Chapter 19
OSPF Configuration Guidelines

To configure OSPF, you include statements at the [edit protocols ospf] hierarchy level of the configuration.

```
protocols {
    ospf {
        disable;
        export [policy-names];
        external-preference preference;
        preference preference;
        rib-group routing-table-group-name;
        reference-bandwidth reference-bandwidth;
        traffic-engineering {
            no-topology;
            shortcuts;
        }
        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-key key <key-id identifier>;
                dead-interval seconds;
                hello-interval seconds;
                interface-type type;
                metric metric;
                neighbor address <eligible>;
                passive;
                poll-interval seconds;
                priority number;
                retransmit-interval seconds;
                transit-delay seconds;
                transmit-interval seconds;
            }
            nssa {
                area-range network/mask-length <restrict>;
                default-metric metric;
                (summaries | no-summaries);
            }
            stub <default-metric metric> <summaries | no-summaries>;
```

```
                virtual-link neighbor-id router-id transit-area area-id {
                    disable;
                    authentication-key key <key-id identifier>;
                    dead-interval seconds;
                    hello-interval seconds;
                    retransmit-interval seconds;
                    transit-delay seconds;
                }
            }
        }
    }
```

By default, OSPF is disabled.

This chapter describes the following tasks for configuring OSPF:

Minimum OSPF Configuration on page 222

Configure the Backbone Area and Other Areas on page 223

Configure OSPF on Router Interfaces on page 225

Configure Authentication on page 227

Configure the Priority for Becoming the Designated Router on page 228

Configure Route Summarization on page 228

Modify the Interface Metric on page 228

Configure Route Preferences on page 229

Configure OSPF Timers on page 230

Advertise Interface Addresses without Running OSPF on page 232

Enable OSPF Traffic Engineering Support on page 232

Configure OSPF Routing Policy on page 233

Install OSPF Multiple Routing Instances on page 233

Trace OSPF Protocol Traffic on page 233

Minimum OSPF Configuration

For OSPF to run on the router, you must configure a backbone area on at least one interface. To do this, include at least the following statements. All other OSPF configuration statements are optional.

```
[edit]
protocols {
    ospf {
        area 0 {
            interface interface-name;
        }
    }
}
```

Configure the Backbone Area and Other Areas

You can group the routers in a single AS into areas to reduce the amount of link-state advertisement 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 NSSAs, which allow external routes to be flooded within an area.

To configure areas, you can perform the following tasks:

Configure the Backbone Area on page 223

Configure a Nonbackbone Area on page 223

Configure a Stub Area on page 223

Configure a Not-So-Stubby Area on page 224

Configure a Virtual Link on page 225

Configure the Backbone Area

To configure OSPF on the router, you must configure a 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 at the [edit protocols ospf] hierarchy level:

```
[edit protocols ospf]
area 0.0.0.0;
```

Configure 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 at the [edit protocols ospf] hierarchy level, specifying the area number for that area. The area number can be any number except 0.0.0.0, which is reserved for the backbone area.

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

Configure 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 at the [edit protocols ospf area *area-id*] hierarchy level:

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

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.

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

Configure a Not-So-Stubby Area

An OSPF stub area has no external routes in it, so you cannot redistribute from another protocol into a stub area. A not-so-stubby area (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 as being both a stub area and an NSSA.

To configure an NSSA, include the nssa statement at the [edit protocols ospf area *area-id*] hierarchy level:

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

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 area border routers from advertising summary routes into an NSSA, include the no-summaries statement. If you include the default-metric option, only the default route is advertised. To flood summary LSAs into the NSSA area, include the summaries 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.

Configure a Virtual Link

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

To configure a 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;
```

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

Example: Configure a Virtual Link

Configure 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}
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 a 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;
}
```

Configure 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 IS-IS on an interface, you must also include the family iso statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level. For more information about the family inet statement, see the *JUNOS Internet Software Configuration Guide: Interfaces and Chassis*.

Configure 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 at the [edit protocols ospf area *area-id*] hierarchy level:

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

Specify the interface by IP address. For more information about interface names, see interface naming in the *JUNOS Internet Softw are Configur ation Guide: Interf aces and Chassis*.

Configure an Interface on a Point-to-Multipoint Network

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

To configure OSPF in point-to-multipoint mode, include the following statements at the [edit protocols ospf area 0.0.0.0] hierarchy level:

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

Specify the interface by IP address. For more information about interface names, see interface naming in the *JUNOS Internet Softw are Configur ation Guide: Interf aces and Chassis*.

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

Configure an Interface on a Nonbroadcast, Multiaccess Network

When configuring OSPF on a nonbroadcast, multiaccess (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 at the [edit protocols ospf area 0] hierarchy level:

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

Specify the interface by IP address. For more information about interface names, see interface naming in the *JUNOS Internet Softw are Configur ation Guide: Interf aces and Chassis*.

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.

Configure Authentication

All OSPF protocol exchanges can be authenticated to guarantee that only trusted routers participate in the AS's routing. By default, OSPF 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.
- 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.

To enable authentication and specify an authentication method, include the authentication-type statement at the [edit protocols ospf area *area-id*] hierarchy level:

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

authentication can be none, simple, or md5.

If you have included the authentication-type statement to select an authentication method, you can configure a key (password) on each interface by including the authentication-key statement at the [edit protocols ospf area *area-id* interface *interface-name*] hierarchy level:

```
[edit protocols ospf area area-id interface interface-name]
authentication-key key key-id identifier;
```

The key (password) can be 1 to 8 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. You can define multiple keys, each with a different key identifier. OSPF uses the key with the highest number.

Configure 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 OSPF 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.

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 at the [edit protocols ospf area *area-id* interface *interface-name*] hierarchy level:

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

Configure 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 at the [edit protocols ospf area *area-id*] hierarchy level. To summarize multiple ranges, include multiple area-range statements.

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

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.

Modify the Interface Metric

All OSPF interfaces have a cost, which is a routing metric that is used in the OSPF 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:

cost = reference-bandwidth / bandwidth *reference-bandwidth* is the reference bandwidth. Its default value is 100 Mbps (which you specify as 100000000), 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 at the [edit protocols ospf area *area-id* interface *interface-name*] hierarchy level:

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

To modify the reference bandwidth, include the reference-bandwidth statement at the [edit protocols ospf] hierarchy level:

```
[edit protocols ospf]
reference-bandwidth reference-bandwidth;
```

For example, if you set the reference bandwidth to 1 Gbps (that is, *reference-bandwidth* is set to 1000000000), 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.

Configure 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 5.

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) at the [edit protocols ospf] hierarchy level:

```
[edit protocols ospf]
external-preference preference;
preference preference;
```

The preference value can range from 0 through 255.

Configure OSPF Timers

OSPF routers constantly track the status of their neighbors, sending and receiving hello packets that indicate that 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 modify the following OSPF timers:

Modify the Hello Interval on page 230

Modify the LSA Transmission Interval on page 231

Control the LSA Retransmission Interval on page 231

Modify the Router Dead Interval on page 231

Specify the Transit Delay on page 231

Modify 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 at the [edit protocols ospf area *area-id* interface *interface-name*] or [edit protocols ospf area *area-id* virtual-link] hierarchy level:

```
hello-interval seconds
```

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 at the [edit protocols ospf area *area-id* interface *interface-name*] hierarchy level.

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

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.

Modify the LSA Transmission Interval

The transmit interval specifies how often OSPF link-state advertisement packets are transmitted on an interface. This interval determines the maximum packet transmission rate on an interface, which affects network stability. Because packets are built at the instant of transmission, only the latest information is sent even if the transmission is delayed.

The default transmit interval is 30 milliseconds. To modify the interval, include the transmit-interval statement at the [edit protocols ospf area *area-id* interface *interface-name*] hierarchy level:

```
[edit protocols ospf area area-id interface interface-name]
transmit-interval milliseconds;
```

Control 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 at the [edit protocols ospf area *area-id* interface *interface-name*] or [edit protocols ospf area *area-id* virtual-link] hierarchy level:

```
retransmit-interval seconds;
```

Modify 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 interv al*. By default, this interval is 40 seconds (four times the default hello interval).

To modify the router dead interval, include the dead-interval statement at the [edit protocols ospf area *area-id* interface *interface-name*] or [edit protocols ospf area *area-id* virtual-link] hierarchy level. This interval must be the same for all routers on a shared network.

```
dead-interval seconds;
```

Specify 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 at the [edit protocols ospf area *area-id* interface *interface-name*] or [edit protocols ospf area *area-id* virtual-link] hierarchy level:

```
transit-delay seconds;
```

Advertise 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 at the [edit protocols ospf interface *interface-name*] hierarchy level:

```
[edit protocols ospf interface interface-name]
passive;
```

Point-to-point interfaces are different than 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 one 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.

Enabling OSPF on an interface (by including the interface statement at the [edit protocols ospf] hierarchy level), 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.

Enable OSPF Traffic Engineering Support

When traffic engineering is enabled on the router, you can enable OSPF's traffic engineering support, which allows OSPF to generate LSAs that carry traffic engineering parameters. These parameters are used to create the traffic engineering database (TED), which is used by CSPF to compute MPLS label-switched paths (LSPs).

By default, the traffic engineering support is disabled. To enable it, include the traffic-engineering statement at the [edit protocols ospf] hierarchy level:

```
[edit protocols ospf]
traffic-engineering {
    no-topology;
    shortcuts;
}
```

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

Configure 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 from the routing table by defining a *routing policy* for that protocol. For information about defining routing policy, see “Configure Routing Policy” on page 35.

For OSPF, you can apply routing policies that affect how the routing table exports routes into OSPF. To do this, include the export statement at the [edit protocols ospf] hierarchy level:

```
[edit protocols ospf]
export [policy-names];
```

> **Note**
> For OSPF, you should not apply routing policies that affect how routes are imported into the routing table; doing so with a link-state protocol could easily lead to an inconsistent topology database.

Install OSPF Multiple Routing Instances

You can install routes learned from OSPF instances into routing tables in the OSPF routing table group. To do this, include the rib-group statement at the [edit protocols ospf] hierarchy level:

```
[edit protocols ospf]
rib-group routing-table-group-name;
```

For a multiple routing instance configuration example, see “Configure Multiple Routing Instances” on page 163.

For information about routing instances, see “Configure Static Routes” on page 93.

Trace OSPF Protocol Traffic

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

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

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

- all—Trace everything.
- database-description—Trace all database description packets, which are used in synchronizing the OSPF topological database.
- error—Trace OSPF errored packets.
- event—Trace OSPF state transitions.
- flooding—Trace link-state flooding packets.
- general—Trace general events.
- hello—Trace hello packets, which are used to establish neighbor adjacencies and to determine whether neighbors are reachable.
- lsa-ack—Trace link-state acknowledgment packets, which are used in synchronizing the OSPF topological database.
- lsa-request—Trace link-state request packets, which are used in synchronizing the OSPF topological database.
- lsa-update—Trace link-state updates packets, which are used in synchronizing the OSPF topological database.
- normal—Trace normal events.
- packets—Trace all OSPF packets.
- packet-dump—Dump the contents of selected packet types.
- policy—Trace policy processing.
- spf—Trace shortest-path-first (SPF) calculations.
- state—Trace state transitions.
- task—Trace routing protocol task processing.
- timer—Trace routing protocol timer processing.

For general information about tracing and global tracing options, see “Trace Global Routing Protocol Operations” on page 128.

Examples: Trace 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;
        }
    }
}
```