area, but RFC 2328 specifies that authentication be configured per interface. This command represents a hybrid of these solutions in order to allow backward compatibility. An authentication type is set up per area to act as a default for all interfaces in the area. The default setting for interfaces is to use the area default, but each interface can be individually set to any authentication method.

The **nssatranslator** parameter sets the NSSA translator role when the switch is acting as an NSSA border router. If you specify **always**, the switch will always translate Type-7 LSAs to Type-5 LSAs, regardless of the translator state of other border routers in the NSSA, as long as it retains border router status. If it loses border router status it will stop translating Type-7 LSAs until it regains border router status. If you specify **candidate**, the switch will participate in the NSSA translator election process. The NSSA border router with the highest router identifier is elected as the translator. Once elected, the switch will translate Type-7 LSAs until it loses border router status or another NSSA border router with a higher router identifier is elected as the translator. The default is **candidate**. If the switch is acting as a translator it will set the Nt bit in router LSAs it originates into the NSSA. The **nssatranslator** parameter is only valid when **stubarea** is set to **nssa**.

The **nssastability** parameter specifies the additional time, in seconds, that the switch will continue to translate Type-7 LSAs after losing the translator role. An elected translator loses its translator role when another NSSA border router with a higher router identifier is elected as translator, or an NSSA router configured to always translate gains border router status. The time interval allows for a more stable transition to the newly elected translator and minimises excessive flushing of translated Type-7 LSAs. The default is 40. The **nssastability** parameter is only valid when **stubarea** is set to **nssa** and **nssatranslator** is set to **candidate**. Changes to **nssastability** will not take effect until the next time the switch loses its NSSA translator role.

The **stubarea** parameter specifies whether the switch treats the area as a stub area. As external advertisements are not flooded into or out of stub areas. The backbone cannot be configured as a stub area, nor can a virtual link be configured through a stub area. If **area** is set to 0.0.0.0 or **backbone**, **stubarea** must be **off**. If **area** specifies the transit area of a virtual link, **stubarea** must be **off**. All switches within a particular area must have the same setting for **stubarea**. The **nssa** value specifies that the area is a not-so-stubby-area (NSSA). External routes can be imported as type 7 advertisements in a NSSA. The default is **off** when **area** is set to 0.0.0.0 or **backbone**; or **on** when **area** is set to any other (non-backbone) value.

If the area has been configured as a stub area, and the switch is to act as the area border router for the stub area, then the **stubmetric** parameter specifies the metric (cost) of the default route as advertised by the switch in the default summary link. The default is 1.

The **summary** parameter controls the generation of summary LSA's into stub areas. By default, the default (0.0.0.0) summary LSA is emitted into a stub area by an area border router. If **send**, summary LSA's from other areas are also emitted into the stub area. If **none**, the default (0.0.0.0) summary LSA is emitted into the stub area. If **stubarea** is **nssa**, then the default is **send**; otherwise, the default is **none**.

Examples To change area 1 into a stub area, use the command:

```
set ospf are=0.0.0.1 stuba=on stubm=10
```

Related Commands

- [add ospf area](#)
- [add ospf range](#)
- [delete ospf area](#)
- [delete ospf range](#)
- [set ospf range](#)
- [show ospf area](#)
- [show ospf range](#)

set ospf host

Syntax SET OSPF HOSt=*ipadd* METric=0..65535

where *ipadd* is an IP address in dotted decimal notation

Description This command changes the metric of a static OSPF host route in the routing table.

OSPF allows host routes to be advertised to the local OSPF area. Such routes are advertised with a route mask of 255.255.255.255. Host routes are normally used for point-to-point networks on which it is undesirable to run OSPF with IP hosts directly connected to the switch.

The **host** parameter specifies the IP address of the host or point-to-point network to be modified. The host must exist in the routing table.

The **metric** parameter specifies the metric for the route to the host.

Examples To change the metric for OSPF host 172.30.1.2 to 5, use the command:

```
set ospf ho=172.30.1.2 met=5
```

Related Commands

- [add ospf host](#)
- [delete ospf host](#)
- [show ospf host](#)

set ospf interface

Syntax SET OSPF INTERface=*interface* [AREa={BACKbone|*area-number*}]
 [AUTHentication={AREadefault|NONE|PASSword|MD5}]
 [BOOST1=0..1023] [DEadinterval=2..2147483647]
 [DEMAND={ON|OFF|YES|NO|True|False}]
 [HELLOinterval=1..65535] [NETwork={BROadcast|
 NON-broadcast}] [PASSive={ON|OFF|YES|NO|True|False}]
 [PASSword={NONE|*password*}] [POLLIinterval=1..2147483647]
 [PRIOrity=0..255] [RXminterval=1..3600]
 [TRAnsitdelay=1..3600] [VIRtuallink=*router-id*]

where:

- *interface* is a valid interface name.
- *area-number* is a 4-byte OSPF area number in dotted decimal notation.
- *password* is a string 1 to 8 characters long. Valid characters are any printable character. If *password* contains spaces, then it must be in double quotes.

Description This command modifies the parameters of an OSPF interface.

The **interface** parameter specifies the name of the OSPF interface to add, and must be a defined IP interface, or a valid virtual interface instance. Valid interfaces are:

- PPP (such as ppp0)
- VLAN (such as vlan1)
- FR (such as fr0)
- virtual interface (such as virt9)

To see a list of current valid interfaces, use the [show ospf interface command on page 29-85](#) or the [show interface command on page 11-87 of Chapter 11, Interfaces](#).

Disabled ghost OSPF interfaces exist for each IP interface that is attached to OSPF for use with SNMP. These interfaces must be added to OSPF with the [add ospf interface command on page 29-31](#) before OSPF can use them. Deleting the interface with the [delete ospf interface command on page 29-43](#) turns it back into a disabled ghost OSPF interface.

Important Addressless PPP links are the only exception to the range checking mechanism since there is no range to check them against.

The **area** parameter specifies the area where the interface belongs. When defining a virtual link, the area number is the area number of the transit area used for the virtual link. The area number 0.0.0.0 is reserved for the backbone area, and can be used interchangeably with **backbone**. The specified area must already exist in the area table.

The **authentication** parameter specifies the type of authentication required for the interface. If **areadefault** is specified the authentication type is taken from the area setting. If **none** is specified, no authentication is used on incoming OSPF packets. If **password** is specified, a simple password up to eight characters in length is used to authenticate each incoming OSPF packet. The password is configured for each interface. If **md5** is specified, MD5 message digests are used to authenticate OSPF packets. Define the MD5 keys used in the

process for every interface using the [add ospf md5key command on page 29-35](#). The default is **areadefault**.

RFC 1583 specifies that authentication be configured per area, but RFC 2328 specifies that authentication be configured per interface. This command represents a hybrid of these solutions, to allow backward compatibility. An authentication type is set up per area to act as a default for all interfaces in the area. The default setting for interfaces is to use the area default, but each interface can be individually set to any authentication method.

The **boost1** parameter increases the type 1 metrics of all OSPF LSAs flooded out the interface by the value specified. Setting the value to anything other than 0 increases the LSA metrics by that amount. For example, setting the parameter in the switch to **boost1=2** boosts the metrics by 2, so an LSA in this switch with a metric of 3 has a value of 5 when it is flooded out the interface. The default is 0.



Caution The **boost1** parameter is a departure from the OSPF specification so should be used with extreme caution.

The **deadinterval** parameter specifies the interval in seconds after which no hello packets are received by a neighbour, the neighbour declares the switch to be down. The timer is advertised in the hello packets. Switches on the same network where the interface is attached must have the same **deadinterval** timer. The value must be at least twice the value of the **hellointerval** timer. The recommended multiplier is four. For example, if the **hellointerval** timer is set to 10 seconds, then the **deadinterval** timer should be set to 40 seconds. The default is four times the value of the **hellointerval** timer.

The **demand** parameter specifies whether the interface connects to a demand circuit. Two switches connecting to the same common network segment need not agree on that segment's demand circuit status. This means that configuring one switch does not require configuring other switches that connect to the same common network segment. If one switch has been configured and the common network is a broadcast or non-broadcast multi-access (NBMA) network, the behaviour (e.g. sending, receiving hello packets) of the network remains the same as if the interface has not been configured as a demand circuit. If one switch has been configured and the common network segment is a point-to-point link, the switch on the other end may agree to treat the link as a demand circuit and the point-to-point network receives the full benefit. When broadcast and non-broadcast multi-access (NBMA) networks are declared as demand circuits (i.e. more than one switch has the network configured as a demand circuit), routing update traffic is reduced but the periodic sending of hellos is not, which requires that the data link connection remain constantly open. The default is **off**.

Virtual links created using the **virtuallink** parameter are always treated as demand circuits, and ignore the setting of the **demand** parameter.

The **hellointerval** parameter specifies the interval in seconds between hello packets transmitted over the interface. Switches on the network to which the interface is attached must have the same **hellointerval** timer. The value must be less than the value of the **pollinterval** timer. The default is 10.

The **network** parameter specifies the OSPF network type of the interface, and is only valid for VLAN interfaces. Specify **broadcast** if you want OSPF to treat the network as a broadcast network. Hello messages are transmitted as broadcast messages, and neighbours are learned dynamically. You can not configure static neighbours or use the **pollinterval** parameter to set the time interval

between hello messages to inactive neighbours. Specify **non-broadcast** if you want OSPF to treat the network as an NBMA network. All OSPF packets are transmitted as unicast messages, so neighbours must be statically defined. You can use the **pollinterval** parameter to set the time interval between hello messages to inactive neighbours. The default is **broadcast**.

When you change the network type of a VLAN interface from **broadcast** to **non-broadcast**:

- All OSPF packets are sent as unicast messages, not broadcast messages, so neighbours need to be statically configured.
- Any existing dynamically learned neighbours are automatically converted to static neighbours.
- Hello messages are not transmitted until at least one static neighbour exists.

When you change the network type of a VLAN interface from **non-broadcast** to **broadcast**:

- Any existing statically defined neighbours are cleared.
- Hello messages are sent as broadcast messages, so neighbours are dynamically learned.

The **passive** parameter specifies whether the interface acts as a passive interface. Specify **on**, **yes** or **true** if you want the interface to act as a passive OSPF interface. The interface does not take part in the OSPF protocol and OSPF packets are not sent or received on the interface. However, the network that the interface belongs to is added to the router LSA as a stub network, so that routing information about the network will be carried through OSPF to the rest of the routing domain. Specify **off**, **no** or **false** if you want the interface to act as a normal OSPF interface. The default is the value of the **passiveinterfacedefault** parameter of the [set ospf command on page 29-59](#), or **off** if the **passiveinterfacedefault** parameter has not been set.

The **password** parameter specifies the password used for authentication. A password is required if the authentication scheme for the area has been set to **password** with the [add ospf area command on page 29-28](#) or the [set ospf area command on page 29-63](#). If **none** is specified, no password is configured on the interface, and any previously set password is removed. The default is **none**.

The **pollinterval** timer defines the time in seconds when hello packets are sent to neighbouring switches that are deemed to be inactive. The **pollinterval** timer is used on NBMA networks (Frame Relay, X.25) and point-to-point networks configured as OSPF on demand circuits. The value must be greater than the value of the **hellointerval** timer. When the neighbour is in state "Down", hellos are sent to the neighbour at the interval of **pollinterval**. The default is four times the value of the **hellointerval** timer.

The **priority** parameter is used on multi-access networks to set the router priority. When two switches attached to a network attempt to become the designated router, the one with the highest priority takes precedence. If the priorities are the same then the switch with the highest router identification number takes precedence. A switch with a priority of zero is ineligible to become the designated router. Router priority can be configured only for switches attached to multi-access networks. The default is 1.

The **rxmtinterval** parameter specifies the interval in seconds between link state retransmissions, for adjacencies on the interface. The timer is also used when

retransmitting database description and link state request packets. The value should be set well above the round trip time between the two switches. It should be set higher for slow serial and virtual links. A typical value is 5 seconds on a local area network. The default is 5.

The **transitdelay** parameter specifies the estimated time in seconds that is required to transmit a link state update packet over this interface. This time is added to link state advertisements sent over the interface. The value of this parameter should take into account the transmission and propagation delays of the interface. This mechanism helps keep the link state advertisement timers synchronised on different switches. A typical value is 1 second on a local area network. The default is 1.

The **virtuallink** parameter specifies the router identification number for another area border router to be included in the backbone using a virtual link. Each end of the virtual link must be configured. This parameter is required when a virtual interface is added to OSPF. Virtual links are always treated as demand circuits. This is because when a virtual link's underlying physical path contains one or more demand circuits, periodic OSPF protocol exchanges over the virtual link would unnecessarily keep the underlying demand circuits open. When configuring virtual links, the **area** parameter specifies the area number of the transit area used for the virtual link.

Examples To set the **hellointerval** to 20 and the **deadinterval** to 80 on OSPF interface ppp0, use the command:

```
set ospf int=ppp0 he=20 de=80
```

To set the OSPF interface PPP0 to a demand circuit over the point-to-point link, use the command:

```
set ospf int=ppp0 dem=on
```

Related Commands

- [set ospf area](#)
- [add ospf interface](#)
- [add ospf range](#)
- [delete ospf interface](#)
- [disable ospf interface](#)
- [enable ospf interface](#)
- [reset ospf interface](#)
- [set ospf area](#)
- [set ospf range](#)
- [show ospf area](#)
- [show ospf interface](#)
- [show ospf range](#)

set ospf neighbour

Syntax SET OSPF NEighbour=*ipadd* PRIOrity=0..255

where *ipadd* is an IP address in dotted decimal notation

Description This command modifies the parameters of a non-broadcast multi-access neighbour. On non-broadcast multi-access networks, such as Frame Relay, the neighbours cannot be discovered by dynamic means. To overcome this problem, neighbours are configured statically. All parameters are required to modify a neighbour.

Static OSPF neighbours should be defined on switches that are attached to an OSPF network and are eligible to become the *Designated Router* (DR) or *Backup Designated Router* (BDR) for that network. To be eligible to become the DR or BDR for an OSPF network, a switch must have at least one interface on the OSPF network with a non-zero priority. The priority of an OSPF interface is set with the **priority** parameter of the [add ospf interface command on page 29-31](#) or the [set ospf interface command on page 29-66](#).

Switches connected to an OSPF network via an OSPF interface with a non-zero priority should be configured with static OSPF neighbours corresponding to all other OSPF routers attached to the OSPF network. OSPF routers without an OSPF interface with a non-zero priority should not be configured with static OSPF neighbours.

The **neighbour** parameter specifies the IP address of the neighbour. The IP address must be within an IP address range associated with an OSPF interface. The specified neighbour must exist in the neighbour table.

The **priority** parameter specifies the priority of the neighbour router. If the priority for the neighbour is set to zero then it is not initially considered eligible to become the designated router. The priority for the switch itself is set in the interface priority.

Examples To change the priority of the neighbour with IP address 172.30.1.2 to 45, use the command:

```
set ospf nei=172.30.1.2 prio=45
```

Related Commands [add ospf neighbour](#)
[delete ospf neighbour](#)
[show ospf neighbour](#)

set ospf range

Syntax SET OSPF RANge=*ipadd* [AREa={BAckbone|*area-number*}]
[MASK=*ipadd*] [EFFect={ADVertise|DONotadvertise}]

where:

- *ipadd* specifies an IP address in dotted decimal notation.
- *area-number* is a 4-byte OSPF area number in dotted decimal notation.

Description This command modifies an OSPF range. An OSPF area is defined as a list of IP address ranges. A IP address range is defined by an IP address and network mask pair. After an area has been created, the ranges that are to be active in the area must be defined. All the OSPF routers within an area must use the same set of ranges.

Configuring a range on an OSPF ABR enables the ABR to aggregate the network advertisements from one area into another area in the form of summary LSAs. However, networks advertised into a transit area should not be aggregated into summary LSAs, so the switch ignores any configured range when it is advertising into a transit area.

The **range** parameter specifies the base IP address of the range to modify.

The **area** parameter specifies the area number of the area to which the range belongs. The area number 0.0.0.0 is reserved for the backbone area, and can be used interchangeably with **backbone**. The specified area must already exist in the area table. If the area to which the range belongs changes, the interfaces that belong to the range also change.

The **mask** parameter specifies a mask used in combination with the base address to determine addresses that belong to the range. The value must be consistent with the value specified for the **range** parameter. The default is the network mask for the address class of the IP address (e.g. 255.255.0.0 for a Class B address and 255.255.255.0 for a Class C address).

The range defined by the combination of **range** and **mask** must not overlap any other range defined on the switch. Either **area** or **mask** must be specified.

The **effect** parameter controls the exporting of ranges into other areas by an area border router. This parameter is valid when the switch is configured as an area border router, and can be used to control where the traffic into and out of an area for a particular network travels. If **advertise**, summary LSAs are created for the range; if **donotadvertise**, they are not created. The default is **advertise**.

Examples To change an address range of 172.30.0.0–172.30.255.255 into an address range of 172.30.0.0–172.30.0.255, use the command:

```
set ospf ran=172.30.0.0 mask=255.255.255.0
```

Related Commands [add ospf range](#)
[delete ospf range](#)
[show ospf range](#)

set ospf redistribute

Syntax SET OSPF REDistribute PROTOcol={BGP|INTERface|RIP|STATIC}
 [LIMit=1..4000] [METric={0..16777214|ORiginal}]
 [ROUTEMap={*routermap*|NONE}] [SUBNET={ON|OFF|YES|NO|True|False}] [TAG={1..65535|ORiginal}] [TYpe={1|2|ORiginal}]

where *routermap* is the name of an IP route map

Description This command modifies a redistribution definition for OSPF to use when it imports BGP, RIP, non-OSPF interface, and statically configured routes into AS external or NSSA AS external LSAs that it advertises. Only one redistribution definition can be created for each route type.

Modifying a BGP, RIP, or static interface route redistribution definition will change the setting of the **bgpimport**, **rip** and **staticexport** parameters of the **set ospf** command on page 29-59. For more information, see “Redistributing External Routes” on page 29-15. If you configure a BGP route filter using the **bgpfilter** parameter of the **set ospf** command, the filter will be applied before any BGP route redistribution definition.

The **protocol** parameter specifies the route redistribution definition to modify. Specify **bgp** or **rip** to modify the redistribution definition for routes derived from BGP or RIP. Specify **interface** to modify the redistribution definition for non-OSPF interface routes. Specify **static** to modify the redistribution definition for statically configured routes.

The **limit** parameter specifies the maximum number of routes that can be redistributed into OSPF for the specified protocol. The **limit** parameter overwrites the setting of the **bgplimit** parameter in the **set ospf** command on page 29-59.

The **metric** parameter specifies the route metric that OSPF assigns to routes that it redistributes. If you specify **original**, the original route metric is preserved in the redistributed route—metric1 for Type-1 routes or metric2 for Type-2 routes. If you assign a route map that sets the metric, the route map overrides the setting in this parameter. The default is 20.

The **routermap** parameter specifies a route map to filter routes before redistributing them through OSPF. This route map can match on interface, prefix list, next hop, metric, or tag. The route map must already exist. To create a route map, use the **add ip routermap** command on page 31-49 of Chapter 31, *Filtering IP Routes*. Specify **none** to remove an existing routermap. If you do not specify a route map, OSPF will import all routes for the protocol, depending on the setting of the **subnet** parameter.

The **subnet** parameter specifies whether OSPF redistributes subnets. Specify **off** if you want OSPF to only redistribute classfull network routes. Specify **on** if you want OSPF to redistribute classless network routes as well. The default is **on**.

The **tag** parameter specifies a number OSPF uses to label routes that it redistributes. If you specify **original**, the original route tag is preserved in the redistributed route. If you assign a route map that sets the tag, the route map overrides the setting in this parameter. The default is **original**.

The **type** parameter specifies the OSPF external route type which OSPF assigns to routes that it redistributes. Use the **type** parameter to ensure that all externally-sourced OSPF routes are the same type and therefore use the same method to calculate route metrics. Specify **1** if you require the routes to have a Type-1 external metric, or **2** if you require the routes to have a Type-2 external metric. If you assign a route map that sets the type, the route map overrides the setting in this parameter. See [“Path types” on page 29-9](#) for more information about the types. The default is **2**.

Example To change the metric from the default of 20 to 30 for static routes that OSPF redistributes, use the command:

```
set ospf red prot=sta met=30
```

Related Commands [add ip routemap in Chapter 31, Filtering IP Routes](#)
[add ospf redistribute](#)
[delete ospf redistribute](#)
[show ospf redistribute](#)

set ospf stub

Syntax SET OSPF STUB=*ipadd* MASK=*ipadd* METric=0..65535

where *ipadd* is an IP address in dotted decimal notation

Description This command changes the metric of a static non-OSPF stub network in the OSPF routing table. Networks, such as an LAN with only one OSPF router, can be added as stub networks to OSPF without running OSPF on the interface to which the network is attached. The stub links are added to the switch's LSA.

The **stub** parameter specifies the base IP address of the stub network. The IP address must fall within one of the OSPF ranges defined on the switch. The stub must not already exist in the routing table.

The **mask** parameter specifies a mask used in combination with the base address to determine addresses that belong to the stub network. The value must be consistent with the value specified for the **stub** parameter.

The **metric** parameter specifies the metric for the route to the stub network.

Examples Ethernet interface vlan1 is attached to an Ethernet LAN with an IP subnet address of 172.30.1.0 and network mask 255.255.255.0. To add the LAN attached to the Ethernet interface as a non-OSPF stub network, use the command:

```
add ospf stub=172.30.1.0 mask=255.255.255.0
```

The stub network is added to the OSPF routing table with a default metric of 1. To change the metric to 5, use the command:

```
set ospf stub=172.30.1.0 mask=255.255.255.0 met=5
```

Related Commands

- [add ospf stub](#)
- [delete ospf stub](#)
- [set ospf host](#)
- [set ospf interface](#)
- [show ospf stub](#)

set ospf summaryaddress

Syntax SET OSPF SUMMmaryaddress=*ipadd* [MASK=*ipadd*] [ADVertise={ON|OFF|YES|NO|True|False}] [TAG=0..65535]

where *ipadd* is an IP address in dotted decimal notation

Description This command modifies a summary address entry in OSPF, configuring a single route to advertise (or not) in an OSPF AS external LSA. When OSPF redistributes routes, networks that fall within the range of this summary address are advertised by a single AS external LSA containing this summary address.

The AS external LSAs for static routes have their metric and metric type set by the parameters in the OSPF redistribution definition for the protocol. See [“Redistributing External Routes” on page 29-15](#).

The AS external LSAs for routes imported from BGP and RIP have their metric and metric type set by the metric and metric type of the original route.

The **summaryaddress** parameter specifies the base IP address of the summary address to modify.

The **mask** parameter specifies a mask used in combination with the base IP address to define a summary address range.

The **advertise** parameter specifies whether OSPF advertises the summary address in an AS external LSA for this network range. Specify **on**, **yes** or **true** if you want OSPF to advertise the summary address in AS external LSAs. Individual networks that fall within the range of the summary address are not advertised. Specify **off**, **no** or **false** if you do not want OSPF to advertise either the summary address or the individual networks within the range of the summary address. The default is **on**.

The **tag** parameter specifies the tag value that OSPF places in the AS external LSAs created as a result of redistributing the summary route. The tag overrides tags set by the original route and by the route map used to select the original routes for redistribution. If you do not specify a value, OSPF places a tag value of 0 in the AS external LSAs.

Examples To stop advertising the summary address entry for all class C networks matching 198.168.1.*, use the command:

```
add ospf summ=192.168.1.0 adv=off
```

Related Commands

- [add ospf summaryaddress](#)
- [delete ospf summaryaddress](#)
- [show ospf summaryaddress](#)

show ospf

Syntax SHow OSPF

Description This command displays information about the general configuration of OSPF routing (Figure 29-6, Table 29-2). This display relates to the MIB group ospfGeneralGroup.

Figure 29-6: Example output from the **show ospf** command

```
Router ID ..... 123.234.143.231
OSPF module status ..... Enabled
Area border router status ..... Yes
NSSA border router status ..... No
AS boundary router status ..... Disabled
PTP stub network generation ..... Enabled
External LSA count ..... 10234
External LSA sum of checksums ... 1002345623
New LSAs originated ..... 10345
New LSAs received ..... 34500
RIP ..... Off
BGP importing:
  Enabled ..... Yes
  Import filter ..... 301
  Routes imported/limit ..... 1763 / 4000
Passive Interface Default ..... No
In Route Map ..... -
Export static routes ..... Yes
Dynamic interface support ..... None
Number of areas added ..... 0
Number of active areas ..... 10
Logging ..... Disabled
Debugging ..... Disabled
AS external default route:
  Status ..... Disabled
  Type ..... 1
  Metric ..... 1
System resources used:
  CPU percentage ..... 5.7%
  Memory ..... 32kB
```

Table 29-2: Parameters in output of the **show ospf** command

Parameter	Meaning
Router ID	Unique OSPF router identification number for this switch.
OSPF module status	Whether OSPF routing is enabled.
Area border router status	Whether the switch is acting as an area border router.
AS boundary router status	Whether the switch is enabled to act as an autonomous system boundary router, one of "Enabled", "Disabled" or "NSSA".
PTP stub network generation	Whether stub network links are created when numbered point-to-point links come up.
External LSA count	Number of external link state advertisements in the topological database.

Table 29-2: Parameters in output of the **show ospf** command (cont)

Parameter	Meaning
External LSA sum of checksums	32-bit checksum of the external link state advertisements. This checksum can be used to compare the topological databases on two different switches. When they have the same topology, the checksum is the same.
New LSAs originated	Number of new link state advertisements that have originated from this switch.
New LSAs received	The number of link state advertisements received that are new or new instances of link state advertisements.
RIP	Whether RIP information is imported into or exported out of the OSPF routing domain; either "Off", "Import", "Export" or "Import/export".
BGP importing	Information about the importing of BGP routes into OSPF.
Enabled	Whether importing BGP routes into OSPF is enabled.
Import filter	IP filter number used to filter routes before they are imported into OSPF, or whether no filters are used.
Routes imported/limit	The number of BGP routes imported into OSPF, and the maximum number of routes that can be imported at a time.
Passive interface default	Whether all OSPF interfaces are treated as passive OSPF interfaces. For more information, see "Passive Interfaces" on page 29-12 .
In Route Map	The route map used to filter OSPF routes being added to the Routing Information Base (RIB) for IP to use, or "-" if a route map has not been set.
Export static routes	Whether static routing information is exported to OSPF routing domain.
Dynamic interface support	Whether dynamic interface routes are imported into the OSPF routing table: Stub AS external None Undefined
Number of areas added	Number of areas added using the add ospf area command.
Number of active areas	Number of OSPF areas defined on the switch.
Auto BW cost calculation	Whether or not the switch will assign OSPF interface metrics based on the available interface bandwidth.
Auto BW cost reference	The reference bandwidth in Mbps used for calculating the OSPF metric.
Logging	Whether OSPF logging is enabled or disabled.
Debugging	Whether OSPF debugging is enabled or disabled.
AS external default route	Information about the generation of a default destination (0.0.0.0) AS external LSA.
Status	Whether the generation of a default destination (0.0.0.0) AS external LSA is enabled or disabled.
Type	LSA type for the default destination (0.0.0.0) AS external LSA; either "1", "2", or "Undefined".
Metric	Metric for the default destination (0.0.0.0) AS external LSA.

Table 29-2: Parameters in output of the **show ospf** command (cont)

Parameter	Meaning
System resources used	Information about the level of system resources used by the OSPF routing protocol.
CPU percentage	A 5 minute running average of the amount of CPU time used by the OSPF routing protocol.
Memory	The amount of memory, in kilobytes, used in the OSPF LSA database and internal route tables.

Related Commands [enable ospf log](#)
[set ospf](#)

show ospf area

Syntax `SHoW OSPF AREa [= {Backbone | area-number}] [{FULl | SUMmary}]`

where *area-number* is a 4-byte OSPF area number in dotted decimal notation

Description This command displays OSPF area information. The **area** parameter specifies the area number. The area number 0.0.0.0 is reserved for the backbone area, and can be used interchangeably with **backbone**. The specified area must exist in the area table. If an area is not specified, a summary of all OSPF areas is displayed (Figure 29-7, Table 29-3). If an area is specified, detailed information about the specified area is displayed (Figure 29-8 on page 29-80, Table 29-4 on page 29-80). This display relates to the MIB entity ospfAreaTable.

The **full** and **summary** parameters override the default behaviour. The **full** parameter displays detailed information about the specified area or all areas. The **summary** parameter displays summary information about the specified area or all areas.

Figure 29-7: Example output from the **show ospf area** command

Area	State	Authentication	StubArea	StubMetric	Summary LSAs
Backbone	Active	Password	No	4	Send
0.0.0.1	Active	Password	Yes	5	None
0.0.0.2	Inactive	None	Yes	20001	None

Table 29-3: Parameters in output of the **show ospf area** command

Parameter	Meaning
Area	Unique 32-bit number for the area, in dotted decimal notation.
State	Whether the area is active or inactive. This depends on having active ranges and interfaces defined for the area.
Authentication	Whether the type of authentication for incoming OSPF packets to this area is password or none.
StubArea	Whether the area imports external link state advertisements. This field relates to the MIB variable ospfImportASExtern.
StubMetric	If the area is a stub area, this is the default metric of the default summary link advertised into the area.
Summary LSAs	Whether summary LSAs other than the default (0.0.0.0) are emitted into the area; either Send, None, or Undefined.

Figure 29-8: Example output from the **show ospf area** command for a specific area

```

Area 0.0.0.1:
  State ..... Active
  Authentication .... Password
  Stub area ..... No
  Stub cost ..... 1
  NSSA ..... Yes
    Role ..... CANDIDATE
    Stability Interval ..... 40
    State ..... DISABLED
  Summary LSAs ..... Send
  SPF runs ..... 23
  Area border router count ..... 3
  AS border router count ..... 2
  LSA count ..... 10
  LSA sum of checksums ..... 345bf

Ranges:
  Range ..... 192.168.25.0
    Mask ..... 255.255.255.0
  Range ..... 192.168.250.0
    Mask ..... 255.255.255.0

Interfaces:
  ppp23:
    Type ..... Point to point
    State ..... ptp

```

Table 29-4: Parameters in output of the **show ospf area** command for a specific area

Parameter	Meaning
Area	Unique 32-bit number for the area, in dotted decimal notation.
State	Whether the area is active or inactive. This depends on having active ranges and interfaces defined for the area.
Authentication	Whether authentication for incoming OSPF packets to this area is password, MD5, or none.
Stub area	Whether the area imports external link state advertisements. This field relates to the MIB variable ospfImportASExtern.
Stub cost	If the area is a stub area, this is the default metric of the default summary link advertised into the area.
NSSA	Whether the area is a not-so-stubby-area (NSSA).
Role	NSSA translator role; one of "CANDIDATE" or "ALWAYS". This field is only displayed when NSSA is "Yes".
Stability Interval	Time period, in seconds, that the switch will continue to translate Type-7 LSAs after losing its elected translator role to another NSSA border router. This field is only displayed when NSSA is "Yes".
State	Current NSSA translator state. If Role is "ALWAYS", one of "DISABLED" or "ENABLED". If Role is "CANDIDATE", one of "DISABLED" or "ELECTED". This field is only displayed when NSSA is "Yes".

Table 29-4: Parameters in output of the **show ospf area** command for a specific area

Parameter	Meaning
Summary LSAs	Whether summary LSAs other than the default (0.0.0.0) are emitted into the area; either Send, None, or Undefined.
SPF runs	Number of times the intra-area routing table has been recalculated using the topological database.
Area border router count	Total number of area border routers reachable within the area.
AS border router count	Total number of Autonomous System boundary routers reachable within the area.
LSA count	Total number of link state advertisements in this area's topological database, excluding external LSA's.
LSA sum of checksums	32-bit checksum of the checksums of the area's link state advertisements. This checksum can be used to compare the topological databases on two different switches.
Ranges	Information about the ranges active for this area.
Range	Base IP address of a range.
Mask	Mask for a range.
Interfaces	Interfaces in use by the area.
Type	Interface type: Unknown Broadcast NBMA Point to Point Virtual
State	Status of the interface: Unknown Down Loopback Waiting Ptp DR BackupDR OtherDR

Examples To display a summary of all configured OSPF areas, use the command:

```
sh ospf are
```

To display detailed information about all configured OSPF areas, use the command:

```
sh ospf are ful
```

Related Commands

- [add ospf area](#)
- [add ospf range](#)
- [delete ospf area](#)
- [delete ospf range](#)
- [reset ospf counter](#)
- [set ospf area](#)
- [set ospf range](#)
- [show ospf range](#)

show ospf debug

Syntax SHow OSPF DEBUg [= {ALL | NRL | PERFormance | SPF | TIMer}]

Description This command displays internal debugging information for the OSPF routing module. It is intended for use under the direction of a technical support engineer.

Related Commands [disable ospf debug](#)
[enable ospf debug](#)
[disable debug active](#) in Chapter 4, Configuring and Monitoring the System
[show debug active](#) in Chapter 4, Configuring and Monitoring the System

show ospf host

Syntax `SHoW OSPF HOst [=ipadd] [AREa={BAckbone | area-number}]`

where:

- *ipadd* is an IP address in dotted decimal notation.
- *area-number* is a 4-byte OSPF area number in dotted decimal notation.

Description This command displays information about statically configured OSPF host routes (Figure 29-9, Table 29-5). This display relates to the MIB entity ospfHostTable.

The **host** parameter specifies the OSPF host route to display. Host routes whose base IP addresses match the specified IP address are displayed. Wildcard addresses with zeros in the right-hand position may be used to match multiple host routes. For example, the value 172.16.0.0 matches (and displays) all OSPF host routes whose base address begins with 172.16.

The **area** parameter specifies the area for which host route information is to be displayed. Host routes associated with the specified area are displayed. Wildcard addresses with zeros in the right-hand position may be used to match multiple areas. For example, the value 172.16.0.0 matches (and displays) OSPF host routes associated with areas whose area number begins with 172.16. The area number 0.0.0.0 is reserved for the backbone area, and can be used interchangeably with **backbone**.

Figure 29-9: Example output from the **show ospf host** command

IP address	Mask	State	Area	Metric	TOS	Type
192.188.150.2	255.255.255.255	Inactive	0.0.0.24	20000	0	Stat
192.188.250.235	255.255.255.255	Active	123.125.230.156	10	0	Stat
192.188.125.240	255.255.255.255	Active	0.0.0.1	1	0	Stat

Table 29-5: Parameters in output of the **show ospf host** command

Parameter	Meaning
IP address	IP address of the host or point-to-point network.
Mask	Mask for the host route.
State	Whether the host route entry is active or inactive. If active, it is being advertised via the switch LSA.
Area	Area number of the area containing the host.
Metric	Metric to be advertised for the host.
TOS	Type of service of the route to the host.
Type	Whether the type of host route entry is permanent static or dynamic.

Examples To display all OSPF host routes, use the command:

```
sh ospf ho
```

To display information for all host routes from 172.30.0.0 to 172.30.255.255 in area 0.0.0.3, use the command:

```
sh ospf ho=172.30.0.0 are=0.0.0.3
```

Related Commands [add ospf host](#)
[delete ospf host](#)
[set ospf host](#)

show ospf interface

Syntax `SHoW OSPF INTeRface[=interface] [AREa={BACkbone|
area-number}] [IPaddress=ipadd] [{FULl|SUMmary}]`

where:

- *interface* is a valid interface name.
- *area-number* is a 4-byte OSPF area number in dotted decimal notation.
- *ipadd* is an IP address in dotted decimal notation.

Description This command displays information about OSPF interfaces. If an interface is not specified, summary information for all OSPF interfaces attached to OSPF is displayed ([Figure 29-10 on page 29-86](#), [Table 29-6 on page 29-86](#)). If an interface is given, detailed information about the specified OSPF interface is displayed ([Figure 29-11 on page 29-87](#), [Table 29-7 on page 29-87](#)). This display relates to the MIB entity ospfIfTable.

The **interface** parameter specifies a valid interface already assigned and configured. Valid interfaces are:

- PPP (such as ppp0)
- VLAN (such as vlan1)
- FR (such as fr0)
- virtual interface (such as virt9)

Disabled ghost OSPF interfaces exist for each IP interface that is attached to OSPF for use with SNMP. These interfaces must be added to OSPF with the [add ospf interface command on page 29-31](#) before OSPF can use them. Deleting the interface with the [delete ospf interface command on page 29-43](#) turns it back into a disabled ghost OSPF interface.

The **area** parameter specifies the area for which interface information is to be displayed. Interfaces associated with the specified area are displayed. Wildcard addresses with zeros in the right-hand position may be used to match multiple areas. For example, the value 172.16.0.0 matches (and displays) OSPF interfaces associated with areas whose area number begins with 172.16. The area number 0.0.0.0 is reserved for the backbone area, and can be used interchangeably with **backbone**.

The **ipaddress** parameter specifies the IP address for which interface information is to be displayed. Interfaces with the specified address are displayed. Wildcard addresses with zeros in the right-hand position may be used to match multiple interfaces. For example, the value 172.16.0.0 matches (and displays) all OSPF interfaces with addresses beginning 172.16.

The parameters **full** and **summary** can override the default behaviour. The **full** parameter displays detailed information about the specified interface or all interfaces. The **summary** parameter displays summary information about the specified interface or all interfaces.

Figure 29-10: Example output from the **show ospf interface** command

Iface	Status	Area	State	Designated rtr / Virtual nbr	Backup DR / Transit area
fr0	Enabled	Backbone	DR	192.168.0.3	192.168.0.4
ppp23	Disabled	0.0.0.1	ptp	-	-
virt0	Enabled	Backbone	ptp	192.168.250.5	0.0.0.1

Table 29-6: Parameters in output of the **show ospf interface** command

Parameter	Meaning
Iface	Interface name.
Status	Whether the administrative status of the interface is enabled.
Area	Area with which the interface is associated.
State	The state of the interface: Unknown Down Loopback Waiting Ptp DR BackupDR OtherDR
Designated rtr	Current designated router for a broadcast or non-broadcast multi-access network, or "None" if a designated router has not yet been selected.
Backup DR	Current backup designated router for a broadcast or non-broadcast multi-access network, or "None" if a backup designated router has not yet been selected.
Virtual nbr	Router identification number of the virtual neighbour for virtual interfaces.
Transit area	Transit area used for virtual interfaces.

Figure 29-11: Example output from the **show ospf interface** command for a specified interface

```

vlan1:
  Status ..... Enabled
  Area ..... Backbone
  IP address ..... 192.168.250.1
  IP net mask ..... 255.255.255.0
  IP network number ..... 192.168.250.0
  Type ..... broadcast
  OSPF on demand ..... ON (OFF)
  Passive ..... No
  State ..... otherDR
  Router priority ..... 5
  Transit delay ..... 1 second
  Retransmit interval ..... 5 seconds
  Hello interval ..... 10 seconds
  Router dead interval ..... 40 seconds
  Poll interval ..... 120 seconds
  Interface events ..... 1
  Authentication ..... Password (area default)
  Password ..... Charlie1
  Designated router ..... 192.168.250.254
  Backup designated router ..... 192.168.250.253
  Metric boost 1 ..... 0

```

Table 29-7: Parameters in output of the **show ospf interface** command for a specific interface

Parameter	Meaning
Status	Whether the administrative status of the interface is enabled.
Area	Area with which the interface is associated.
IP address	IP address assigned to the interface when the interface is not an addressless interface. This is defined by the IP interface where OSPF is attached.
IP net mask	IP network mask for the associated IP interface, if the interface is not an addressless interface.
IP network number	IP network address for the interface when the interface is not an addressless interface. The network address is determined from the IP address and network mask.
Addressless index	Addressless interface index for addressless interfaces.
Type	Type of network associated with the interface: Broadcast NBMA (non-broadcast multi-access) Point-to-Point Unknown Virtual
OSPF on demand	Whether the interface is configured as a demand circuit. If the interface is a virtual interface, the parameter is always on. For a point-to-point interface, the status of the demand circuit after negotiation with the remote end is displayed within the parenthesis.
Passive	Whether the interface is acting as a passive interface. For more information, see “Passive Interfaces” on page 29-12.

Table 29-7: Parameters in output of the **show ospf interface** command for a specific interface (cont)

Parameter	Meaning
State	State of the interface: Unknown Down Loopback Waiting Ptp DR BackupDR OtherDR
Router priority	Router priority for the interface. A value of zero means the switch is not eligible to become the designated router on the network associated with the interface.
Transit delay	Estimated time in seconds to transmit a link state advertisement over the interface.
Retransmit interval	Seconds between retransmissions of link state advertisement, database description, and link state request packets.
Hello interval	Seconds between hello packets transmitted from the interface. This value must be the same for all switches attached to a common network.
Router dead interval	Seconds advertised by this switch after which it should be considered as down if another switch on the network has not heard from this switch. This value must be the same for all switches attached to a common network.
Poll interval	On non-broadcast multi-access networks, the time in seconds when hello packets are sent to neighbouring switches that are deemed to be inactive. This should be set to a larger value than the hello interval timer.
Interface events	Number of times the interface has changed its OSPF interface state, or the number of times an error has occurred.
Authentication	Whether the authentication type for this interface is password, MD5, or none; and followed by "(area default)" if the interface type is actually determined by the area's default authentication type.
Password	Authentication key if the area is set to authentication by password. This password is a security measure care must be taken with it. The field is empty if no password is defined.
Designated router	Current designated router for a broadcast or non-broadcast multi-access network, or "None" if a designated router has not yet been selected.
Backup designated router	Current backup designated router for a broadcast or non-broadcast multi-access network, or "None" if a backup designated router has not yet been selected.
Virtual neighbour	Identification number for the virtual neighbour when the interface is a virtual interface.
Transit area	Transit area used when the interface is a virtual interface.
Metric boost 1	The boost applied to type 1 metrics of all OSPF LSAs flooded out the interface.

Examples To display summary information for all OSPF interfaces, use the command:

```
sh ospf int
```

To display detailed information for all OSPF interfaces, use the command:

```
sh ospf int ful
```

To display detailed information for OSPF interface ppp0, use the command:

```
sh ospf int=ppp0
```

Related Commands

- [add ospf interface](#)
- [delete ospf interface](#)
- [set ospf interface](#)
- [reset ospf counter](#)

show ospf lsa

Syntax `SHoW OSPF LSA=link-id [AREa={BAckbone|area-number}]`
`[{FULl|SUMmary}] [TYpe={ASBrsummary|ASExternal|ASNSSA|`
`ASSummary|IPsummary|SUMmary|NETwork|ROUTer}]`

where:

- *link-id* is an IP address in dotted decimal notation.
- *area-number* is a 4-byte OSPF area number in dotted decimal notation.

Description This command displays the current link state advertisements within the topological database. This display relates to the MIB entity ospfLsdbTable.

The display is split into two sections, link state advertisements by area, and for if the area is not a stub area, the external link state advertisements.

The LSA parameter specifies the unique ID of the network link, or link state identifier, associated with the LSA to display. Only LSAs whose link state identifiers match the specified link state identifier are displayed. If a link state identifier is not specified, summary information about all LSAs is displayed (Figure 29-12 on page 29-91, Table 29-8 on page 29-91). If a link state identifier is specified, detailed information about the specified LSA is displayed (Figure on page 29-92, Table 29-9 on page 29-93). Wildcard addresses with zeros in the right-hand position may be used to match multiple LSAs. For example, the value 172.16.0.0 matches (and displays) all LSAs whose link state identifier begins with 172.16.

The **area** parameter specifies the area for which LSA information is to be displayed. LSAs associated with the specified area are displayed. Wildcard addresses with zeros in the right-hand position may be used to match multiple areas. For example, the value 172.16.0.0 matches (and displays) LSAs associated with areas whose area number begins with 172.16. The area number 0.0.0.0 is reserved for the backbone area, and can be used interchangeably with **backbone**.

The **full** and **summary** parameters override the default behaviour. The **full** parameter displays detailed information about the specified LSA or all LSAs. The **summary** parameter displays summary information about the specified LSA or all LSAs.

The **type** parameter specifies the type of link state advertisements to display. If **router** is specified, only router link state advertisements are displayed. If **network** is specified, only network link state advertisements are displayed. If **ipsummary** or **summary** is specified, only summary IP advertisements are displayed. If **asbrsummary** or **assummary** is specified, only autonomous system border router summary links are displayed. If **asexternal** is specified, only autonomous system external links are displayed. The **asexternal** parameter is not valid if **area** has been specified. If **asnssa** is specified, only autonomous system external type 7 (NSSA) links are displayed.

Figure 29-12: Example output from the **show ospf lsa summary** command

Type	LS ID	Router ID	Sequence	Age	Len	Csum

Area 0.0.0.0:						
Router	172.28.2.8	172.28.2.8	80000001	401	36	cd27
Summary	172.30.0.0	172.28.2.8	80000001	384	28	fcd2
Area 0.0.0.1:						
Router	172.28.2.8	172.28.2.8	800000a6	381	36	482b
Router	172.28.2.9	172.28.2.9	80000096	*6	36	de4a
Router	172.28.10.4	172.28.10.4	80000093	387	36	aadc
Summary	172.28.0.0	172.28.2.8	80000024	391	28	c6ea
AsNssa	192.32.0.0	172.30.2.4	80000005	392	36	72f9
External:						
AsExternal	0.0.0.0	172.28.2.8	80000001	383	36	b1d4
AsExternal	172.16.0.0	172.28.2.8	80000001	383	36	4741
AsExternal	172.20.0.0	172.28.2.8	80000001	383	36	e644
AsExternal	172.21.0.0	172.28.2.8	80000001	383	36	da4f
AsExternal	172.23.0.0	172.28.2.8	80000001	383	36	ff1
AsExternal	172.24.0.0	172.28.2.8	80000001	383	36	72d9
AsExternal	172.26.0.0	172.28.2.8	80000001	383	36	218f
AsExternal	172.27.0.0	172.28.2.8	80000001	383	36	ba6d

Table 29-8: Parameters in output of the **show ospf lsa summary** command

Parameter	Meaning
Type	Type of link state advertisement.
LS ID	The link state advertisement identification field and its contents are specific to the type of link state advertisement. This field identifies the piece of the routing domain being described by the advertisement.
RouterID	Router identification number of the OSPF router that originated the link state advertisement.
Sequence	Sequence number of the link state advertisement. The sequence number is a 32-bit signed integer. It starts with the value 0x80000001 and increments by one up to 0x7FFFFFFF.
Age	Age of the link state advertisement. The maximum age for any link state advertisement is 3600 seconds. If the age field has a "*", it means the LSA is a DoNotAge LSA.
Len	Length of the LSA in bytes, including the 20 byte LSA header.
Csum	Checksum of the complete link state advertisement, except the age field. This can be used to compare two instances of the same link state advertisement.

Figure 29-13: Example output from the **show ospf lsa full** command

Type	LS ID	Router ID	Sequence	Age	Len	Csum

Area backbone:						
Router	1.1.1.1	1.1.1.1	8000008e	34	36	9c42
Options: --B Number of links: 1						
Link 1: Type: Transit ID: 192.168.3.4 Data: 192.168.3.1						
TOS 0 metric: 1 Number of other metrics: 0						
Router	2.2.2.2	2.2.2.2	80000114	28	36	50fe
Options: --B Number of links: 1						
Link 1: Type: Transit ID: 192.168.3.4 Data: 192.168.3.2						
TOS 0 metric: 1 Number of other metrics: 0						
Router	4.4.4.4	4.4.4.4	800000f7	26	36	0f4b
Options: --B Number of links: 1						
Link 1: Type: Transit ID: 192.168.3.4 Data: 192.168.3.4						
TOS 0 metric: 1 Number of other metrics: 0						
Network	192.168.3.4	4.4.4.4	80000002	26	36	5365
Network Mask: 255.255.255.0						
Attached router: 4.4.4.4						
Attached router: 1.1.1.1						
Attached router: 2.2.2.2						
Summary	192.168.1.0	1.1.1.1	80000004	1208	28	9456
Network Mask: 255.255.255.0						
TOS: 0 Metric: 1						
Summary	192.168.2.0	2.2.2.2	80000005	147	28	697b
Network Mask: 255.255.255.0						
TOS: 0 Metric: 1						
Summary	192.168.4.0	4.4.4.4	80000004	1580	28	19c2
Network Mask: 255.255.255.0						
TOS: 0 Metric: 1						
Area 3.3.3.3:						
Router	4.4.4.4	4.4.4.4	8000000c	1342	36	fc47
Options: --B Number of links: 1						
Link 1: Type: Transit ID: 192.168.4.4 Data: 192.168.4.4						
TOS 0 metric: 1 Number of other metrics: 0						
Router	5.5.5.5	5.5.5.5	80000013	83	36	ad87
Options: --- Number of links: 1						
Link 1: Type: Transit ID: 192.168.4.4 Data: 192.168.4.5						
TOS 0 metric: 1 Number of other metrics: 0						
Network	192.168.4.4	4.4.4.4	80000004	1338	32	8031
Network Mask: 255.255.255.0						
Attached router: 4.4.4.4						
Attached router: 5.5.5.5						
Summary	192.168.1.0	4.4.4.4	80000005	41	28	429a
Network Mask: 255.255.255.0						
TOS: 0 Metric: 2						
Summary	192.168.2.0	4.4.4.4	80000001	14	28	3fa0
Network Mask: 255.255.255.0						
TOS: 0 Metric: 2						
Summary	192.168.3.0	4.4.4.4	80000004	1560	28	24b8
Network Mask: 255.255.255.0						
TOS: 0 Metric: 1						
AsNssa	192.32.0.0	172.30.2.4	80000005	392	36	72f9
Network Mask: 255.255.255.0 Options: --N---						
TOS: 0 Metric: E1- 2 Foward: 192.30.5.2 Tag: 00000000						

Figure 29-13: Example output from the **show ospf lsa full** command (cont)

```

External:
  AsExternal  202.49.72.0      5.5.5.5      80000004      83   36   12e2
    Network Mask: 255.255.255.0  Options: ----- NSSA Translated: N
      TOS: 0  Metric: E1-      2  Foward: 0.0.0.0      Tag: 00000000

  AsNssa      192.30.0.0      172.30.2.4    80000004      392  36   1f52
    Network Mask: 255.255.255.0  Options: ----- NSSA Translated: Y
      TOS: 0      Metric:  E1-      2  Foward: 192.30.5.2  Tag: 00000000
-----

```

Table 29-9: Parameters in output of the **show ospf lsa full** command

Parameter	Meaning
Type	Type of link state advertisement: Router Network Summary AsSummary AsExternal AsNssa
LS ID	Link state advertisement identification field and its contents are specific to the type of link state advertisement. This field identifies the piece of the routing domain being described by the advertisement.
Router ID	Router identification number of the OSPF router that originated the link state advertisement.
Sequence	Sequence number of the link state advertisement. The sequence number is a 32-bit signed integer. It starts with the value 0x80000001 and increments by one up to 0x7FFFFFFF
Age	Age of the link state advertisement. The maximum age for any link state advertisement is 3600 seconds.
Len	Length of the LSA in bytes, including the 20 byte LSA header.
Csum	Checksum of the complete link state advertisement, except the age field. This is useful for comparing two instances of the same link state advertisement.
Router	Information about a router link state advertisement.
Options	The option bits set in the LSA packet header; zero or more of "B" (switch is an area border router), "E" (switch is an autonomous system boundary router), "V" (switch is the end point of a virtual link), "N" (AsNssa LSA can be translated) or "-" (bit not set).
Number of links	Number of links in the LSA packet.
Link	Link number of a link in the LSA.
Type	Type of link for the associated link in the LSA.
ID	Router ID of the OSPF router at the remote end of the link.
Data	Data pertaining to the link type.
TOS 0 metric	Metric for TOS (Type Of Service) 0 for the link.
Number of other metrics	Number of other metrics defined for the link.

Table 29-9: Parameters in output of the **show ospf lsa full** command (cont)

Parameter	Meaning
TOS	TOS (Type Of Service) entry in the LSA.
Metric	Metric for the accompanying TOS (Type Of Service).
Network	Information about a network link state advertisement.
Network mask	IP network mask for the advertised network.
Attached router	Router ID of the attached OSPF router for the network LSA.
Summary	Information about an autonomous system border router summary link state advertisement.
AsNssa	Information about an autonomous system external link state advertisement (type 7) within an NSSA.
AsExternal	Information about an autonomous system external link state advertisement.
NSSA Translated	Whether the AsExternal link state advertisement has been translated from an AsNssa (type 7) link state advertisement.
Forward	Forwarding IP address for the accompanying TOS entry.
Tag	Field for communication between AS boundary routers. Not used by OSPF.

Examples To display all LSAs for the backbone area, use the command:

```
sh ospf lsa are=ba
```

show ospf md5key

Syntax `SHoW OSPF MD5key [INTErface=interface]`

Description This command displays information about the OSPF MD5 keys for all interfaces, or for a specific interface (Figure 29-14, Table 29-10). Each interface has its own set of keys.

The **interface** parameter specifies the OSPF interface for the MD5 keys to be displayed. If an interface is specified, only the MD5 keys for the specified interface are displayed. If an interface is not specified, the MD5 keys for all OSPF interfaces are displayed.

Figure 29-14: Example output from the **show ospf md5key** command

OSPF MD5 keys			
Interface	ID	Key	Active
vlan1	1	O3jf87Pls	No
	2	xm39s2F28	Yes
vlan2	3	ba2958d2x	Yes

Table 29-10: Parameters in output of the **show ospf md5key** command

Parameter	Meaning
Interface	OSPF interface to which the keys belong.
ID	The key's identification number.
Key	The MD5 key.
Active	Whether or not the key is currently being used to authenticate packets being received from one or more neighbours.

Example To show the MD5 keys for the vlan1 interface, use the command:

```
show ospf md5key interface=vlan1
```

Related Commands [add ospf md5key](#)
[delete ospf md5key](#)

show ospf neighbour

Syntax SHow OSPF NEIghbour [=ipadd] [INTerface=interface]

where:

- *ipadd* is an IP address in dotted decimal notation.
- *interface* is a valid interface name.

Description This command displays information about OSPF neighbours (Figure 29-15 on page 29-96, Table 29-11 on page 29-96). If a neighbour is specified, only information about the specified neighbour is displayed. This display relates to the MIB entity ospfNbrTable.

The **interface** parameter specifies a valid interface already assigned and configured. Valid interfaces are:

- PPP (such as ppp0)
- VLAN (such as vlan1)
- FR (such as fr0)
- virtual interface (such as virt9)

The **neighbour** parameter specifies the IP address for which neighbour information is to be displayed. Neighbours with the specified address are displayed. Wildcard addresses with zeros in the right-hand position may be used to match multiple neighbours. For example, the value 172.16.0.0 matches (and displays) all OSPF neighbours with addresses beginning 172.16.

The **interface** parameter specifies the interface for which neighbour information is to be displayed. Neighbours associated with the specified interface are displayed.

Figure 29-15: Example output from the **show ospf neighbour** command

IP address	State	Interface	Router ID	Priority	LSRxmtQ	Type
192.168.67.1	full	vlan1	10.34.5.78	10	23	Dyn
192.168.200.12	full	fr0	0.0.0.5	5	0	Stat
192.168.234.123	full	fr0	124.123.234.143	255	1034	Stat

Table 29-11: Parameters in output of the **show ospf neighbour** command

Parameter	Meaning
IP address	IP address of the neighbour.
State	Neighbour state.
Interface	Interface associated with the neighbour.
Router ID	Neighbour's assigned OSPF router identification number.
Priority	Priority currently being advertised by the neighbour in its hello packets.
LSRetransQ	Current length of the link state retransmission queue.
Type	Whether the entry is dynamically learned or statically defined.

Related Commands [add ospf neighbour](#)
 [delete ospf neighbour](#)
 [reset ospf counter](#)
 [set ospf neighbour](#)

show ospf range

Syntax SHow OSPF RANge [=ipadd] [AREa={BAckbone|area-number}]

where:

- *ipadd* specifies an IP address in dotted decimal notation.
- *area-number* is a 4-byte OSPF area number in dotted decimal notation.

Description This command displays information about the OSPF ranges that have been defined (Figure 29-16 on page 29-98, Table 29-12 on page 29-98). This display relates to the MIB entity ospfAreaRangeTable. The range table shows which ranges have been defined in which areas.

The **range** parameter specifies the OSPF range or ranges to display. Ranges whose base IP address matches the specified range are displayed. Wildcard addresses with zeros in the right-hand position may be used to match multiple ranges. For example, the value 172.16.0.0 matches (and displays) all OSPF ranges whose base address begins with 172.16.

The **area** parameter specifies the area for which range information is to be displayed. Ranges associated with the specified area are displayed. Wildcard addresses with zeros in the right-hand position may be used to match multiple areas. For example, the value 172.16.0.0 matches (and displays) OSPF ranges associated with areas whose area number begins with 172.16. The area number 0.0.0.0 is reserved for the backbone area, and can be used interchangeably with **backbone**.

Figure 29-16: Example output from the **show ospf range** command

Base IP address	State	Mask	Area	Effect
10.0.0.0	Active	255.0.0.0	0.0.0.1	Advertise
192.168.10.0	Inactive	255.255.255.0	10.34.143.234	Advertise
192.168.123.240	Active	255.255.255.240	123.234.243.125	Advertise

Table 29-12: Parameters in output of the **show ospf range** command

Parameter	Meaning
Base IP address	Base IP address of the network or subnet for the range.
State	Whether the range is active or inactive. A range is active if it is attached to an active area
Mask	Subnet mask for the network or subnet.
Area	Area to which the range is attached.
Effect	Whether the range is advertised to other areas via a summary LSA.

Examples To display OSPF range information for all ranges whose base address begins with 172.16, use the command:

```
sh ospf ran=172.16.0.0
```

To display OSPF range information for all areas numbered 192.168.10.0 to 192.168.10.255, use the command:

```
sh ospf ran are=192.168.10.0
```

Related Commands [add ospf range](#)
[delete ospf range](#)
[set ospf area](#)

show ospf redistribute

Syntax SHow OSPF REDistribute

Description This command displays the redistribution definitions that OSPF uses when it imports routes from other protocols into the set of routes that it advertises (Figure 29-17, Table 29-13).

Figure 29-17: Example output from the **show ospf redistribute** command

OSPF Redistribute						
Protocol	Metric	RouteMap	Subnet	Tag	Type	Limit/Redistributed
Static	20	-	YES	10	Ext2	500/ 201
BGP	20	-	NO	20	Ext2	2000/ 1600
Interface	Original	rmi	NO	Original	Original	1000/ 10

Table 29-13: Parameters in the output of the **show ospf redistribute** command

Parameter	Meaning
Protocol	The routing source from which OSPF imports the routes for this redistribution definition; one of "BGP", "Interface", "RIP", or "Static".
Metric	The route metric that OSPF assigns to routes that it redistributes from this protocol, or "Original" if the original route metric is preserved.
RouteMap	The name of the route map that OSPF uses to filter routes from this protocol before redistributing them, or "-" if a routemap has not been set.
Subnet	Whether OSPF redistributes subnets; one of "NO" (OSPF only redistributes classfull network routes) or "YES" (OSPF redistributes classless and classfull network routes).
Tag	The numeric tag that OSPF uses to label routes that it imports from this protocol, or "Original" if the original tag is preserved.
Type	The OSPF external route type which OSPF assigns to routes that it redistributes from this protocol; one of "Ext1" (External Type 1), "Ext2" (External Type 2), or "Original" (original route type is preserved). See "Path types" on page 29-9 for more information.
Limit	The maximum number of routes that OSPF will import and redistribute from this protocol.
Redistributed	The number of routes that OSPF has imported and redistributed from this protocol.

Example To see how OSPF will modify statically-configured routes when it imports them, use the command:

```
sh ospf red
```

Related Commands [add ospf redistribute](#)
[delete ospf redistribute](#)
[set ospf redistribute](#)
[show ip routemap](#) in Chapter 31, Filtering IP Routes

show ospf route

Syntax `SHoW OSPF ROUte[=ipadd] [AREa={BACkbone|area-number}]`
`[TYpe={AB|ASbr}]`

where:

- *ipadd* specifies an IP address in dotted decimal notation.
- *area-number* is a 4-byte OSPF area number in dotted decimal notation.

Description This command displays the current state of all OSPF internal routes within the OSPF module (Figure 29-18 on page 29-101, Table 29-14 on page 29-102). The display has two sections: the Area Border Router (AB) routes within particular areas, and the Autonomous System Boundary Router (ASBR) routes, if the area is not a stub area.

The **route** parameter specifies the unique ID of the route to be displayed. If a route is not specified, all OSPF routes matching the other criteria are displayed. Wildcard addresses with zeros in the right-hand position may be used to match multiple routes. For example, the value 172.16.0.0 matches (and displays) OSPF routes whose address begins with 172.16.

The **area** parameter specifies the area for which route information is to be displayed. Routes associated with the specified area are displayed. Wildcard addresses with zeros in the right-hand position may be used to match multiple areas. For example, the value 172.16.0.0 matches (and displays) OSPF routes associated with areas whose area number begins with 172.16. The area number 0.0.0.0 is reserved for the backbone area, and can be used interchangeably with **backbone**.

The **type** parameter specifies the type of route to display. If **AB** is specified, Area Border Router routes are displayed. If **asbr** is specified, Autonomous System Boundary Router routes are displayed.

Figure 29-18: Example output from the **show ospf route** command

OSPF Routes					
Destination DLCI/Circ.	Mask Type	Policy	NextHop Protocol	Interface Metrics	Age Preference

Area backbone AB routes:					
1.1.1.1	255.255.255.255		192.168.3.1	fr0	0
20	ospfAB	0	ospf	1	10
2.2.2.2	255.255.255.255		192.168.3.2	fr0	0
22	ospfAB	0	ospf	1	10
ASBR routes:					
5.5.5.5	255.255.255.255		192.168.4.5	vlan1	0
-	ospfAS	0	ospf	1	10

Table 29-14: Parameters in output of the **show ospf route** command

Parameter	Meaning
Destination	Router ID of the BA router or ASBR router to which this route points.
DLCI/Circ.	Frame Relay or X25T circuit that the route uses
Mask	Mask for the route. Should always be 255.255.255.255.
Type	Whether the route is ospfAB (AB route) or ospfAS (ASBR route).
Policy	Forwarding policy of the route, which is always 0.
NextHop	Forwarding address of the next hop for this route, or 0.0.0.0 if the destination is directly attached.
Protocol	Protocol of the route, which is always OSPF.
Interface	Interface with which the route is associated.
Metrics	Metrics (costs) associated with the route.
Age	Age of the route in seconds.
Preference	Forwarding preference of the route. This is related to the route's OSPF path type. 10 for Intra Area routes and 11 for Inter area routes.

Examples To display information about all OSPF routes, use the command:

```
sh ospf rou
```

To display OSPF routes whose base address begins with 172.16, use the command:

```
sh ospf rou=172.16.0.0
```

To display OSPF range information for all areas numbered 192.168.10.0 to 192.168.10.255, use the command:

```
sh ospf rou are=192.168.10.0
```

Related Commands [show ospf area](#)
[show ospf interface](#)
[show ospf range](#)

show ospf stub

Syntax `SHoW OSPF STUB [=ipadd] [AREa={BAckbone | area-number}]`

where:

- *ipadd* is an IP address in dotted decimal notation.
- *area-number* is a 4-byte OSPF area number in dotted decimal notation.

Description This command displays information about statically configured non-OSPF stub network routes (Figure 29-19, Table 29-15).

The **stub** parameter specifies the stub network or networks to display. Only stub networks whose base IP address matches the specified IP address are displayed. Wildcard addresses with zeros in the right-hand position may be used to match multiple stub networks.

The **area** parameter specifies the area for which stub network information is to be displayed. Stub networks associated with the specified area are displayed. Wildcard addresses with zeros in the right-hand position may be used to match multiple areas. For example, the value 172.16.0.0 matches (and displays) OSPF stub networks associated with areas whose area number begins with 172.16. The area number 0.0.0.0 is reserved for the backbone area, and can be used interchangeably with **backbone**.

Figure 29-19: Example output from the **show ospf stub** command

IP address	Mask	State	Area	Metric	TOS	Type
192.168.150.2	255.255.255.255	Active	Backbone	5	0	Stat
192.168.8.0	255.255.255.0	Active	Backbone	4	0	Stat

Table 29-15: Parameters in output of the **show ospf stub** command

Parameter	Meaning
IP address	IP address of the host or point-to-point network.
Mask	Mask for the stub network.
State	Whether the stub network is active or inactive. When a stub network is active, it is advertised via the switch LSA.
Area	Area number of the area containing the stub network.
Metric	Metric to be advertised for the stub network.
TOS	Type of service of the route to the stub network.
Type	Whether the stub network entry is permanent static or dynamic.

Examples To display all OSPF stub networks, use the command:

```
sh ospf stub
```

To display information for all stub networks from 172.30.0.0 to 172.30.255.255 in area 0.0.0.3, use the command:

```
sh ospf stub=172.30.0.0 are=0.0.0.3
```

Related Commands [add ospf stub](#)
 [delete ospf stub](#)
 [set ospf stub](#)

show ospf summaryaddress

Syntax `SHow OSPF SUMMaryaddress`

Description This command displays the list of summary addresses configured for OSPF (Figure 29-20, Table 29-16).

Figure 29-20: Example output from the **show ospf summaryaddress** command

OSPF summary addresses			
Base IP address	Mask	Advertise	Tag
192.168.1.0	255.255.255.0	Yes	13
10.3.0.0	255.255.0.0	No	0

Table 29-16: Parameters in the output of the **show ospf summaryaddress** command

Parameter	Meaning
Base IP address	The base IP address of this summary address entry. In combination with the Mask field, it defines the routes that are summarized, and the route that is generated as the summary.
Mask	Subnet mask for the network or subnet.
Advertise	Whether the summary address is advertised.
Tag	The tag inserted in the AS external LSA for this summary address. It overrides tag information in the original route or set by the route map used to select the original route for redistribution.

Examples To display the list of currently configured summary addresses, use the command:

```
sh ospf summ
```

Related Commands

- [add ospf summaryaddress](#)
- [delete ospf summaryaddress](#)
- [set ospf summaryaddress](#)

