

## Chapter 15

# OSPF Configuration Guidelines

To configure Open Shortest Path First version 2 (OSPF), you include the following statements:

```
protocols {
  ospf {
    disable;
    export [ policy-names ];
    external-preference preference;
    graceful-restart {
      disable;
      helper-disable;
      notify-duration seconds;
      restart-duration seconds;
    }
    import [ policy-names ];
    overload {
      <timeout seconds>;
    }
    preference preference;
    prefix-export-limit;
    rib-group group-name;
    reference-bandwidth reference-bandwidth;
    spf-delay;
    traffic-engineering {
      no-topology;
      shortcuts {
        lsp-metric-into-summary;
      }
    }
  }
  traceoptions {
    file name <replace> <size size> <files number> <no-stamp>
      <(world-readable | no-world-readable)>;
    flag flag <flag-modifier> <disable>;
  }
  area area-id {
    area-range network/mask-length <restrict>;
    authentication-type authentication;
    interface interface-name {
      disable;
      authentication {
        md5 key-id;
        simple-password key-id;
      }
    }
  }
}
```



For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

To configure OSPF version 3 (OSPFv3), you include the following statements:

```

protocols {
  ospf3 {
    disable;
    export [ policy-names ];
    external-preference preference;
    graceful-restart {
      disable;
      helper-disable;
      notify-duration seconds;
      restart-duration seconds;
    }
    import [ policy-names ];
    overload {
      <timeout seconds>;
    }
    preference preference;
    prefix-export-limit;
    reference-bandwidth reference-bandwidth;
    rib-group group-name;
    spf-delay;
    traceoptions {
      file name <replace> <size size> <files number> <no-stamp>
        <(world-readable | no-world-readable)>;
      flag flag <flag-modifier> <disable>;
    }
    area area-id {
      area-range network/mask-length <restrict>;
      interface interface-name {
        disable;
        dead-interval seconds;
        hello-interval seconds;
        metric metric;
        passive;
        priority number;
        retransmit-interval seconds;
        transit-delay seconds;
      }
    }
    nssa {
      area-range network/mask-length <restrict>;
      default-lsa {
        default-metric metric;
        metric-type type;
        type-7;
      }
      (no-summaries | summaries);
    }
  }
  stub <default-metric metric> <summaries | no-summaries>;
}

```

```

        virtual-link neighbor-id router-id transit-area area-id {
            disable;
            dead-interval seconds;
            hello-interval seconds;
            retransmit-interval seconds;
            transit-delay seconds;
        }
    }
}

```

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

By default, OSPFv2 and OSPFv3 are disabled.



**NOTE:** In this manual, the term *OSPF* refers to both OSPFv2 and OSPFv3.

---

This chapter describes the following tasks for configuring OSPF:

- Minimum OSPF Configuration on page 303
- Configuring the Backbone Area and Other Areas on page 304
- Configuring OSPF on Router Interfaces on page 307
- Configuring Authentication on page 309
- Configuring a Prefix Export Limit on page 310
- Configuring the Priority for Becoming the Designated Router on page 310
- Configuring Route Summarization on page 311
- Modifying the Interface Metric on page 311
- Configuring Route Preferences on page 312
- Configuring OSPF Timers on page 313
- Configuring the BFD Protocol on page 315
- Configuring Graceful Restart on page 315
- Configuring the SPF Delay on page 316
- Advertising Interface Addresses Without Running OSPF on page 316
- Advertising Label-Switched Paths into OSPF on page 317
- Configuring the Router to Appear Overloaded on page 317
- Enabling OSPF Traffic Engineering Support on page 318

Modifying the Traffic Engineering Metric on page 319

Configuring OSPF Routing Policy on page 320

Configuring OSPF Routing Table Groups on page 320

Configuring a Sham Link on page 321

Configuring a Peer Interface on page 321

Configuring Point-to-Point Adjacency over LAN Interfaces on page 322

Tracing OSPF Protocol Traffic on page 322

## Minimum OSPF Configuration

---

You must create a backbone area if your network consists of multiple areas. An area border router (ABR) must have at least one interface in the backbone area, or it must have a virtual link to a router in the backbone area. To do this, include at least the following statements. All other OSPF configuration statements are optional.

```
[edit]
protocols {
  (ospf | ospf3 ) {
    area 0 {
      interface interface-name;
    }
  }
}
```

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.



**NOTE:** When you configure OSPFv2 on an interface, you must also include the family inet statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level. When you configure OSPFv3 on an interface, you must also include the family inet6 statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level. For more information about the family inet statement, see the *JUNOS Network Interfaces and Class of Service Configuration Guide*.



**NOTE:** OSPFv3 does not support routing instances.

---

## Configuring the Backbone Area and Other Areas

---

You can group the routers in a single autonomous system (AS) into areas to reduce the amount of link-state advertisement (LSA) traffic on the network and to reduce the size of the topological databases that OSPF routers must maintain. If you do this, the AS must contain a single backbone area and optionally can contain any number of nonbackbone areas. The routers that make up the backbone must be physically contiguous. If they are not, you must configure virtual links to create the appearance of connectivity. You also can configure stub areas, which are areas through which AS external advertisements are not flooded, and not-so-stubby areas (NSSAs), which allow external routes to be flooded within an area.

To configure areas, you can perform the following tasks:

- Configuring the Backbone Area on page 304

- Configuring a Nonbackbone Area on page 304

- Configuring a Stub Area on page 305

- Configuring a Not-So-Stubby Area on page 305

- Configuring an OSPF Virtual Link on page 306

See page 307 for a virtual link configuration example.

### **Configuring the Backbone Area**

You must create a backbone area if your network consists of multiple areas. An ABR must have at least one interface in the backbone area, or it must have a virtual link to a router in the backbone area. The backbone comprises all area border routers and all routers that are not included in any other area. You configure all these routers by including the following area statement:

```
[edit protocols (ospf | ospf3)]
area 0.0.0.0;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

### **Configuring a Nonbackbone Area**

Each OSPF area consists of routers configured with the same area number. To configure a router to be in an area, include the area statement. The area number can be any number except 0.0.0.0, which is reserved for the backbone area.

```
[edit protocols (ospf | ospf3)]
area area-id;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

## Configuring a Stub Area

Stub areas are areas into which OSPF does not flood AS external advertisements. You might want to configure stub areas when much of the topological database consists of AS external advertisements and you want to minimize the size of the topological databases on an area's routers.

You cannot configure an area as being both a stub area and an NSSA.

To configure a stub area, include the `stub` statement:

```
[edit protocols (ospf | ospf3) area area-id]
  stub <default-metric metric> <(no-summaries | summaries)>;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

To inject a default route with a specified metric value into the area, include the `default-metric` option and a metric value. The default route matches any destination that is not explicitly reachable from within the area.

To have the stub areas not advertise summary routes into the stub area, include the `no-summaries` option. Only the default route is advertised, and only if you include the `default-metric` option. The default route injected into the not-so-stubby area (NSSA) is a Type 3 LSA.

You must include the `stub` statement when configuring all routers that are in the stub area.

## Configuring a Not-So-Stubby Area

An OSPF stub area has no external routes, so you cannot redistribute from another protocol into a stub area. An NSSA allows external routes to be flooded within the area. These routes are then leaked into other areas. However, external routes from other areas still do not enter the NSSA.

You cannot configure an area to be both a stub area and an NSSA.

To configure an NSSA, include the `nssa` statement:

```
[edit protocols (ospf | ospf3) area area-id]
  nssa {
    area-range network/mask-length <restrict>;
    default-lsa {
      default-metric metric;
      metric-type type;
      type-7;
    }
    (no-summaries | summaries);
  }
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

By default, a default route is not advertised. To advertise a default route with the specified metric within the area, include the `default-metric` statement. You can configure this option only on area border routers.

To prevent an ABR from advertising summary routes into an NSSA, include the `no-summaries` statement. If you include the `default-metric` option in addition to the `no-summaries` statement, only the default route is advertised. The default route is a Type 3 LSA injected into the NSSA. To flood summary LSAs into the NSSA area, include the `summaries` statement. When `summaries` is configured (which is the default if the `no-summaries` statement is not specified), a Type-7 LSA is sent. To define the type of metric, include the `metric-type` statement.

To aggregate external routes learned within the area when a route is advertised to other areas, include one or more `area-range` statements. If you also include the `restrict` option, the aggregate is not advertised, effectively creating a route filter. All external routes learned within the area that do not fall into the range of one of the prefixes are advertised individually to other areas.

### **Configuring an OSPF Virtual Link**

If any router on the backbone is not physically connected to the backbone itself, you must establish a virtual connection between that router and the backbone. You can establish a virtual connection between area border routers by configuring a OSPF virtual link.

To configure an OSPF virtual link, include the `virtual-link` statement when configuring the backbone area (area 0):

```
[edit protocols ospf area 0.0.0.0]
  virtual-link neighbor-id router-id transit-area area-id;
```

To configure an OSPFv3 virtual link, include the `virtual-link` statement when configuring the backbone area (area 0):

```
[edit protocols ospf3 area 0.0.0.0]
  virtual-link neighbor-id router-id transit-area area-id;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

Specify the router ID (as an IPv4 address) of the router at the other end of the virtual link. This router must be an area border router that is physically connected to the backbone. Also, specify the number of the area through which the virtual link transits.

For the virtual connection to work, you also must configure a link to the backbone area on the remote area border router (the router at the other end of the LSP).

**Example: Configuring an OSPF Virtual Link**

Configure an OSPF virtual link on the local router. This router must be an area border router that is physically connected to the backbone.

```
[edit protocols ospf]
area 0.0.0.0 {
  virtual-link neighbor-id 192.168.0.3 transit-area 1.1.1.1;
  interface t3-1/0/0 {
    hello-interval 1;
    dead-interval 3;
  }
}
```

You must also configure an OSPF virtual link on the remote area border router:

```
[edit protocols ospf]
area 0.0.0.0 {
  virtual-link neighbor-id 192.168.0.5 transit-area 1.1.1.1;
}
```

## Configuring OSPF on Router Interfaces

---

To enable OSPF on the router, you must configure OSPF on at least one of the router's interfaces. How you configure an interface depends on whether the interface is connected to a broadcast or point-to-point network, a point-to-multipoint network, or a nonbroadcast, multiaccess network.



**NOTE:** When you configure OSPFv2 on an interface, you must also include the family inet statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level. When you configure OSPFv3 on an interface, you must also include the family inet6 statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level. For more information about the family inet statement, see the *JUNOS Network Interfaces and Class of Service Configuration Guide*.

To configure OSPF on an interface, you can perform the following tasks:

Configuring an Interface on a Broadcast or Point-to-Point Network on page 307

Configuring an Interface on a Point-to-Multipoint Network on page 308

Configuring an Interface on a Nonbroadcast, Multiaccess Network on page 308

### Configuring an Interface on a Broadcast or Point-to-Point Network

If the interface on which you are configuring OSPF supports broadcast mode (such as a LAN), or if the interface supports point-to-point mode (such as a PPP interface or a point-to-point logical interface on Frame Relay), include the following form of the interface statement:

```
[edit protocols (ospf | ospf3) area area-id]
interface interface-name;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

Specify the interface by IP address or interface name for OSPFv2, or only the interface name for OSPFv3. For more information about interface names, see the *JUNOS Network Interfaces and Class of Service Configuration Guide*.

### Configuring an Interface on a Point-to-Multipoint Network

When you configure OSPFv2 on a nonbroadcast multiaccess (NBMA) network, such as a multipoint ATM or Frame Relay, OSPFv2 operates by default in point-to-multipoint mode. In this mode, OSPFv2 treats the network as a set of point-to-point links. Because there is no autodiscovery mechanism, each neighbor must be configured.

To configure OSPFv2 in point-to-multipoint mode, include the following statement:

```
[edit protocols ospf area 0.0.0.0]
interface interface-name {
    neighbor address;
}
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

Specify the interface by IP address or interface name. For more information about interface names, see the *JUNOS Network Interfaces and Class of Service Configuration Guide*.

To configure multiple neighbors, include a neighbor statement for each neighbor.

### Configuring an Interface on a Nonbroadcast, Multiaccess Network

When configuring OSPFv2 on an NBMA network, you can use nonbroadcast mode rather than point-to-multipoint mode. Using this mode offers no advantages over point-to-multipoint mode, but it has more disadvantages than point-to-multipoint mode. Nevertheless, you might occasionally find it necessary to configure nonbroadcast mode to interoperate with other equipment.

Nonbroadcast mode treats the NBMA network as a partially connected LAN, electing designated and backup designated routers. All routers must have a direct connection to both the designated and backup designated routers, or unpredictable results occur.

To configure nonbroadcast mode, include the following statements:

```
[edit protocols ospf area 0]
interface interface-name {
    interface-type nbma;
    neighbor address <eligible>;
    poll-interval seconds;
}
```

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

Specify the interface by IP address or interface name. For more information about interface names, see the *JUNOS Network Interfaces and Class of Service Configuration Guide*.



**NOTE:** For nonbroadcast interfaces, specify the IP address of the nonbroadcast interface as the *interface-name*.

---

To configure multiple neighbors, include a neighbor statement for each neighbor.

OSPF routers normally discover their neighbors dynamically by listening to the broadcast or multicast hello packets on the network. Because an NBMA network does not support broadcast (or multicast), the router cannot discover its neighbors dynamically, so you must configure all the neighbors statically. Do this by including the neighbor statement and specifying the IP address of each neighboring router in the *address* option. To configure multiple neighbors, include multiple neighbor statements. If the neighbor is allowed to become the designated router, include the eligible keyword.

By default, the router sends hello packets out the interface every 120 seconds before it establishes adjacency with a neighbor. To modify this interval, include the poll-interval statement.

## Configuring Authentication

---

All OSPFv2 protocol exchanges can be authenticated to guarantee that only trusted routers participate in the AS's routing. By default, OSPFv2 authentication is disabled. You can configure one of the following authentication methods. Each area must use the same method.

Simple authentication uses a text password that is included in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet.



**NOTE:** OSPFv3 does not support authentication.

---

The MD5 algorithm creates an encoded checksum that is included in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet.

For MD5 authentication to work, both the receiving and transmitting routers must have the same MD5 key. Define an MD5 key for each interface. If MD5 is enabled on an interface, that interface accepts routing updates only if MD5 authentication succeeds; otherwise, updates are rejected. The key ID can be set to any value between 0 and 255, with a default value of 0. The router only accepts OSPFv2 packets sent using the same key ID that is defined for that interface.

To enable authentication and specify an authentication method, include the `authentication-type` statement:

```
[edit protocols ospf area area-id]
authentication-type authentication;
```

The authentication type can be `none`, `simple`, or `md5`.

If you include the `authentication-type` statement to select an authentication method, you can configure a key (password) on each interface by including the `authentication` statement:

```
[edit protocols ospf area area-id interface interface-name]
authentication {
  md5 key-id;
  simple-password key-id;
}
```

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

The simple key (password) can be from 1 through 8 characters long. The MD5 key can be from 1 through 16 characters long. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").

The key identifier, which is required for MD5 authentication, specifies the identifier associated with the MD5 key.

## Configuring a Prefix Export Limit

---

By default, there is no limit to the number of prefixes that can be exported into OSPF. To limit the number of prefixes, include the `prefix-export-limit` statement:

```
[edit protocols (ospf | ospf3) level level-number]
prefix-export-limit number;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

The number can be a value from 0 through 4,294,967,295.

## Configuring the Priority for Becoming the Designated Router

---

A router advertises its priority to become a designated router in its hello packets. On all multiaccess networks, the Hello protocol uses the advertised priorities to elect a designated router for the network. This router is responsible for sending network link advertisements, which describe all the routers attached to the network. These advertisements are flooded throughout a single area.

At least one router on each logical IP network or subnet must be eligible to be the designated router for OSPFv2. At least one router on each logical link must be eligible to be the designated router for OSPFv3.

A router's priority for becoming the designated router is indicated by an arbitrary number from 0 through 255, with a higher value indicating a greater likelihood of becoming the designated router. By default, routers have a priority value of 128. A value of 1 means that the router has the least chance of becoming a designated router. A value of 0 marks the router as ineligible to become the designated router.

To modify the router's priority value, include the priority statement:

```
[edit protocols (ospf | ospf3) area area-id interface interface-name]
priority number;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

## Configuring Route Summarization

---

Area border routers send summary link advertisements to describe the routes to other areas. To minimize the number of these advertisements that are flooded, you can configure the router to coalesce, or summarize, a range of IP addresses and send reachability information about these addresses in a single link-state advertisement.

To summarize a range of IP addresses, include the area-range statement. To summarize multiple ranges, include multiple area-range statements.

```
[edit protocols (ospf | ospf3) area area-id]
area-range network/mask-length <restrict >;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

All routes that match the specified area range are filtered at the area boundary, and the summary is advertised in their place. If you specify the restrict option, the routes are filtered but no summary is advertised.

## Modifying the Interface Metric

---

All OSPF interfaces have a cost, which is a routing metric that is used in the link-state calculation. Routes with lower total path metrics are preferred over those with higher path metrics.

When several equal-cost routes to a destination exist, traffic is distributed equally among them.

The cost of a route is described by a single dimensionless metric that is determined using the following formula:

$$\text{cost} = \text{reference-bandwidth} / \text{bandwidth}$$

*reference-bandwidth* is the reference bandwidth. Its default value is 100 Mbps (which you specify as 100,000,000), which gives a metric of 1 for any bandwidth that is 100 Mbps or greater.

To modify the metric for routes advertised from an interface, include the metric statement:

```
[edit protocols (ospf | ospf3) area area-id interface interface-name]
metric metric;
```

To modify the reference bandwidth, include the reference-bandwidth statement:

```
[edit protocols (ospf | ospf3)]
reference-bandwidth reference-bandwidth;
```

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

For example, if you set the reference bandwidth to 1 Gbps (that is, *reference-bandwidth* is set to 1,000,000,000), a 100-Mbps interface has a default metric of 10.

By default, the loopback interface (lo0) metric is 0. No bandwidth is associated with the loopback interface.

## Configuring Route Preferences

---

Route preferences are used to select which route is installed in the forwarding table when several protocols calculate routes to the same destination. The route with the lowest preference value is selected. For more information about route preferences, see “Route Preferences” on page 6.

By default, internal OSPF routes have a preference value of 10, and external OSPF routes have a value of 150. To change the preference values, include the preference statement (for internal routes) or the external-preference statement (for external routes):

```
[edit protocols (ospf | ospf3)]
external-preference preference;
preference preference;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

The preference can be a value from 0 through 255.

## Configuring OSPF Timers

---

OSPF routers constantly track the status of their neighbors, sending and receiving hello packets that indicate whether the neighbor still is functioning, and sending and receiving link-state advertisement and acknowledgment packets. OSPF sends packets and expects to receive packets at specified intervals.

You can perform the following tasks when modifying the OSPF timers:

Modifying the Hello Interval on page 313

Controlling the LSA Retransmission Interval on page 314

Modifying the Router Dead Interval on page 314

Specifying the Transit Delay on page 314

### **Modifying the Hello Interval**

Routers send hello packets at a fixed interval on all interfaces, including virtual links, to establish and maintain neighbor relationships. This interval, which must be the same on all routers on a shared network, is advertised in the hello interval field in the hello packet. By default, the router sends hello packets every 10 seconds.

To modify how often the router sends hello packets out of an interface, include the `hello-interval` statement:

```
hello-interval seconds;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

On nonbroadcast networks, the router sends hello packets every 120 seconds until active neighbors are detected by default. This interval is long enough to minimize the bandwidth required on slow WAN links. To modify this interval, include the `poll-interval` statement:

```
[edit protocols ospf area area-id interface interface-name]  
poll-interval seconds;
```



**NOTE:** The `poll-interval` statement is valid for OSPFv2 only.

---

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

Once the router detects an active neighbor, the hello packet interval changes from the time specified in the `poll-interval` statement to the time specified in the `hello-interval` statement.

### **Controlling the LSA Retransmission Interval**

When a router sends link-state advertisements to its neighbors, the router expects to receive an acknowledgment packet from the neighbor within a certain amount of time. If the router does not receive an acknowledgment, it retransmits the advertisement.

By default, the router waits 5 seconds for an acknowledgment before retransmitting the link-state advertisement. To modify this interval, include the `retransmit-interval` statement:

```
retransmit-interval seconds;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

### **Modifying the Router Dead Interval**

If a router does not receive a hello packet from a neighbor within a fixed amount of time, the router modifies its topological database to indicate that the neighbor is nonoperational. The time that the router waits is called the *router dead interval*. By default, this interval is 40 seconds (four times the default hello interval).

To modify the router dead interval, include the `dead-interval` statement. This interval must be the same for all routers on a shared network.

```
dead-interval seconds;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

### **Specifying the Transit Delay**

Before a link-state update packet is propagated out of an interface, the router must increase the age of the packet. If you have a very slow link (for example, one with an average propagation delay of multiple seconds), the age of the packet must be increased by a similar amount. Doing this ensures that you do not receive a packet back that is younger than the original copy.

The default transit delay is 1 second. You should never have to modify the default value. However, if you need to specify the approximate transit delay to use to age update packets, include the `transit-delay` statement:

```
transit-delay seconds;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

## Configuring the BFD Protocol

---

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. The failure detection timers for BFD have shorter time limits than the failure detection mechanisms of OSPF, providing faster detection. These timers are also adaptive and can be adjusted to be more or less aggressive.

To enable failure detection, include the `bfd-liveness-detection` statement:

```
bfd-liveness-statement {
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    minimum-transmit-interval milliseconds;
    multiplier number;
}
```

To specify the minimum transmit and receive interval for failure detection, include the `minimum-interval` statement:

```
minimum-interval milliseconds;
```

To specify the minimum receive interval for failure detection, include the `minimum-receive-interval` statement:

```
minimum-receive-interval milliseconds;
```

To specify the minimum transmit interval for failure detection, include the `minimum-transmit-interval` statement:

```
minimum-transmit-interval milliseconds;
```

To specify the detection time multiplier for failure detection, include the `multiplier` statement:

```
multiplier number;
```

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

## Configuring Graceful Restart

---

Graceful restart is disabled by default. You can globally enable graceful restart for all routing protocols at the `[edit routing-options]` hierarchy level.

To configure graceful restart parameters specifically for OSPF, include the `graceful-restart` statement:

```
[edit protocols (ospf | ospf3)]
graceful-restart {
    notify-duration seconds;
    restart-duration seconds;
}
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

To disable graceful restart, specify the `disable` statement. To configure a time period for complete reacquisition of OSPF neighbors, specify the `restart-duration` statement. To configure a time period for sending out purged grace LSAs over all interfaces, specify the `notify-duration` statement. To disable the graceful restart helper capability, specify the `helper-disable` statement.

OSPF supports two types of graceful restart: planned and unplanned. During a planned restart, the restarting router informs the neighbors before restarting. The neighbors act as if the router is still within the network topology, and continue forwarding traffic to the restarting router. A grace period is set to specify the time period for which the neighbors should consider the restarting router as part of the topology. During an unplanned restart, the router restarts without warning.

## Configuring the SPF Delay

---

You can configure the shortest path first (SPF) algorithm delay. The SPF algorithm delay is the amount of time, in milliseconds, between the detection of a topology change and when the SPF algorithm actually runs to achieve convergence. The shorter the delay, the shorter the convergence time.

To configure the SPF delay, include the `spf-delay` statement:

```
[edit protocols (ospf | ospf3)]
spf-delay milliseconds;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

The time can range from 50 through 1000 milliseconds.

## Advertising Interface Addresses Without Running OSPF

---

By default, OSPF must be configured on an interface for direct interface addresses to be advertised as interior routes. To advertise the direct interface addresses without actually running OSPF on that interface, include the `passive` statement:

```
[edit protocols (ospf | ospf3) interface interface-name]
passive;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

Point-to-point interfaces are different from multipoint in that only one OSPF adjacency is possible. (A LAN, for instance, can have multiple addresses and can run OSPF on each subnet simultaneously.) As such, when you configure a numbered point-to-point interface to OSPF by name, multiple OSPF interfaces are created. One, which is unnumbered, is the interface on which the protocol is run. An additional OSPF interface is created for each address configured on the interface, if any, which is automatically marked as passive.

For OSPFv3, one OSPF-specific interface must be created per interface name configured under OSPFv3. OSPFv3 does not allow interfaces to be configured by IP address.

Enabling OSPF on an interface (by including the interface statement), disabling it (by including the disable statement), and not actually having OSPF run on an interface (by including the passive statement) are mutually exclusive states.

## Advertising Label-Switched Paths into OSPF

---

You can advertise label-switched paths (LSPs) into OSPFv2 as point-to-point links so that all participating routers can take the LSP into account when performing SPF calculations. The advertisement contains a local address (the from address of the label-switched path), a remote address (the to address of the label-switched path), and a metric with the following precedence:

1. Use the label-switched path metric defined under OSPFv2.
2. Use the label-switched path metric configured for the label-switched path under MPLS.
3. If you do not configure any of the above, use the default OSPFv2 metric of 1.

To advertise LSPs, include the label-switched-path statement, with a specified name and metric:

```
[edit protocols ospf area area-id
 label-switched-path name metric metric;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.



**NOTE:** If you want an LSP that is announced into OSPFv2 to be used in SPF calculations, there must be a reverse link (that is, a link from the tail end of the LSP to the head end). You can accomplish this by configuring an LSP in the reverse direction and also announcing it in OSPFv2.

---

For more information about advertising label-switched paths, see the *JUNOS MPLS Applications Configuration Guide*.

## Configuring the Router to Appear Overloaded

---

If the time elapsed after the OSPF instance is enabled is less than the specified timeout, overload mode is set.

You can configure the local router so that it appears to be overloaded. You might do this when you want the router to participate in OSPF routing, but do not want it to be used for transit traffic. (Traffic to directly attached interfaces continues to transit the router.)

You configure or disable overload mode in OSPF with or without a timeout. Without a timeout, overload mode is set until it is explicitly deleted from the configuration. With a timeout, overload mode is set if the time elapsed since the OSPF instance started is less than the specified timeout.

A timer is started for the difference between the timeout and the time elapsed since the instance started. When the timer expires, overload mode is cleared. In overload mode, the router LSA is originated with all the transit router links (except stub) set to a metric of 0xFFFF. The stub router links are advertised with the actual cost of the interfaces corresponding to the stub. This causes the transit traffic to avoid the overloaded router and take paths around the router. However, the overloaded router's own links are still accessible.

To mark the router as overloaded, include the overload statement:

```
[edit protocols (ospf | ospf3)]
overload;
```

To specify the number of seconds at which overload is reset, include the timeout option when specifying the overload statement:

```
[edit protocols (ospf | ospf3)]
overload timeout <seconds>;
```

The time can be a value from 60 through 1800 seconds.

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

## Enabling OSPF Traffic Engineering Support

---

When traffic engineering is enabled on the router, you can enable the OSPFv2 traffic engineering support, which allows OSPFv2 to generate LSAs that carry traffic engineering parameters. These parameters are used to create the Traffic Engineering Database (TED), which is used by Constrained Shortest Path First (CSPF) to compute MPLS LSPs.



**NOTE:** Traffic engineering is not supported for OSPFv3.



**NOTE:** Whenever possible, use IS-IS IGP shortcuts instead of traffic engineering shortcuts.

---

By default, traffic engineering support is disabled. To enable it, include the traffic-engineering statement:

```
[edit protocols ospf]
traffic-engineering {
  no-topology;
  shortcuts {
    lsp-metric-into-summary;
  }
}
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

To disable the dissemination of the link-state topology information, specify the `no-topology` statement. To use LSPs as next hops, specify the `shortcuts` statement.

When traffic engineering is enabled for OSPF, the SPF algorithm takes into account the various LSPs configured under MPLS. These routes are installed into the primary routing table, `inet.0`. To advertise the LSP metric for a prefix in a summary LSA, specify the `lsp-metric-into-summary` statement.

For more information about configuring LSPs and MPLS, see the *JUNOS MPLS Applications Configuration Guide*.

### Example: Enabling OSPF Traffic Engineering Support

Enable OSPF traffic engineering support by configuring a virtual link on the local router. This router must be an area border router that is physically connected to the backbone.

```
[edit protocols]
ospf {
  traffic-engineering {
    shortcuts {
      lsp-metric-into-summary;
    }
  }
}

[edit protocols]
mpls {
  traffic-engineering bgp-igp;
  label-switched-path xxxx {
    to yy.yy.yy.yy
  }
}
```

## Modifying the Traffic Engineering Metric

---

When traffic engineering is enabled on the router, you can configure an OSPF metric that is used exclusively for traffic engineering. The traffic engineering metric is used for information injected into the Traffic Engineering Database (TED). Its value does not affect normal OSPF forwarding.

To modify the default value, include the `te-metric` statement:

```
[edit protocols ospf area area interface interface-name]
te-metric metric;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

## Configuring OSPF Routing Policy

---

All routing protocols store the routes that they learn in the routing table. The routing table uses this collected route information to determine the active routes to destinations. The routing table then installs the active routes into its forwarding table and also exports them back into the routing protocols. It is these exported routes that the protocols advertise.

For each protocol, you control which routes the protocol stores in the routing table and which routes the routing table exports into the protocol by defining a *routing policy* for that protocol. For information about defining a routing policy, see the *JUNOS Policy Framework Configuration Guide*.

By default, if a router has multiple OSPF areas, learned routes from other areas are automatically installed into area 0 of the routing table.

To apply routing policies that affect how the routing table exports routes into OSPF, include the export statement:

```
[edit protocols (ospf | ospf3)]
  export [ policy-names ];
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

OSPF import policy allows users to define policy to prevent adding OSPF routes to the routing table. This filtering happens when OSPF installs the route in the routing table. You can filter the routes, but not LSA flooding. The import policy can filter on any attribute of the OSPF route.

To filter OSPF routes from being added to the routing table, include the import statement:

```
[edit protocols (ospf | ospf3)]
  import [ policy-names ];
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.



**NOTE:** For OSPF, you cannot apply routing policies that affect how routes are imported into the routing table; doing so with a link-state protocol can result in an inconsistent topology database.

---

## Configuring OSPF Routing Table Groups

---

To install routes learned from OSPF routing instances into routing tables in the OSPF routing table group, include the rib-group statement:

```
[edit protocols (ospf | ospf3)]
  rib-group group-name;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

For a multiple routing instance configuration example, see “Routing Instances Configuration Guidelines” on page 165.

## Configuring a Sham Link

---

You can create an intra-area link or sham link between two provider edge (PE) routers so that the VPN backbone is preferred over the back-door link. Each sham link is identified by the combination of a local endpoint address and a remote endpoint address.

To configure a sham link, include the sham-link statement:

```
[edit protocols ospf area area-id]
sham-link {
  local-endpoint address;
  remote-endpoint address {
    metric metric;
  }
}
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

To configure the local endpoint address, specify the local-address statement. To configure the remote endpoint address, specify the remote-address statement. To configure the remote endpoint metric value, specify the metric statement.

## Configuring a Peer Interface

---

You can configure a peer interface for OSPF routers. Generalized Multiprotocol Label Switching (GMPLS) requires traffic engineering information to be transported through a link separate from the control channel. You establish this separate link by configuring a peer interface.

To configure a peer interface, include the peer-interface statement:

```
[edit protocols ospf area area-id]
peer-interface interface-name {
  disable;
  dead-interval seconds;
  hello-interval seconds;
  retransmit-interval seconds;
  transit-delay seconds;
}
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

To disable the peer interface, specify the disable statement. To modify the peer interface dead interval, specify the dead-interval statement. To modify how often the router sends hello packets out of the peer interface, specify the hello-interval statement. To modify how often the peer interface retransmits the link-state advertisement, specify the retransmit-interval statement. To specify the approximate transit delay to use to age update packets, include the transit-delay statement.

For more information about configuring GMPLS, see the *JUNOS MPLS Applications Configuration Guide*.

## Configuring Point-to-Point Adjacency over LAN Interfaces

---

You can configure OSPF point-to-point adjacency over Ethernet and aggregated Ethernet interfaces. OSPF treats point-to-point interfaces and broadcast interfaces differently. The differences comprise the establishment of neighbor adjacencies, flooding link-state information, representation of the network topology, the SPF calculation, and protocol packets destination addresses.

You can configure OSPF so that it treats LAN interfaces with only two routers connected to them as point-to-point interfaces. The point-to-point extension modifies the shortest-path-first (SPF) calculation so that the correct next hop is calculated. The next-hop address is null for point-to-point interfaces by default and is rejected.

To configure an interface as a point-to-point interface in OSPF, include the point-to-point statement:

```
[edit protocols ospf area area-id interface interface-name]
p2p;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

Two routers are supported for each subnet. The LAN interface must be numbered.

For more information about point-to-point interfaces, see “Configuring an Interface on a Point-to-Multipoint Network” on page 308.

## Tracing OSPF Protocol Traffic

---

To trace OSPF protocol traffic, you can specify options with the global traceoptions statement at the [edit routing-options] hierarchy level, and you can specify OSPF-specific options by including the traceoptions statement:

```
[edit protocols (ospf | ospf3)]
traceoptions {
  file name <replace> <size size> <files number> <no-stamp>
    <(world-readable | no-world-readable)>;
  flag flag <flag-modifier> <disable>;
}
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

You can specify the following OSPF-specific trace flags in the OSPF traceoptions statement:

all—Everything

database-description—All database description packets, which are used in synchronizing the OSPF topological database

error—OSPF error packets

event—OSPF state transitions

flooding—Link-state flooding packets

general—General events

hello—Hello packets, which are used to establish neighbor adjacencies and to determine whether neighbors are reachable

lsa-ack—Link-state acknowledgment packets, which are used in synchronizing the OSPF topological database

lsa-request—Link-state request packets, which are used in synchronizing the OSPF topological database

lsa-update—Link-state updates packets, which are used in synchronizing the OSPF topological database

normal—Normal events

packets—All OSPF packets

packet-dump—Dump the contents of selected packet types

policy—Policy processing

spf—Shortest path first (SPF) calculations

state—State transitions

task—Routing protocol task processing

timer—Routing protocol timer processing



**NOTE:** Use the traceoption flags detail and all with caution. These flags may cause the CPU to become very busy.

---

For general information about tracing and global tracing options, see “Tracing Global Routing Protocol Operations” on page 99.

**Examples: Tracing OSPF Protocol Traffic**

Trace only unusual or abnormal operations to the file routing-log, and trace detailed information about all OSPF packets to the file ospf-log:

```
[edit]
routing-options {
  traceoptions {
    file routing-log;
  }
}
protocols {
  ospf {
    traceoptions {
      file ospf-log size 10k files 5;
      flag lsa-ack;
      flag database-description;
      flag hello;
      flag lsa-update;
      flag lsa-request;
    }
    area 0.0.0.0 {
      interface 10.0.0.1;
    }
  }
}
```

Trace SPF calculations:

```
[edit]
protocols {
  ospf {
    traceoptions {
      file ospf-log;
      flag spf;
    }
    area 0.0.0.0 {
      interface 10.0.0.1;
    }
  }
}
```

Trace the creation, receipt, and retransmission of all link-state advertisements:

```
[edit]
protocols {
  ospf {
    traceoptions {
      file ospf-log;
      flag lsa-request;
      flag lsa-update;
      flag lsa-ack;
    }
    area 0.0.0.0 {
      interface 10.0.0.1;
    }
  }
}
```