

Chapter 14

OSPF Overview

The Open Shortest Path First version 2 (OSPF) protocol is an interior gateway protocol (IGP) that routes packets within a single autonomous system (AS). OSPF uses link-state information to make routing decisions.

In this document, the term *OSPF* refers to both OSPFv2 and OSPFv3.

This chapter discusses the following topics that provide background information about OSPF:

OSPF Protocol Overview on page 289

OSPF Standards on page 291

OSPF Area Terminology on page 292

OSPF Routing Algorithm on page 294

OSPF Packets on page 295

External Metrics on page 298

Designated Router on page 298

OSPF Extensions to Support Traffic Engineering on page 298

OSPF Protocol Overview

OSPF is an IGP that routes packets within a single AS. OSPF uses link-state information to make routing decisions, making route calculations using the shortest path first (SPF) algorithm (also referred to as the Dijkstra algorithm). Each router running OSPF floods link-state advertisements throughout the AS that contain information about that router's attached interfaces and routing metrics. Each router takes the information in these link-state advertisements and creates a complete routing table for the network.

The JUNOS software supports OSPF version 2, including virtual links, stub areas, and authentication. The JUNOS software does not support type-of-service (ToS) routing.

OSPF was designed for the Transmission Control Protocol/Internet Protocol (TCP/IP) environment and as a result explicitly supports IP subnetting and the tagging of externally derived routing information. OSPF also provides for the authentication of routing updates.

OSPF routes IP packets based solely on the destination IP address contained in the IP packet header. OSPF quickly detects topological changes, such as when router interfaces become unavailable, and calculates new loop-free routes quickly and with a minimum of routing overhead traffic.

Each interface running OSPF is assigned a cost, which is a unitless number based on factors such as throughput, round-trip time, and reliability, which are used to determine how easy or difficult it is to reach a destination. If two or more routes to a destination have the same cost, OSPF distributes traffic equally among the routes, a process that is called *load balancing*.

Each router maintains a database that describes the topology of the AS. Each OSPF router has an identical topological database so that all routers in the area have a consistent view of the network. All routers maintain summarized topologies of other areas within an AS. Each router distributes information about its local state by flooding link-state advertisements throughout the AS. When the AS topology changes, OSPF ensures that the contents of all routers' topological databases converge quickly.

All OSPF protocol exchanges can be authenticated. This means that only trusted routers can participate in the AS's routing. A variety of authentication schemes can be used; a single authentication scheme is configured for each area, which enables some areas to use stricter authentication than others.

Externally derived routing data (for example, routes learned from BGP) is passed transparently throughout the AS. This externally derived data is kept separate from the OSPF link-state data. Each external route can be tagged by the advertising router, enabling the passing of additional information between routers on the boundaries of the AS.

OSPF Version 3

OSPFv3 is a modified version of OSPF that supports Internet Protocol version 6 (IPv6) addressing. OSPFv3 differs from OSPFv2 in the following ways:

- All neighbor ID information is based on a 32-bit router ID.

- The protocol runs per link rather than per subnet.

- Router and network link-state advertisements (LSAs) do not carry prefix information.

- Two new LSA types are included: link-LSA and intra-area-prefix-LSA.

Flooding scopes are as follows:

Link-local

Area

AS

Link-local addresses are used for all neighbor exchanges except virtual links.

Authentication is removed; the IPv6 authentication header relies on the IP layer.

The packet format has changed as follows:

Version number 2 is now version number 3.

The db option field has been expanded to 24 bits.

Authentication information has been removed.

Hello messages do not have address information.

Two new option bits are included: R and V6.

Type 3 summary LSAs have been renamed inter-area-prefix-LSAs.

Type 4 summary LSAs have been renamed inter-area-router-LSAs.

OSPF Standards

OSPF is defined in the following documents:

RFC 1587, *The OSPF NSS A Option*

RFC 2328, *OSPF Version 2*

RFC 2740, *OSPF for IPv6*

RFC 3623, *OSPF Graceful Restart*

Internet draft draft-katz-yeung-ospf-traffic-01.txt, *Traffic Engineering Extensions to OSPF* (expires April 2000)

Internet draft draft-ietf-katz-ward-bfd-00.txt, *Bidirectional Forwarding Detection* (except the transmission of echo packets) (expires January 2005)

Internet draft draft-ietf-ccamp-ospf-gmpls-extensions-09.txt, *OSPF Extensions in Support of Generalized MPLS* (only interface switching) (expires June 2003)

Internet draft draft-ietf-isis-igp-p2p-over-lan-03.txt, *Point-to-point operation over LAN in link-state routing protocols* (expires February 2004)

To access Internet RFCs and drafts, go to the IETF Web site at <http://www.ietf.org>.

OSPF Area Terminology

In OSPF, a single AS can be divided into smaller groups called *areas*. This reduces the number of link-state advertisements and other OSPF overhead traffic sent on the network, and it reduces the size of the topological database that each router must maintain.

This section discusses the following topics:

Areas on page 292

Area Border Routers on page 292

Backbone Areas on page 293

AS Boundary Routers on page 293

Stub Areas on page 293

Not-So-Stubby Areas on page 293

Transit Areas on page 293

Areas

An *area* is a set of networks and hosts within an AS that have been administratively grouped together. We recommend that you configure an area as a collection of contiguous IP subnetted networks. Routers that are wholly within an area are called *internal routers*. All interfaces on internal routers are directly connected to networks within the area.

The topology of an area is hidden from the rest of the AS, thus significantly reducing routing traffic in the AS. Also, routing within the area is determined only by the area's topology, providing the area with some protection from bad routing data.

All routers within an area have identical topological databases.

Area Border Routers

Routers that belong to more than one area are called *area border routers*. They maintain a separate topological database for each area to which they are connected.

Backbone Areas

An OSPF *backbone area* consists of all networks in area ID 0.0.0.0, their attached routers, and all area border routers. The backbone itself does not have any area border routers. The backbone distributes routing information between areas. The backbone is simply another area, so the terminology and rules of areas apply: a router that is directly connected to the backbone is an internal router on the backbone, and the backbone's topology is hidden from the other areas in the AS.

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 backbone connectivity. You can create virtual links between any two area border routers that have an interface to a common nonbackbone area. OSPF treats two routers joined by a virtual link as if they were connected to an unnumbered point-to-point network.

AS Boundary Routers

Routers that exchange routing information with routers in other ASs are called *AS boundary routers*. They advertise externally learned routes throughout the AS. Any router in the AS—an internal router, an area border router, or a backbone router—can be an AS boundary router.

Every router within the AS knows the path to the AS boundary routers.

Stub Areas

Stub areas are areas through which or into which AS external advertisements are not flooded. You might want to create stub areas when much of the topological database consists of AS external advertisements. Doing so reduces the size of the topological databases and therefore the amount of memory required on the internal routers in the stub area.

When an area border router is configured for a stub area, the router automatically advertises a default route in place of the external routes that are not being advertised within the stub area so that routers in the stub area can reach destinations outside the area.

The following restrictions apply to stub areas: you cannot create a virtual link through a stub area, and a stub area cannot contain an AS boundary router.

Not-So-Stubby Areas

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.

Transit Areas

Transit areas are used to pass traffic from one adjacent area to the backbone (or to another area if the backbone is more than two hops away from an area). The traffic does not originate in, nor is it destined for, the transit area.

OSPF Routing Algorithm

OSPF uses the shortest path first (SPF) algorithm, also referred to as the Dijkstra algorithm, to determine the route to reach each destination. All routers in an area run this algorithm in parallel, storing the results in their individual topological databases. Routers with interfaces to multiple areas run multiple copies of the algorithm. This section provides a brief summary of how the SPF algorithm works.

When a router starts, it initializes OSPF and waits for indications from lower-level protocols that the router interfaces are functional. The router then uses the OSPF hello protocol to acquire neighbors, doing this by sending hello packets to its neighbors and receiving their hello packets.

On broadcast or nonbroadcast multiaccess networks (physical networks that support the attachment of more than two routers), the OSPF hello protocol elects a designated router for the network. This router is responsible for sending *link-state advertisements* that describe the network, which reduces the amount of network traffic and the size of the routers' topological databases.

The router then attempts to form *adjacencies* with some of its newly acquired neighbors. (On multiaccess networks, only the designated router and backup designated router form adjacencies with other routers.) Adjacencies determine the distribution of routing protocol packets: routing protocol packets are sent and received only on adjacencies, and topological database updates are sent only along adjacencies. When adjacencies have been established, pairs of adjacent routers synchronize their topological databases.

A router sends LSA packets to advertise its state periodically and when the router's state changes. These packets include information about the router's adjacencies, which allows detection of nonoperational routers.

Using a reliable algorithm, the router floods LSAs throughout the area, which ensures that all routers in an area have exactly the same topological database. Each router uses the information in its topological database to calculate a shortest-path tree, with itself as the root. The router then uses this tree to route network traffic.

The description of the SPF algorithm up to this point has explained how the algorithm works within a single area (*intra-area routing*). For internal routers to be able to route to destinations outside the area (*interarea routing*), the area border routers must inject additional routing information into the area. Because the area border routers are connected to the backbone, they have access to complete topological data about the backbone. They use this information to calculate paths to all destinations outside its area and then advertise these paths to the area's internal routers.

AS boundary routers flood information about external ASs throughout the AS, except to stub areas. Area border routers are responsible for advertising the paths to all AS boundary routers.

OSPF Packets

This section contains the following topics:

OSPF Packet Header on page 295

Hello Packets on page 296

Database Description Packets on page 296

Link-State Request Packets on page 296

Link-State Update Packets on page 297

Link-State Acknowledgment Packets on page 297

There also are several types of link-state advertisement packets, which are discussed in “Link-State Advertisement Packet Types” on page 297.

OSPF Packet Header

All OSPF packets have a common 24-byte header that contains all information necessary to determine whether OSPF should accept the packet. The header consists of the following fields:

Version number—The current OSPF version number. This can be either 2 or 3.

Type—Type of OSPF packet.

Packet length—Length of the packet, in bytes, including the header.

Router ID—IP address of the router from which the packet originated.

Area ID—Identifier of the area in which the packet is traveling. Each OSPF packet is associated with a single area. Packets traveling over a virtual link are labeled with the backbone area ID, 0.0.0.0. You configure the area ID with the area statements.

Checksum—Fletcher checksum.

Authentication type—Authentication scheme to use for the packet. You configure the authentication type with the authentication-type statement. This is valid for OSPFv2 only.

Authentication—The authentication information itself. This is valid for OSPFv2 only.

Hello Packets

Routers periodically send hello packets on all interfaces, including virtual links, to establish and maintain neighbor relationships. Hello packets are multicast on physical networks that have a multicast or broadcast capability, which enables dynamic discovery of neighboring routers. (On nonbroadcast networks, dynamic neighbor discovery is not possible, so you must configure all neighbors statically using the neighbor statement.)

Hello packets consist of the OSPF header plus the following fields:

Network mask—Network mask associated with the interface.

Hello interval—How often the router sends hello packets. All routers on a shared network must use the same hello interval. You configure this interval with the hello-interval statement.

Options—Optional capabilities of the router.

Router priority—The router's priority to become the designated router. You can configure this value with the priority statement.

Router dead interval—How long the router waits without receiving any OSPF packets from a router before declaring that router to be down. All routers on a shared network must use the same router dead interval. You can configure this value with the dead-interval statement.

Designated router—IP address of the designated router.

Backup designated router—IP address of the backup designated router.

Neighbor—IP addresses of the routers from which valid hello packets have been received within the time specified by the router dead interval.

Database Description Packets

When initializing an adjacency, OSPF exchanges database description packets, which describe the contents of the topological database. These packets consist of the OSPF header, packet sequence number, and the link-state advertisement's header.

Link-State Request Packets

When a router detects that portions of its topological database are out of date, it sends a link-state request packet to a neighbor requesting a precise instance of the database. These packets consist of the OSPF header plus fields that uniquely identify the database information that the router is seeking.

Link-State Update Packets

Link-state update packets carry one or more link-state advertisements one hop farther from their origin. The router multicasts (floods) these packets on physical networks that support multicast or broadcast mode. The router acknowledges all link-state update packets and, if retransmission is necessary, sends the retransmitted advertisements unicast.

Link-state update packets consist of the OSPF header plus the following fields:

Number of advertisements—Number of link-state advertisements included in this packet.

Link-state advertisements—The link-state advertisements themselves.

Link-State Acknowledgment Packets

The router sends link-state acknowledgment packets in response to link-state update packets to verify that the update packets have been received successfully. A single acknowledgment packet can include responses to multiple update packets.

Link-state acknowledgment packets consist of the OSPF header plus the link-state advertisement header.

Link-State Advertisement Packet Types

Link-state request, link-state update, and link-state acknowledgment packets are used to reliably flood link-state advertisement packets. OSPF sends the following types of link-state advertisements:

Router link advertisements—Are sent by all routers to describe the state and cost of the router's links to the area. These link-state advertisements are flooded throughout a single area only.

Network link advertisements—Are sent by designated routers to describe all the routers attached to the network. These link-state advertisements are flooded throughout a single area only.

Summary link advertisements—Are sent by area border routers to describe the routes that they know about in other areas. There are two types of summary link advertisements: those used when the destination is an IP network, and those used when the destination is an AS boundary router. Summary link advertisements describe interarea routes; that is, routes to destinations outside the area but within the AS. These link-state advertisements are flooded throughout the advertisement's associated areas.

AS external link advertisement—Are sent by AS boundary routers to describe external routes that they know about. These link-state advertisements are flooded throughout the AS (except for stub areas).

Each link-state advertisement type describes a portion of the OSPF routing domain. All link-state advertisements are flooded throughout the AS.

Each link-state advertisement packet begins with a common 20-byte header.

External Metrics

When OSPF exports route information from external ASs, it includes a cost, or *external metric*, in the route. There are two types of external metrics: Type 1 and Type 2. Type 1 external metrics are equivalent to the link-state metric; that is, the cost of the route used in the internal AS. Type 2 external metrics are greater than the cost of any path internal to the AS.

Designated Router

Each multiaccess network has a designated router, which performs two main functions:

- Originate network link advertisements on behalf of the network.

- Establish adjacencies with all routers on the network, thus participating in the synchronizing of the link-state databases.

The OSPF hello protocol elects a designated router for the network based on the priorities advertised by all the routers. In general, when an interface first becomes functional, it checks whether the network currently has a designated router. If there is one, the router accepts that designated router regardless of its own router priority. Otherwise, if the router has the highest priority on the network, it becomes the designated router. If router priorities tie, the router with the highest router ID (which is typically the router's IP address) is chosen as the designated router.

OSPF Extensions to Support Traffic Engineering

To help provide traffic engineering and MPLS with information about network topology and loading, extensions have been added to the JUNOS implementation of OSPF. Specifically, OSPF generates opaque LSAs, which carry traffic engineering parameters. These parameters are used to populate the Traffic Engineering Database (TED), which is used by the Constrained Shortest Path First (CSPF) algorithm to compute the paths that MPLS LSPs will take. This path information is used by RSVP to set up LSPs and reserve bandwidth for them.

Configuring OSPF IGP Shortcuts

In OSPF, you can configure shortcuts, which allow OSPF to use an LSP as the next hop as if it were a logical interface from the ingress router to the egress router. The address specified on the `to` statement at the `[edit protocols mpls label-switched-path lsp-path-name]` hierarchy level must match the router ID of the egress router for the LSP to function as a direct link to the egress router and to be used as input to OSPF SPF calculations. When used in this way, LSPs are no different than Asynchronous Transfer Mode (ATM) and Frame Relay virtual circuits (VCs), except that LSPs carry only IPv4 traffic.



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