Chapter 13
OSPF Commands

area

Assigns an area ID value to which the routing switch will be attached. Area membership is assigned using the Interface Level command, **ip ospf area** <area-id>.

- When an area is defined as **normal**, all external routes will be advertised into the area.
- When an area is defined as **stub**, external routes will not be advertised into the area.
- When an area is defined as **nssa**, OSPF does not flood external routes from other areas into the area, but does translate and flood route information from the area into other areas, such as the backbone.

NOTE: You can assign one area per routing switch port. If the routing switch has 64 ports, 64 areas are supported on that routing switch.

By default, the OSPF feature is disabled. OSPF must be enabled and active on the routing switch for this command to be operational.

EXAMPLE:

To define a normal area, enter the following commands:

```
HP9300(config)# router ospf
HP9300(config-ospf-router)# area 192.53.0.0
```

EXAMPLE:

To define an area as a stub area, enter the following commands:

```
HP9300(config)# router ospf
HP9300(config-ospf-router)# area 192.53.0.0 stub 1
```

EXAMPLE:

To define an area as an NSSA, enter the following commands:

```
HP9300(config)# router ospf
HP9300(config-ospf-router)# area 192.53.0.0 nssa 1
```

EXAMPLE:

To disable summary LSAs for a stub area, enter commands such as the following:

```
HP9300(config-ospf-router)# area 40 stub 1 no-summary
```

Syntax: area <num> | <ip-addr> [nssa <cost> | stub <cost> [no-summary]]

The <num> | <ip-addr> parameter specifies the area number, which can be a number or in IP address format. If you specify an number, the number can be from 0 – 2,147,483,647.

The **nssa** parameter specifies that this is a Not-So-Stubby Area (NSSA).

The <cost> specifies an additional cost for using a route to or from this area and can be from 1 – 16777215. If you configure a stub area or NSSA, you must specify the cost. There is no default. Normal areas do not use the cost parameter.

The **no-summary** parameter applies only to stub areas and disables summary LSAs from being sent into the area.

NOTE: You can assign one area on a routing switch interface. For example, if the system or chassis module has 16 ports, 16 areas are supported on the chassis or module.

Possible values: Area ID: Valid IP address; Stub cost: 1 – 16777215

Default value: If you do not specify **stub** or **nssa**, a normal area is defined.

area <num> | <ip-addr> virtual-link

Provides an area-border routing switch a logical connection to the backbone area (0.0.0.0) when a physical connection to the backbone area does not exist.

The **area** <num> **|** <ip-addr> represents the shared area of the two area border routers—the one with a physical connection to the backbone and the router that requires a logical connection to the backbone. The defined area serves as the connection point between the two routers.

The **virtual-link** <ip-addr> is the Router ID of the router physically connected to the backbone, when assigned from the router interface requiring a logical connection. When assigning the parameters from the router with the physical connection, the router ID is the IP address of the router requiring a logical connection to the backbone.

NOTE: When establishing an area virtual link, it must be configured on both of the routers (both ends of the virtual link).

By default, the OSPF feature is disabled. OSPF must be enabled and active on the routing switch for this command to be operational. In addition, the neighbor area border router that has a physical connection to the backbone and the common area, must have connections to both the neighbor area border router and the area border router requiring a logical connection.

EXAMPLE:

To define the virtual link for area 195.22.0.0, enter the following:

```
HP9300(config)# router ospf
HP9300(config-ospf-router)# area 195.22.0.0 virt 201.44.53.44
```

Other parameters that can be modified with this command, as seen in the syntax, are summarized below:

authentication-key: A password used to validate action

dead-interval: The number of seconds that a neighbor router will wait for a hello packet from the current router, before declaring the router down

hello-interval: The length of time between the transmission of hello packets.

md5-authentication: The MD5 key-activation wait time, key ID, and key string.

retransmit-interval: The time between retransmits of link state advertisements to router adjacencies for this interface.

transmit-delay: The time it takes to transmit Link State Update packets on this interface

Syntax: area <num> | <ip-addr> virtual-link <ip-addr> [authentication-key <string>] [dead-interval <num>] [hello-interval <num>] [md5-authentication key-activation-wait-time <num> | key-id <num> key <string>] [retransmit-interval <num>] [transmit-delay <num>]

Possible values and Default values:

Parameter	Possible Values	Default
authentication-key:	up to 8 alphanumeric characters	none
dead-interval:	1 – 65535 seconds	40 seconds
hello-interval:	1 – 65535	10 seconds
md5-authentication key-activation-wait-time:	0 – 14400	300 seconds (5 minutes)
md5-authentication key ID:	1 – 255	none
md5-authentication key string:	up to 16 alphanumeric characters	none
retransmit-interval:	0 – 3600 seconds.	5 seconds
transmit-delay:	0 – 3600	1

area range

Assigns representative values to a range of IP addresses within an area, so that only those reference range addresses are advertised to the network, instead of all the addresses within that range. Up to four ranges can be assigned to an area.

NOTE: Range assignment is optional.

EXAMPLE:

```
HP9300(config)# router ospf
HP9300(config-ospf-router)# area 192.53.0.0 range 193.45.0.0 255.255.0.0
```

Syntax: area <ip-addr> range <ip-addr> <ip-mask>

Possible values: area id (ip address), ip address and ip mask

Default value: N/A

database-overflow-interval

Configures how often a routing switch will check to see if the OSPF external link state database overflow condition has been eliminated by removal of entries originated on the routing switch.

If the configured value of the data-base-overflow-interval is zero, then the routing switch will never leave the database overflow condition. The default value for the database overflow interval is zero.

EXAMPLE:

```
HP9300(config-ospf-router)# data-base-overflow-interval 60
```

Syntax: database-overflow-interval <value>

Possible values: 0 – 86,400 seconds

Default value: 0

default-information-originate

Disables or re-enables origination of default routes.

With this feature enabled, the routing switch advertises a type 5 default route that is flooded throughout the AS (except stub areas and NSSAs). In addition, internal NSSA ASBRs advertise their default routes as translatable type 7 default routes.

If you disable default information origination, the default route originated by the routing switch is flushed. Default routes generated by other OSPF routers are not affected. If you re-enable the feature, the feature takes effect immediately and thus does not require you to reload the software.

The ASBR advertises its default route, if one is configured on the ASBR. You can use the “always” option when you enable the default route origination. The always option causes the ASBR to create and advertise a default route if it does not already have one configured.

If you want to disable default information originate, use the following CLI method.

EXAMPLE:

To disable default information originate, enter the following command:

```
HP9300(config-ospf-router)# no default-information-originate
```

To re-enable the feature or change the feature’s parameters, enter a command such as the following:

```
HP9300(config-ospf-router)# default-information-originate always metric 2
metric-type 1
```

Syntax: [no] default-information-originate [always] [metric <value>] [metric-type <type>]

The **always** parameter advertises the default route regardless of whether the routing switch has a default route. This option is disabled by default.

The **metric** <value> parameter specifies a metric for the default route. If this option is not used, the default metric is used for the route.

The **metric-type** <type> parameter specifies the external link type associated with the default route advertised into the OSPF routing domain. The <type> can be one of the following:

- 1 – Type 1 external route
- 2 – Type 2 external route

If you do not use this option, the default redistribution metric type is used for the route type.

NOTE: If you specify a metric and metric type, the values you specify are used only if the routing switch does not have a default route, but still wants to advertises one because the always option is configured.

Possible values: N/A

Default value: Enabled

default-metric

Sets the global default-metric value that will be adopted by all external routes imported into OSPF.

EXAMPLE:

To set a default metric of 1000 to be applied to all external routes imported into OSPF, enter the following command.

```
HP9300(config-ospf-router)# def 1000
```

Syntax: default-metric <value>

Possible values: 1 – 16,777,215

Default value: 10

deny redistribute

Defines the route(s) upon which you do not want to perform OSPF redistribution.

EXAMPLE:

To deny redistribution on incoming routes received from the 192.95.0.0 network, enter the following:

```
HP9300(config-ospf-router)# deny redis all 2 192.95.0.0 255.255.0.0
```

Syntax: deny redistribute <filter-num> all | bgp | rip | static
[address <ip-addr> <ip-mask> [match-metric <value> [set-metric <value>]]]

Possible values: see below

all	apply redistribution to all route types
bgp	apply redistribution to BGP4 routes only
rip	apply redistribution to RIP routes only
static	apply redistribution to the static route only
ip address	network and sub-net addresses
match-metric	applies redistribution only to those incoming routes that match a specific metric value; Possible values: 1 – 15
set-metric	OSPF metric value that will be applied to all routes imported into OSPF

Default value: N/A

NOTE: If a **set-metric** value is not set using the set-metric parameter, then the value configured for the global parameter default-metric will be applied.

distance

Configures an administrative distance for a specific OSPF route type. For example, you can use this feature to prefer a static route over an OSPF inter-area route but you also want to prefer OSPF intra-area routes to static routes.

The distance you specify influences the choice of routes when the routing switch has multiple routes for the same network from different protocols. The routing switch prefers the route with the lower administrative distance.

You can specify unique default administrative distances for the following route types:

- Intra-area routes
- Inter-area routes
- External routes

The default for all these OSPF route types is 110.

NOTE: This feature does not influence the choice of routes within OSPF. For example, an OSPF intra-area route is always preferred over an OSPF inter-area route, even if the intra-area route's distance is greater than the inter-area route's distance.

To configure administrative distances for OSPF route types, use the following CLI method.

EXAMPLE:

To change the default administrative distances for inter-area routes, intra-area routes, and external routes, enter the following command:

```
HP9300(config-ospf-router)# distance external 100
HP9300(config-ospf-router)# distance inter-area 90
HP9300(config-ospf-router)# distance intra-area 80
```

Syntax: distance external | inter-area | intra-area <distance>

The **external | inter-area | intra-area** parameter specifies the route type for which you are changing the default administrative distance.

The <distance> parameter specifies the new distance for the specified route type. Unless you change the distance for one of the route types using commands such as those shown above, the default is 110.

To reset the administrative distance to its system default (110), enter a command such as the following:

```
HP9300(config-ospf-router)# no distance external 100
```

Possible values: see above

Default value: 110

end

Moves activity to the privileged EXEC level from any level of the CLI except the user EXEC level.

EXAMPLE:

To move to the privileged level, enter the following from any level of the CLI.

```
HP9300(config-ospf-router)# end
HP9300#
```

Syntax: end

Possible values: N/A

Default value: N/A

exit

Moves activity up one level from the current level. In this case, activity will be moved to the global level.

EXAMPLE:

```
HP9300(config-ospf-router)# exit
HP9300(config)#
```

Syntax: exit

Possible values: N/A

Default value: N/A

external-lsdb-limit

Provides compliance with RFC 1765 in the handling of OSPF external link-state database (LSDB) overflow. HP routing switches, by default, support 2000 IP OSPF external link state advertisements before a database overflow condition exists.

EXAMPLE:

To decrease this value to 1500 seconds, enter the following:

```
HP9300(config-ospf-router)# external-lsdb-limit 1500
```

Syntax: database-external-lsdb-interval <value>

Possible values: 0 to 2000

Default value: 2000

maximum-number-of-lsa

Adjusts the table sizes for Link-State Advertisements (LSAs) on routing switches with 32MB or more memory.

NOTE: This feature applies only to devices with 32MB memory or greater.

The following table lists the types of LSAs for which you can configure the table size, the default number of entries the tables can hold, and the range of maximum values you can specify. You cannot configure the LSA tables globally; you must configure them for individual LSA types.

LSA Type	Default Maximum Number of Entries	Range of Values
External	2000	500 – 8000
Network	2000	200 – 2000
Router	2200	200 – 2200
Summary	2000	500 – 18000

EXAMPLE:

To change the maximum number of summary LSA entries from 2000 to 18000, enter the following commands:

```
HP9300(config-ospf-router)# maximum-number-of-lsa summary 18000
HP9300(config-ospf-router)# write memory
HP9300(config-ospf-router)# end
HP9300# reload
```

Syntax: maximum-number-of-lsa external | network | router | summary <value>

You must save the change to the startup-config file and reload or reboot. The change does not take effect until your reload or reboot.

NOTE: When you allocate high values for the LSA types, you reduce the amount of memory available to other protocols. If you need to run more than one memory-intensive protocol simultaneously, make sure you consider the memory needs of each. For example, if you are planning to run OSPF and BGP4 on the same routing switch, you might not want to allocate the maximum memory capacity for all the LSA types. Doing so could take memory away from BGP4 and thus limit configurable BGP4 table sizes such as the maximum number of routes the BGP route table can hold.

Possible values: see above

Default value: see above

max-routes

Changes the maximum number of OSPF routes the routing switch can hold.

EXAMPLE:

To change the maximum number of OSPF routes to 32000, enter the following command:

```
HP9300(config-ospf-router)# max-routes 32000
HP9300(config-ospf-router)# end
HP9300# reload
```

Syntax: max-routes <num>

Possible values: 4000 – 32000

Default value: 16000

metric-type

Specifies the type of OSPF metric to be used for routes imported into OSPF. Type 2 specifies a big metric (3 bytes). Type 1 specifies a small metric (2 bytes).

EXAMPLE:

```
HP9300(config-ospf-router)# metric-type type1
```

Syntax: metric-type type1 | type2

Possible values: type1, type2

Default value: type2

no

Disables other commands. To disable a command, place the word **no** before the command.

permit redistribute

Defines the route types upon which you want to perform OSPF redistribution.

OSPF must be enabled and active for this command to be operational. OSPF is disabled by default.

EXAMPLE:

```
HP9300(config-ospf-router)# permit redis rip 1 address 192.147.72.0 255.255.255.0
```

Syntax: permit redistribute <filter-num> all | bgp | connected | rip | static
[address <ip-addr> <ip-mask> [match-metric <value> [set-metric <value>]]]

Possible values:

all	apply redistribution to all route types
bgp	apply redistribution to BGP4 routes only
connected	apply redistribution to directly connected routes only
rip	apply redistribution to RIP routes only
static	apply redistribution to the static route only
ip address	network and sub-net addresses
match-metric	match a specific metric value; Possible values: are 1 – 16777215
set-metric	OSPF metric value that will be applied to all routes imported into OSPF

Default value: N/A

quit

Returns you from any level of the CLI to the User EXEC mode.

EXAMPLE:

```
HP9300(config-ospf-router)# quit
HP9300>
```

Syntax: quit

Possible values: N/A

Default value: N/A

redistribution

Enables the OSPF route redistribution function on the routing switch. When enabled, OSPF will import external routes into the OSPF domain. The routing switch will behave as an Autonomous System Boundary Router (ASBR). You can select the types of routes you want to redistribute for each protocol.

EXAMPLE:

To enable redistribution of RIP routes into OSPF:

```
HP9300(config-ospf-router)# redistribution rip
```

Syntax: [no] redistribution bgp | connected | rip | static

Possible values: bgp allows the routing switch to redistribute BGP4 routes into OSPF.

connected allows the routing switch to redistribute routes to directly attached devices into OSPF

rip allows the routing switch to redistribute RIP routes into OSPF.

static allows the routing switch to redistribute static IP routes into OSPF.

Default value: disabled

rfc1583-compatibility

HP routing switches are configured by default to be compliant with RFC 1583 OSPF V2 specification. Routing switches can also be configured to operate with the latest OSPF standard, RFC 2178, by entering the **no rfc1583-compatibility** command.

EXAMPLE:

```
HP9300(config-ospf-router)# rfc1583-compatibility
```

Syntax: [no] rfc1583-compatibility

Possible values: N/A

Default value: disabled

show

Displays a variety of configuration and statistical information about the switch or routing switch. See “Show Commands” on page 20-1.

summary-address

Configures external route summarization.

When the routing switch is an OSPF Autonomous System Boundary Router (ASBR), you can configure it to advertise one external route as an aggregate for all redistributed routes that are covered by a specified address range.

When you configure an address range, the range takes effect immediately. All the imported routes are summarized according to the configured address range. Imported routes that have already been advertised and that fall within the range are flushed out of the AS and a single route corresponding to the range is advertised.

If a route that falls within a configured address range is imported by the routing switch, no action is taken if the routing switch has already advertised the aggregate route; otherwise the routing switch advertises the aggregate route. If an imported route that falls with in a configured address range is removed by the routing switch, no action is taken if there are other imported route(s) that fall with in the same address range; otherwise the aggregate route is flushed.

You can configure up to 32 address ranges. The routing switch sets the forwarding address of the aggregate route to zero and sets the tag to zero.

If you delete an address range, the advertised aggregate route is flushed and all imported routes that fall within the range are advertised individually.

If an external LSDB overflow condition occurs, all aggregate routes are flushed out of the AS, along with other external routes. When the routing switch exits the external LSDB overflow condition, all the imported routes are summarized according to the configured address ranges.

NOTE: If you use redistribution filters in addition to address ranges, the routing switch applies the redistribution filters to routes first, then applies them to the address ranges. If you disable redistribution, all the aggregate routes are flushed, along with other imported routes.

NOTE: This option affects only imported, type 5 external routes. A single type 5 LSA is generated and flooded throughout the AS for multiple external routes. Type 7-route redistribution is not affected by this feature. All type 7 routes will be imported (if redistribution is enabled). To summarize type 7 LSAs or exported routes, use NSSA address range summarization. See the "Configuring OSPF" chapter of the *Advanced Configuration and Management Guide*.

EXAMPLE:

To configure a summary address for OSPF routes, enter commands such as the following:

```
HP9300(config-ospf-router)# summary-address 10.1.0.0 255.255.0.0
```

The command in this example configures summary address 10.1.0.0, which includes addresses 10.1.1.0, 10.1.2.0, 10.1.3.0, and so on. For all of these networks, only the address 10.1.0.0 is advertised in external LSAs.

Syntax: summary-address <ip-addr> <ip-mask>

The <ip-addr> parameter specifies the network address.

The <ip-mask> parameter specifies the network mask.

To display the configured summary addresses, enter the following command at any level of the CLI:

```
HP9300(config-ospf-router)# show ip ospf config

OSPF Redistribution Address Ranges currently defined:

Range-Address    Subnetmask
1.0.0.0          255.0.0.0
1.0.1.0          255.255.255.0
1.0.2.0          255.255.255.0
```

Syntax: show ip ospf config

Possible values: see above

Default value: no summarization

timers

Changes the Shortest Path First (SPF) timers.

The routing switch uses the following timers when calculating the shortest path for OSPF routes:

- SPF delay - When the routing switch receives a topology change, the software waits before it starts a Shortest Path First (SPF) calculation. By default, the software waits five seconds. You can configure the SPF delay to a value from 0 – 65535 seconds. If you set the SPF delay to 0 seconds, the software immediately begins the SPF calculation after receiving a topology change.
- SPF hold time - The routing switch waits for a specific amount of time between consecutive SPF calculations. By default, the routing switch waits ten seconds. You can configure the SPF hold time to a value from 0 – 65535 seconds. If you set the SPF hold time to 0 seconds, the software does not wait between consecutive SPF calculations.

You can set the delay and hold time to lower values to cause the routing switch to change to alternate paths more quickly in the event of a route failure. Note that lower values require more CPU processing time.

You can change one or both of the timers. To do so, use the following CLI method.

EXAMPLE:

To change the SPF delay and hold time, enter commands such as the following:

```
HP9300(config-ospf-router)# timers spf 10 20
```

The command in this example changes the SPF delay to 10 seconds and changes the SPF hold time to 20 seconds.

Syntax: timers spf <delay> <hold-time>

The <delay> parameter specifies the SPF delay.

The <hold-time> parameter specifies the SPF hold time.

To set the timers back to their default values, enter a command such as the following:

```
HP9300(config-ospf-router)# no timers spf 10 20
```

Possible values: see above

Default value: delay 5 seconds; hold time 10 seconds

trap

Generation of OSPF traps is enabled, by default, on the routing switch when OSPF is enabled. To disable all traps, use the global level CONFIG command **no snmp-server trap ospf**.

To stop a specific OSPF trap from being collected, use the CLI command **no trap** <trap>.

EXAMPLE:

To stop changes in the state of neighbors being generated by a routing switch, enter the following command:

```
HP9300(config-ospf-router)# no trap neighbor-state-change-trap
```

To reinstate the command, enter the following command:

```
HP9300(config-ospf-router)# trap neighbor-state-change-trap
```

Syntax: [no] trap <trap>

Possible values: see below: traps are from RFC 1850

Default value: All traps are active when OSPF is enabled.

interface-state-change-trap	[MIB object: OspfIfstateChange]
virtual-interface-state-change-trap	[MIB Object: OspfVirtIfStateChange
neighbor-state-change-trap	[MIB object:ospfNbrStateChange]
virtual-neighbor-state-change-trap	[MIB object: ospfVirtNbrStateChange]
interface-config-error-trap	[MIB object: ospfIfConfigError]
virtual-interface-config-error-trap	[MIB object: ospfIfConfigError]
interface-authentication-failure-trap	[MIB object: ospfIfAuthFailure]
virtual-interface-authentication-failure-trap	[MIB object: ospfVirtIfAuthFailure]
interface-receive-bad-packet-trap	[MIB object: ospfIfrxBadPacket]
virtual-interface-receive-bad-packet-trap	[MIB object: ospfVirtIfRxBadPacket]
interface-retransmit-packet-trap	[MIB object: ospfTxRetransmit]
virtual-interface-retransmit-packet-trap	[MIB object: ospfVirtIfTxRetransmit]
originate-lsa-trap	[MIB object: ospfOriginateLsa]
originate-maxage-lsa-trap	[MIB object: ospfMaxAgeLsa]
link-state-database-overflow-trap	[MIB object: ospfLsdbOverflow]
link-state-database-approaching-overflow-trap	[MIB object: ospfLsdbApproachingOverflow]

write memory

Saves the running configuration into the startup-config file.

EXAMPLE:

```
HP9300(config-bgp-router)# wr mem
```

Syntax: write memory

Possible values: N/A

Default value: N/A

write terminal

Displays the running configuration of the HP switch or routing switch on the terminal screen.

NOTE: This command is equivalent to the **show running-config** command.

EXAMPLE:

```
HP9300(config-bgp-router)# wr term
```

Syntax: write terminal

Possible values: N/A

Default value: N/A