


7 The Open Shortest Path First (OSPF) Protocol

Open Shortest Path First routing (OSPF) is a shortest path first (SPF) or link-state protocol. OSPF is an interior gateway protocol (IGP) that distributes routing information between routers in a single autonomous system (AS). OSPF chooses the least-cost path as the best path. OSPF is suitable for complex networks with a large number of routers because it provides equal-cost multi-path routing where packets to a single destination can be sent to more than one interface simultaneously.

Each participating router distributes its local state (i.e., the router's usable interfaces and reachable neighbors) throughout the AS by flooding. In a link-state protocol, each router maintains a database describing the entire AS topology. This database is built from the collected link state advertisements of all routers. Each multi-access network that has at least two attached routers has a designated router and a backup designated router. The designated router floods a link state advertisement for the multi-access network and has other special responsibilities.

When a router starts, it uses the OSPF Hello Protocol to acquire neighbors. The router sends Hello packets to its neighbors, and in turn receives their Hello packets. On broadcast and point-to-point networks, the router dynamically detects its neighboring routers by sending Hello packets to a multicast address. On nonbroadcast networks, some configuration information is necessary in order to discover neighbors. On all multi-access networks (broadcast or nonbroadcast), the Hello Protocol also elects a Designated router for the network.



OSPF Hello Protocol

The router will attempt to form adjacencies with some of its newly acquired neighbors. Topological databases are synchronized between pairs of adjacent routers. On multi-access networks, the designated router determines which routers should become adjacent.

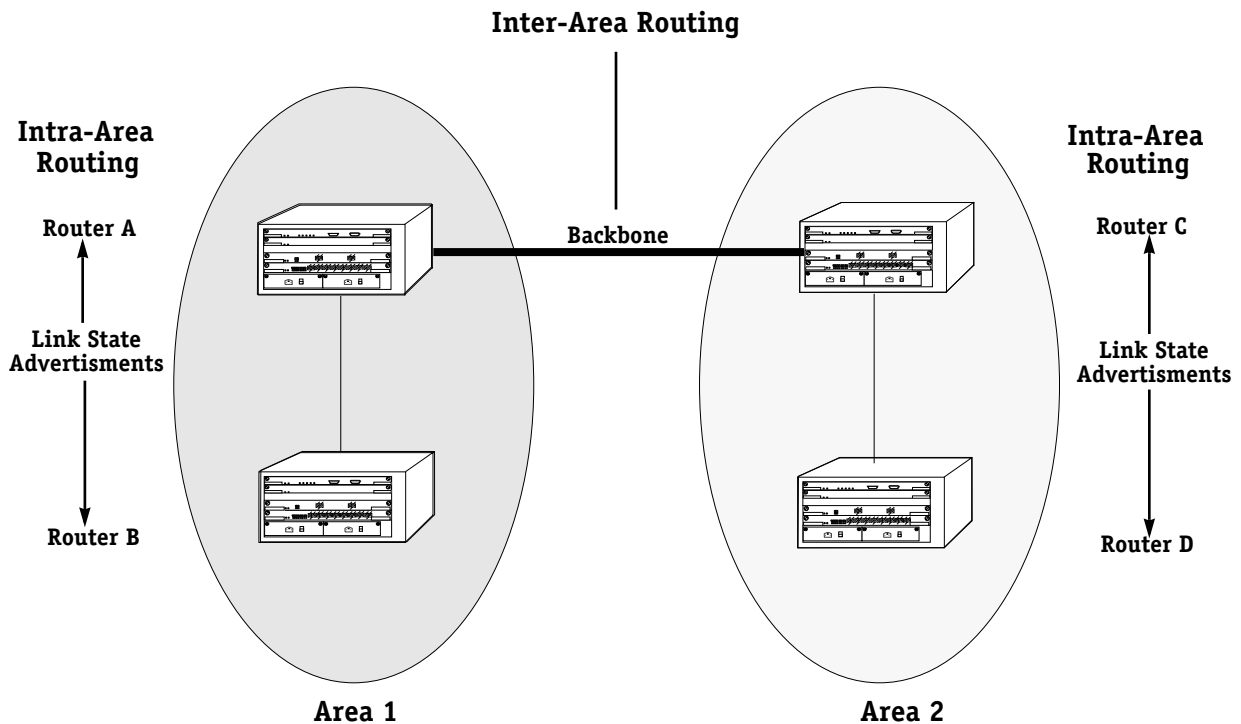
Adjacencies control the distribution of routing protocol packets. Routing protocol packets are sent and received only on adjacencies. In particular, distribution of topological database updates proceeds along adjacencies.

A router periodically advertises its state, which is called its *link state*. Link state is also advertised when a router's state changes. A router's adjacencies are reflected in the contents of its link state advertisements. This relationship between adjacencies and link state allows the protocol to detect downed routers in a timely fashion.

Link state advertisements are flooded throughout the AS. The flooding algorithm ensures that all routers have exactly the same topological database. This database consists of the collection of link state advertisements received from each router belonging to the area. From this database each router calculates a shortest-path tree, with itself as root. This shortest-path tree in turn yields a routing table for the protocol.

OSPF Areas

OSPF allows collections of contiguous networks and hosts to be grouped together. A group, together with the routers having interfaces to any one of the included networks, is called an *area*. Each area runs a separate copy of the basic link-state routing algorithm. This means that each area has its own topological database, as explained in the previous section.



OSPF Intra-Area and Inter-Area Routing

An area's topology is visible only to the members of the area. Conversely, routers internal to a given area know nothing of the detailed topology external to the area. This isolation of knowledge enables the protocol to reduce routing traffic by concentrating on small areas of an AS, as compared to treating the entire AS as a single link-state domain.

With the introduction of areas, it is no longer true that all routers in the AS have an identical topological database. A router actually has a separate topological database for each area it is connected to. (Routers connected to multiple areas are called *area border routers*). Two routers belonging to the same area have identical area topological databases.

Different areas communicate with each other through a *backbone*. The backbone consists of routers with contacts between multiple areas. A backbone must be contiguous (i.e., it must be linked to all areas).

The backbone is responsible for distributing routing information between areas. The backbone itself has all of the properties of an area. The topology of the backbone is invisible to each of the areas, while the backbone itself knows nothing of the topology of the areas.

For information on specifying an area or a backbone, see *OSPF Backbone and Areas* on page 7-18.

Routing in the Autonomous System takes place on two levels, depending on whether the source and destination of a packet reside in the same area (intra-area routing is used) or different areas (inter-area routing is used). In intra-area routing, the packet is routed solely on information obtained within the area. This protects intra-area routing from the injection of bad routing information from other sources outside the area.

All routers in an area must agree on that area's parameters. Since a separate copy of the link-state algorithm is run in each area, most configuration parameters are defined on a per-area basis. All routers belonging to an area must agree on that area's configuration. Misconfiguration will keep neighbors from forming adjacencies between themselves, and routing information could be wrong or confusing for the network.

Classification of Routers


When an AS is split into OSPF areas, the routers are further divided according to function into the following four overlapping categories:

- **Internal routers.** A router with all directly connected networks belonging to the same area. Routers with only backbone interfaces also belong to this category. These routers run a single copy of the basic routing algorithm.
- **Area border routers.** A router that attaches to multiple areas. Area border routers run multiple copies of the basic algorithm, one copy for each attached area and an additional copy for the backbone. Area border routers condense the topological information of their attached areas for distribution to the backbone. The backbone in turn distributes the information to the other areas.
- **Backbone routers.** A router that has an interface to the backbone. This includes all routers that interface to more than one area (i.e., area border routers). However, backbone routers do not have to be area border routers. Routers with all interfaces connected to the backbone are considered to be internal routers.
- **AS boundary routers.** A router that exchanges routing information with routers belonging to other Autonomous Systems. Such a router has AS external routes that are advertised throughout the Autonomous System. The path to each AS boundary router is known by every router in the AS. This classification is completely independent of the previous classifications: AS boundary routers may be internal or area border routers, and may or may not participate in the backbone.

Virtual Links

It is possible to define areas in such a way that the backbone is no longer contiguous. In this case the system administrator must restore backbone connectivity by configuring *virtual links*.

Virtual links can be configured between any two backbone routers that have an interface to a common non-backbone area. Virtual links belong to the backbone. The protocol treats two routers joined by a virtual link as if they were connected by an unnumbered point-to-point network. The routing protocol traffic that flows along the virtual link uses intra-area routing only.



OSPF Areas connected with a Virtual Link


For information on configuring virtual links, see *OSPF Statement Example* on page 7-26.

Stub Areas

OSPF allows certain areas to be configured as *stub areas*. A stub area is an area that has no AS external routes, and AS external advertisements are not flooded into or throughout stub areas.

In order to take advantage of the OSPF stub area support, default routing must be used in the stub area. This is accomplished by configuring one or more of the stub area's border routers to advertise a default route into the stub area. The default routes will match any destination that is not explicitly reachable by an intra-area or inter-area path (i.e., AS external destinations).

An area can be configured as stub when there is a single exit point from the area, as shown.



OSPF Stub Area

Area 1 and Area 3 could be configured as stub areas by configuring Area 1 and Area 3's border routers. Stub areas are configured using the OSPF statement in the **gated.conf** file. For more specific information, see *OSPF Backbone and Areas* on page 7-18.

The OSPF protocol ensures that all routers belonging to an area agree on whether the area has been configured as a stub. This guarantees that no confusion will arise in the flooding of AS external advertisements.

Two restrictions on the use of stub areas are:

- Virtual links cannot be configured through stub areas.
- AS boundary routers cannot be placed internal to stub areas.

OSPF UI Commands

The switch's User Interface (UI) contains several commands for viewing OSPF configuration parameters. The Gated UI commands operate slightly differently from normal switch commands, as you must modify the root command with a parameter. For example, to see statistics for OSPF interfaces, enter the following:

ospf stat

To view the possible OSPF command parameters, enter **ospf** at the system prompt. The following menu appears:

```
parameters for ospf root command are:
stat      OSPF status
nl        List of OSPF neighbors
ifl       List of OSPF interfaces
if <IP addr> Stats on an OSPF interface
al        Information about areas
lsdb      Dump the OSPF Link State Database
```

The parameters are described in the following sections. OSPF is configured like other Gated routing protocols (i.e., through the **gated.conf** file command statements). These command statements are described in *The OSPF Statement* on page 7-14.

Displaying OSPF Status

To display general information about OSPF routing, enter the following command:

ospf stat

A screen similar to the following displays:

```
***** OSPF Status *****

OSPF running with 3 neighbor(s), Router ID = 198.206.184.32

                                Transmit      Receive
                                =====
Hello Packets                   :          1414      5568
Database Description packets    :           8        13
Link State Requests             :          15         2
Link State Updates              :           0       172
Link State Acknowledgements    :          142      140
```

The first line of the display indicates that OSPF is running and that this switch has three neighbors. It also displays the router ID for this interface as configured in the **gated.conf** file (198.206.184.32).

Hello Packets. The number of HELLO packets transmitted and received on this switch. HELLO packets are used to announce a router's presence to other OSPF routers in the network.

Database Description packets. The number of database description packets transmitted and received on this switch. Database description packets contain the contents of the topological database.

Link State Requests. The number of link state requests transmitted and received on this switch. Link state requests are sent out to other switches and contain a complete picture of a switch's network topology.

Link State Updates. The number of link state update packets transmitted and received on this switch. Link state updates are sent out when a change occurs in a switch's link state database (e.g., when a new route is added).

Link State Acknowledgements. The number of link state acknowledgement packets transmitted and received by this switch. Link state acknowledgements are sent out by the switch when a request or update packet has been received.

Displaying OSPF Neighbors

A neighbor is any router that communicates through the OSPF protocol with the local switch. To display a list of the configured OSPF neighbors, enter the following command:

```
ospf nl
```

A screen similar to the following displays:

```
***** OSPF Neighbor List *****
```

4 neighbor(s) currently active

Area	IP Address	Rtr Id	Priority	State	(M/S)
0.0.0.1	198.206.184.10	192.206.184.10	1	Full	(M)
0.0.0.1	198.206.184.31	198.206.184.31	1	Full	(M)
0.0.0.1	198.206.184.40	198.206.184.40	1	Init	(M)
0.0.0.1	198.206.184.43	198.206.184.43	0	2way	(M)

This screen shows specific information about the switch's currently active neighbors.

Area. The area of the neighbor, identified in dotted decimal format. OSPF allows collections of contiguous networks and hosts to be grouped together. Such a group, together with the routers having interfaces to any one of the included networks, is called an area.

IP Address. The IP address of the neighbor interface.

Rtr Id. The Router ID for the neighbor. This number is specified in the **gated.conf** file in the definition statement. For more information on the router ID, see Chapter 4 of this manual.

Priority. The OSPF priority number assigned to the neighbor. Priority numbers are used to resolve the contention between multiple routers attached to a network for designated router status. The range of numbers is 0 to 255. Routers with a priority number of 0 are ineligible for designated router status; higher numbers indicate higher (i.e., more preferred) priorities.

State. The operating state of the neighbor. The following are the possible neighbor states:

Down	The initial state of a neighbor conversation. It indicates that there has been no recent information received from the neighbor. On nonbroadcast networks, Hello packets may still be sent to down neighbors, although at a reduced frequency.
Attempt	This state is only valid for neighbors attached to nonbroadcast networks. It indicates that no recent information has been received from the neighbor, but that a more concerted effort should be made to contact the neighbor. This is done by sending the neighbor Hello packets at intervals of HelloInterval (set in the gated.conf file).
Init	A Hello packet has recently been sent from the neighbor. However, bidirectional communication has not yet been established with the neighbor. All neighbors in this state (or higher) are listed in the Hello packets sent from the associated interface.

2Way	Communication between the two routers is bidirectional. This has been assured by the operation of the Hello Protocol. This is the most advanced state short of beginning adjacency establishment. The designated router is selected from the set of neighbors in 2-Way state or greater.
ExStart	The first step in creating an adjacency between the two neighboring routers. The goal of this step is to decide which router is the master, and to decide upon the initial database description sequence number. Neighbor conversations in this state or greater are called adjacencies.
Exchange	The router is describing its entire link state database by sending database description packets to the neighbor. Each database description packet has a database description sequence number, and is explicitly acknowledged. Only one database description packet is allowed outstanding at any one time. In this state, link state request packets may also be sent asking for the neighbor's more recent advertisements. All adjacencies in exchange state or greater are used by the flooding procedure. In fact, these adjacencies are fully capable of transmitting and receiving all types of OSPF routing protocol packets.
Loading	Link state request packets are sent to the neighbor asking for the more recent advertisements that have been discovered (but not yet received) in the exchange state.
Full	The neighboring routers are fully adjacent. These adjacencies will now appear in router links and network links advertisements.

(M/S). The master/slave status of the neighbor. OSPF requires adjacent routers to maintain synchronized topological databases. Using a procedure called the Database Exchange Process, database description packets are sent and received between two routers to maintain the required synchronization. The process requires the two routers to form a master/slave relationship. Database description packets sent by the master (polls) are acknowledged (with responses) by the slave through an echoing of the packet's sequence number. Both polls and their responses contain summaries of link state data. Only the master is allowed to retransmit database description packets, which it does at a definable, fixed interval (called the Retransmit Interval).

Displaying the OSPF Interface List

The OSPF software keeps track of all active interfaces in the network. To display a list of the interfaces currently configured to run OSPF, enter the following command:

```
ospf ifl
```

A screen similar to the following displays:

```

***** OSPF Interface List *****

1 OSPF interfaces currently active
DR = Designated Router, BDR = Backup Designated Router

  IP address    IF type    state    PRI    BDR    DR
=====
198.206.184.32 broadcast  DR        1    198.206.184.31  198.206.184.10

```

This screen lists all of the currently active OSPF interfaces. The first line indicates how many interfaces the switch is aware of in the network. In the example above, the switch sees one interface.

IP Address. The IP address of the router interface.

IF Type. The type of network to which the interface attaches. The possible values for this column are:

broadcast	A broadcast network.
non-broadcast	A nonbroadcast network, but still multi-access.
point to point	A point-to-point network.
virtual link	A virtual link between backbone routers.

Network types can be specified in the **gated.conf** file. See *OSPF Backbone and Areas* on page 7-18.

State. The functional level of this router interface. State determines whether full adjacencies are allowed to form over the interface. The possible values for the column are:

down	The initial interface state. In this state, the lower-level protocols have indicated that the interface is unusable. No protocol traffic will be sent or received on such an interface.
loopback	The router's interface to the network is looped back. The interface may be looped back in hardware or software. The interface will be unavailable for regular data traffic. However, it may still be desirable to gain information on the quality of this interface, either by sending ICMP pings to the interface or through an error test. For this reason, IP packets may still be addressed to an interface in Loopback state. To facilitate this, such interfaces are advertised in router links advertisements as single host routes, with a destination that is the IP interface address.

waiting	In this state, the router is trying to determine the identity of the designated router for the network. To determine the identity, the router monitors the Hello packets it receives. The router is not allowed to elect a backup designated router or a designated router until it transitions out of Waiting state. This prevents unnecessary changes of the backup or designated router.
P2P	The interface is operational, and connects either to a physical point-to-point network or to a virtual link. Upon entering this state, the router attempts to form an adjacency with the neighboring router. Hello packets are sent to the neighbor at the end of every HelloInterval seconds.
DR other	The interface is to a multi-access network on which another router has been selected to be the designated router. In this state, the local router has not been selected as a backup designated router either. The router forms adjacencies to both the designated router and the backup designated router (if they exist).
BDR	In this state, the local router is the backup designated router on the attached network. It will be promoted to designated router if the present designated router fails. The router establishes adjacencies to all other routers attached to the network. The backup designated router performs slightly different functions during the flooding procedure, compared with the designated router.
DR	In this state, the local router is the designated router on the attached network. Adjacencies are established to all other routers attached to the network. The router must also originate a network links advertisement for the network node. The advertisement will contain information on links to all routers (including the designated router itself) attached to the network.

Priority. The OSPF priority assigned to this interface. Priority numbers are used to resolve the contention between multiple routers attached to a network for designated router status. The range of numbers is 0 to 255. Routers with a priority number of 0 are ineligible for designated router status; higher numbers indicate higher (i.e., more preferred) priorities.

Backup Designated Router (BDR). The IP address of the backup designated router for this interface.

Designated Router (DR). The IP address of the designated router for this interface.

Displaying Details of a Specific OSPF Interface

To display details regarding a specific OSPF interface, enter the following command:

```
ospf if <ipAddress>
```

where **<ipAddress>** is the IP address of the specific interface. A screen similar to the following displays:

```
***** OSPF Interface Stats for 198.206.184.32 *****
```

Area Id	:	0.0.0.1	Transit Delay	:	1
If Type	:	broadcast	Retransmit Interval	:	5
Admin Status	:	enabled	Router Dead Interval	:	40
IF events	:	2	Poll Interval	:	0
Router Priority	:	1	Hello Interval	:	10

The information contained on this screen is derived from the GateD configuration file. The screen shows the state of the various interface parameters as set by the OSPF statement.

Area ID. The Area ID for this interface, identified in dotted decimal format.

IF Type. The kind of network to which the interface attaches. The possible values for this column are:

broadcast	A broadcast network.
non-broadcast	A nonbroadcast network, but still multi-access.
point to point	A point-to-point network.
virtual link	A virtual link between backbone routers.

Network types can be specified in the **gated.conf** file. See *OSPF Backbone and Areas* on page 7-18.

Admin Status. Describes whether the interface is enabled or disabled for OSPF.

IF Events. The number of times an interface has changed its state. You can display the state using the **ifl** subcommand.

Router Priority. The OSPF priority this neighbor uses. Priority numbers are used to resolve the contention between multiple routers attached to a network for designated router status. The range of numbers is 0 to 255. Routers with a priority number of 0 are ineligible for designated router status; higher numbers indicate higher (i.e., more preferred) priorities.

Transit Delay. The estimated number of seconds required to transmit a link state update over this interface. This parameter takes into account transmission and propagation delays and must be greater than 0.

Retransmit Interval. The number of seconds that lapse between retransmissions of Link State Advertisements between the adjacent routers on this interface.

Router Dead Interval. If the router does not receive a Hello packet from a neighbor in the number of seconds shown in this field, then the neighbor is declared down.

Poll Interval. Before adjacency is established with a neighbor, or after a neighboring router has become inactive, OSPF packets are sent periodically at the interval specified here. This parameter is only valid for nonbroadcast interfaces.

Hello Interval. The length of time, in seconds, between Hello packets sent on the interface.

Displaying the OSPF Area List

OSPF allows you to subdivide a network into specific areas. To display a list of the configured OSPF areas, enter the following command:

ospf al

A screen similar to the following displays:

```
***** Area List *****

ASBDRs = Autonomous System Border Routers
ABDRs = Area Border Routes

  Area  Transit/Stub  SPF runs  ASBDRs  ABDRs  LSA count  Full neighbors
=====
0.0.0.0   Neither        3         1       0       4          1
0.0.0.1   Neither        7         1       0       4          1
```

Area. The Area ID for this interface, identified in dotted decimal format.

Transit/Stub. Indicates whether this is a Transit area, a Stub area, or neither. Transit areas can have traffic not destined for, or originating from, the area. Stub areas have only one exit path. See *Stub Areas* on page 7-5 for more information on Stub areas.

SPF runs. The number of times the Shortest Path First (SPF) algorithm has been run.

Autonomous System Border Routers (ASBDRs). The number of Autonomous System Border Routers (ASBDR) detected in this area. ASBDRs are routers that exchange information with routers from another autonomous system (AS).

Area Border Routers (ABDRs). The number of Area Border Routers (ABDR) detected in this area. ABDRs are routers with attachments to more than one area.

Link State Advertisement (LSA) Count. The number of link state advertisements in the area's link state database, excluding external LSAs.

Full Neighbors. The number of neighbors in the full connection state with the local router.

Displaying the OSPF Link State Database

The link state database is a record of the switch's learned and configured links in the network. To display the OSPF link state database, enter the following command:

```
ospf lsdb
```

A screen similar to the following displays:

***** Link State Database *****					
Area	Type	LSID	Router ID	Sequence	Age
0.0.0.1	Stub	192.168.10.0	198.206.184.40	0	0
0.0.0.1	Stub	192.206.184.0	198.206.184.40	0	0
0.0.0.1	routerLink	193.206.184.11	193.206.184.11	2147483540	3
0.0.0.1	routerLink	192.206.184.10	192.206.184.10	2147483582	1
0.0.0.1	routerLink	198.206.184.31	198.206.184.31	2147483594	1
0.0.0.1	routerLink	198.206.184.32	198.206.184.32	2147483579	0
0.0.0.1	routerLink	198.206.184.40	198.206.184.40	2147483626	9
0.0.0.1	routerLink	198.206.184.43	198.206.184.43	2147483506	10
0.0.0.1	networkLink	193.206.184.11	193.206.184.11	2147483597	3
0.0.0.1	networkLink	198.206.184.10	192.206.184.10	2147483568	1
0.0.0.1	networkLink	198.206.184.31	198.206.184.31	2147483647	1

Area. The Area ID for this link state advertisement, identified in dotted decimal format.

Type. Indicates the type of the link state advertisement. The possible values for this column are:

routerLink	Router link advertisements.
networkLink	Network link advertisements.
summaryLink	Summary link advertisements of routes to networks.
asSummaryLink	Summary link advertisements of routes to AS boundary routers.
asExternalLink	AS external link advertisements.
Stub	Stub area address advertisements.

Link State ID (LSID). Identifies the piece of the routing domain that is being described by the advertisement. For Type 1 link state advertisements, it is the originating router's RouterID. For Type 2, it is the IP interface address of the network's Designated Router. For Type 3, it is the destination network's IP address. For Type 4, it is the RouterID of the described AS boundary router. For Type 5, it is the destination network's IP address.

Router ID. The Router ID of the switch in the area.

Sequence. The database description sequence number established after the router leaves the **ExStart** state. The number is incremented as database description packets are exchanged. Used by OSPF to detect, and subsequently correct, old and duplicate link state advertisements.

Age. The age, in seconds, of the link state advertisement. Used by OSPF to limit the time LSAs can remain in the link state database.

The OSPF Statement

The following pages show the full OSPF command statement. This statement is entered in the **gated.conf** file. Each part of the statement is described in the sections following this full statement. Commands must be entered in the order shown:

```
ospf yes | no | on | off [ {
  defaults {
    preference preference;
    cost cost;
    tag [ as ] tag;
    type 1 | 2;
    inherit-metric;
  };
  exportlimit routes;
  exportinterval time;
  traceoptions trace_options;
  syslog [ first pktcnt ] [ every every_pktcnt ];
  backbone | ( area area ) {
    stub [ cost cost ];
    networks {
      network [ restrict ];
      network mask mask [ restrict ];
      network masklen number [ restrict ];
      host host [ restrict ];
    };
    stubhosts {
      host cost cost;
    };
  };
  interface interface_list, [ cost cost ] [ {
    enable | disable;
    passive;
    retransmitinterval time;
    transitdelay time;
    priority priority;
    hellointerval time;
    routerdeadinterval time;
    auth [ none | simple ] auth_key;
    auth md5 key auth_key id id-number [ {
      start-generate date-time;
      stop-generate date-time;
      start-accept date-time;
      stop-accept date-time;
    } ];
    secondary auth [ none | simple ] auth_key;
  } ];
};
```

```

interface interface_list nonbroadcast [ cost cost ] [ {
    pollinterval time;
    routers {
        gateway [ eligible ];
    };
    enable | disable ;
    passive ;
    retransmitinterval time;
    transitdelay time;
    priority priority;
    hellointerval time;
    routerdeadinterval time;
    auth [none | simple] auth_key;
    auth md5 key auth_key id id-number [ {
        start-generate date-time;
        stop-generate date-time;
        start-accept date-time;
        stop-accept date-time;
    }];
    secondary auth [none | simple] auth_key;
};
};
};

```

If the router is specified in the above statement as a backbone, then the following extra statement can be added to the end of the OSPF statement in the configuration file. This statement allows you to set up virtual links:

```

virtuallink neighborid router_id transitarea area {
    enable | disable ;
    passive ;
    retransmitinterval time;
    transitdelay time;
    priority priority;
    hellointerval time;
    routerdeadinterval time;
    auth [none | simple] auth_key;
    auth md5 key auth_key id id-number [ {
        start-generate date-time;
        stop-generate date-time;
        start-accept date-time;
        stop-accept date-time;
    }];
    secondary auth [none | simple ] auth_key;
};

```

The OSPF statement is described in the following three parts:

- **OSPF General Statements.** These statements affect OSPF operation for the entire switch, and are covered in *OSPF General Statements* on page 7-16.
- **OSPF Backbone and Areas.** These statements allow you to define and configure the backbone and areas for this switch, and for specific interfaces. These statements are covered in *OSPF Backbone and Areas* on page 7-18.
- **Using Authentication with OSPF.** These statements define the use of authentication with OSPF, and are described in *Using Authentication in OSPF* on page 7-23.

OSPF General Statements

OSPF general statements control the overall operation of OSPF on the entire switch. All general statements are optional, and should follow the statement that activates OSPF (**ospf on**). The following section describes the commands available in the OSPF general statements.

defaults

These parameters specify the defaults used when importing OSPF AS External (ASE) routes into the GateD routing table, and when exporting routes from the GateD routing table into OSPF ASEs.

preference	The preference is used to determine how OSPF routes compete with routes from other protocols in the GateD routing table. The default value is 150.
cost <i>cost</i>	The cost used when exporting a non-OSPF route from the GateD routing table into OSPF as an ASE. The default value is 1. This may be explicitly overridden in an export policy.
tag [as] <i>tag</i>	OSPF ASE routes have a 32-bit tag field that is not used by the OSPF protocol, but may be used by an export policy to filter routes. When OSPF is interacting with an EGP, the tag field may be used to propagate AS path information, in which case the as keyword is specified and the tag is limited to 12 bits of information. If not specified, the tag is set to 0.
type 1 2	Routes exported from the GateD routing table into OSPF default to type 1 ASEs. This default may be explicitly changed here and overridden in an export policy. By setting the type to 2, routes exported from the GateD routing table into OSPF are labeled type 2 .
inherit-metric	This command allows an OSPF ASE route to inherit the metric of the external route when no metric is specified on the export. This option maintains compatibility with all the current export functions. A metric specified on the export will take precedence. The cost specified in the default will be used if inherit-metric is not specified.

exportinterval *time*

This command specifies how often a batch of ASE link state advertisements will be generated and flooded into OSPF. The default is 1 time per second.

exportlimit *routes*

This command specifies how many ASEs will be generated and flooded in each batch. The default is 100.

traceoptions *trace_options*

This command specifies the tracing options for OSPF. Tracing options specifically for OSPF are:

lsabuild	Create the link state advertisement.
lsatransmit (or lsatx)	Trace the link state packets transmitted.
lsareceive (or lsarx)	Trace the link state packets received.
spf	Trace the Shortest Path First (SPF) calculations.
debug	Trace OSPF at the debugging level of detail.

Packet tracing options (which may be modified with **detail**, **send** and **recv**) include:

hello	OSPF Hello packets are used to determine neighbor reachability.
dd	OSPF database description (DD) packets are used in synchronizing OSPF databases.
request	OSPF link state request packets are used in synchronizing OSPF databases.
lsu	OSPF link state update packets are used in synchronizing OSPF databases.
ack	OSPF link state acknowledge packets are used in synchronizing OSPF databases.

For specific information on tracing options and the trace statement, see Chapter 12 of this manual.

syslog [**first** *pckcnt*] [**every** *every_pckcnt*]

This command specifies the tracing options for logging OSPF packets. If only **syslog** is specified, then every OSPF packet received will generate a report on the console screen. If the first and every parameters are specified, then the log will contain the first *pckcnt* packets for every type of OSPF. After the first *pckcnt* packets, the syslog will only save one message per every *every_pckcnt* packets.

For example, if **syslog** was specified as shown:

syslog first 10 every 30

then the first ten OSPF packets received would generate a report on the console screen. After the first ten, every thirtieth OSPF packet would generate a report.

OSPF Backbone and Areas

OSPF allows collections of contiguous networks and hosts to be grouped together. These groups, together with the routers having interfaces to any one of the included networks, is called an *area*. Each area runs a separate copy of the basic link-state routing algorithm. This means that each area has its own topological database.

Different areas communicate with each other through a *backbone*. The backbone consists of routers with contacts between multiple areas. A backbone must be contiguous (i.e., it must be linked to all areas).

The following is the backbone/area section of the OSPF statement:

```
backbone | ( area area ) {
    stub [ cost cost ];
    networks {
        network [ restrict ];
        network mask mask [ restrict ];
        network masklen number [ restrict ];
        host host [ restrict ];
    };
    stubhosts {
        host cost cost;
    };
    interface interface_list; [ cost cost ] [ {
        enable | disable ;
        passive ;
        retransmitinterval time;
        transitdelay time;
        priority priority;
        hellointerval time;
        routerdeadinterval time;
        auth [none | simple] auth_key;
        auth md5 key auth_key id id-number [ {
            start-generate date-time;
            stop-generate date-time;
            start-accept date-time;
            stop-accept date-time;
        } ];
        secondary auth [none | simple ] auth_key;
    } ];
};
```

```

interface interface_list nonbroadcast [ cost cost ] [ {
    pollinterval time;
    routers {
        gateway [ eligible ];
    };
    enable | disable ;
    passive ;
    retransmitinterval time;
    transitdelay time;
    priority priority;
    hellointerval time;
    routerdeadinterval time;
    auth [none | simple] auth_key;

    auth md5 key auth_keyid id-number[ {
        start-generate date-time;
        stop-generate date-time;
        start-accept date-time;
        stop-accept date-time;
    }];
    secondary_auth [none | simple ] auth_key;
};
};
};

```

If the router is specified in the above statement as a backbone, then the following extra statement can be added to the end of the OSPF statement in the configuration file. This statement allows you to set up virtual links:

```

virtuallink neighborid router_id transitarea area {
    enable | disable ;
    passive;
    retransmitinterval time;
    transitdelay time;
    priority priority;
    hellointerval time;
    routerdeadinterval time;
    auth [none | simple] auth_key;
    auth md5 key auth_keyid id-number[ {
        start-generate date-time;
        stop-generate date-time;
        start-accept date-time;
        stop-accept date-time;
    }];
    secondary_auth [none | simple ] auth_key;
};

```

Interfaces, depending on whether they are simple interfaces, point-to-point links, or nonbroadcast multi-access links, have different parameters that can be assigned to the interface. Parameters for the different types of interfaces are described in:

- *Common Interface Parameters* on page 7-22
- *Interface Parameters for Point-to-Point Links* on page 7-22
- *Interface Parameters for Non-Broadcast Multi-Access Links* on page 7-23

The following section describes the commands available in the OSPF backbone and area statements.

backbone | area *area*

Each OSPF router must be configured into at least one OSPF area. If more than one area is configured, a backbone must be configured. The backbone may only be configured using the backbone keyword. It may not be specified as area 0. The backbone interface may be a virtual link. The commands that modify this statement are:

stub [cost *cost*]

A stub area is one in which there are no ASE routes. If a cost is specified, this is used to inject a default route into the area with the specified cost. See *Stub Areas* on page 7-5 for more information on stub areas.

networks

The networks list describes the scope of an area. Intra-area link state advertisements (LSAs) that fall within the specified ranges are not advertised into other areas as inter-area routes. Instead, the specified ranges are advertised as summary network LSAs. If **restrict** is specified, the summary network LSAs are not advertised. Intra-area LSAs that do not fall into any range are also advertised as summary network LSAs. This option is very useful on well-designed networks in reducing the amount of routing information propagated between areas. The entries in this list are either networks, or a subnetwork and mask pair.

stubhosts

The stubhosts list specifies directly attached hosts that should be advertised as reachable from this router, and the cost with which they should be advertised. Point-to-point interfaces on which it is not desirable to run OSPF should be specified here. It is also useful to assign an additional address to the loopback interface (one not on the 127 network) and advertise it as a stubhost. If this address is the same one used as the **router-id**, it enables routing to OSPF routers by **router-id**, instead of by interface address. This is more reliable than routing to one of the router's interface addresses, which may not always be reachable.

interface [**cost** *cost*]

This form of the interface clause is used to configure a broadcast (which requires IP multi-cast support) or a point-to-point interface. See the section on interface list specification for the description of the *interface_list*. Each interface has a cost. The costs of all interfaces that a packet must cross to reach a destination are summed to get the cost to that destination. The default cost is 1, but another non-zero value may be specified.

interface nonbroadcast [**cost** *cost*]

This form of the interface clause is used to specify a nonbroadcast interface on a nonbroadcast multi-access (NBMA) medium. Since an OSPF broadcast medium must support IP multicasting, a broadcast-capable medium that does not support IP multicasting (such as Ethernet) must be configured as a nonbroadcast interface.

virtuallink neighborid *router_id* **transitarea** *area*

Virtual links are used to establish or increase connectivity of the backbone area. The **neighborid** is the **router_id** of the other end of the virtual link. The area specified with the **transitarea** parameter must also be configured on the system. All standard interface parameters defined *Common Interface Parameters* on page 7-22 can be specified on a virtual link.

Common Interface Parameters

Common interface parameters are used to modify interface statements in the backbone section of the OSPF statement. These parameters are placed after the interface portion of the OSPF statement, and can be used on any type of interface (general, point-to-point, and nonbroadcast multi-access). Interface parameters common to all types of interfaces are:

enable

Enables the interface. This is the default.

disable

Disable the interface.

passive

Do not send or receive OSPF packets on this interface. This has the affect of originating a stub link to the interface. For more on stub links, see *Stub Areas* on page 7-5.

retransmitinterval *time*

The number of seconds between link state advertisement retransmissions for adjacencies (same area routers) belonging to this interface.

transitdelay *time*

The estimated number of seconds required to transmit a link state update over this interface. This command takes into account transmission and propagation delays and must be greater than 0.

priority *priority*

A number between 0 and 255 specifying the priority for becoming the designated router on this interface. When two routers attached to a network both attempt to become the designated router, the one with the highest priority wins. A router whose router priority is set to 0 is ineligible to become the designated router. OSPF supports both NBMA and point-to-point interfaces. The priority for these interfaces must be manually configured to elect the designated router.

hellointerval *time*

The length of time, in seconds, between Hello packets that the router sends from the interface.

routerdeadinterval *time*

The number of seconds between router Hello packets before the router's neighbors declare it down.

Interface Parameters for Point-to-Point Links

The following parameter (along with the common interface parameters described above) is only valid for a point-to-point link. The interface parameters for point-to-point links are:

nomulticast

By default, OSPF packets to neighbors on point-to-point interfaces are sent using the IP multicast mechanism. If the use of IP multicasting is not desired because the remote neighbor does not support it, the **nomulticast** parameter may be specified to force the use of unicast OSPF packets.

Interface Parameters for Non-Broadcast Multi-Access Links

The following parameters (along with the common interface parameters described above) is only valid with a nonbroadcast multi-access link. The interface parameters for nonbroadcast multi-access links are:

pollinterval *time*

Before adjacency is established with a neighbor, OSPF packets are sent periodically at the specified **pollinterval** in seconds.

routers *routerid*

By definition, it is not possible to send broadcast packets to discover OSPF neighbors on a nonbroadcast medium. The list includes one or more neighbors and an indication of their eligibility to become a designated router.

Using Authentication in OSPF

By definition, all OSPF protocol exchanges are authenticated. Authentication can help to guarantee that routing information is only imported from trusted routers. A variety of authentication schemes can be used, but a single scheme must be configured for each interface. The use of different schemes enables some interfaces to use much stricter authentication than others. The three authentication schemes available are: none, simple, and MD5 authentication.

The commands used to activate authentication are as follows:

auth [**none** | **simple**] *auth_key*

OSPF uses authentication to generate and verify the authentication field in the OSPF header. The authentication is configured on a per-interface basis. It is specified by 1 to 8 decimal digits separated by periods (i.e., 1.2.3.4.5.6.7.8), a 1 to 8 byte hexadecimal string preceded by 0x (i.e., 0x12345678), or a 1 to 8 character string in double quotes (i.e., "eightcha").

secondary auth [**none** | **simple**] *auth_key*

Used by OSPF authentication to generate and verify the secondary authentication field in the OSPF header. The authentication key can be configured on a per-interface basis. It is specified by 1 to 8 decimal digits separated by periods (i.e., 1.2.3.4.5.6.7.8), a 1 to 8 byte hexadecimal string preceded by 0x (i.e., 0x12345678), or a 1 to 8 character string in double quotes (i.e., "eightcha").

auth md5 key *auth_key* **id** *id-number* [{

start-generate *date-time*;

stop-generate *date-time*;

start-accept *date-time*;

stop-accept *date-time*;

};]

Used by OSPF authentication to generate MD5 authentication. MD5 is an authentication method that uses encryption to prevent information from being accessed without the correct key.

The MD5 authentication parameters are as follows:

auth_key	A one to eight character text string.
id-number	An integer between 1-255.
start-generate	When sending authenticated material, it must originate after the time specified.
stop-generate	When sending authenticated material, it must originate before the time specified.
start-accept	When receiving authenticated material, it must originate after the time specified.
stop-accept	When receiving authenticated material, it must originate before the time specified.
<i>date-time</i>	A timestamp in the format YYYY/MM/DD HH:MM . All time fields are required.

Examples of these authentication schemes are described in the sections below.

No Authentication

When no authentication is required, use authentication type **none**. To use authentication type **none**, add the following line to the appropriate OSPF interface statements.

```
auth none "";
```

Note that even though no authentication is specified, you must still include a set of quotes.

Simple Authentication Key

When you wish to keep certain routers from exchanging OSPF packets, use the **simple** form of authentication. To specify authentication type as **simple**, add the following lines to your OSPF interface statements:

```
auth simple your-key;
```

where *your-key* is a simple text string.

MD5 Authentication

MD5 authentication can be used to encrypt information sent over the network. The system works by using shared secret keys listed in the configuration file. Keys are used to sign the packets with an MD5 checksum, and they cannot be forged or tampered with. Since the keys are not included in the packet, snooping the key is not possible.

GateD's MD5 authentication uses an authentication key of up to 16 characters, and allows multiple MD5 keys per interface. Each key has two associated time ranges.

For example, the following **auth md5** entry would enable MD5 encryption using a **key** of **test**, and an **id** of **5**:

```
auth md5 key "test" id 5 {
    start-generate 1999/8/7 10:45;
    stop-generate 1999/8/7 11:45;
    start-accept 1999/8/7 10:45;
    stop-accept 1999/8/7 11:45;
};
```

Authentication for transmitted packets would start at 10:45 AM on August 7 and end an hour later. The same is true for authentication packets received.

If no value is given for the time ranges, the default values are that an MD5 key is always generated, and an MD5 key is always accepted.

To specify multiple MD5 keys on an interface, add the following to the appropriate OSPF interface statements.

```
auth md5 {
    key auth_keyid id-number{
        start-generate date-time;
        stop-generate date-time;
        start-accept date-time;
        stop-accept date-time;
    };
    key auth_keyid id-number{
        start-generate date-time;
        stop-generate date-time;
        start-accept date-time;
        stop-accept date-time;
    };
    key auth_keyid id-number{
        start-generate date-time;
        stop-generate date-time;
        start-accept date-time;
        stop-accept date-time;
    };
};
```

OSPF Statement Example

The following is an example of an OSPF statement:


```
ospf yes
  defaults {
    preference 50 ;
    cost 10 ;
  };
  exportlimit 150 ;
  exportinterval 10 ;
  area 1 {
    interface 1.1.1.1 {
      priority 10 ;
      auth simple "test" ;
    };
    virtuellink neighborid 2.2.2.2 transitarea 1 {
      hellointerval 10 ;
      routerdeadinterval 50;
    };
  };
};
```

In the above statement:

- OSPF has been activated on this switch, with a default interface **preference** of **50**, and a default **cost** of **10**. This is added to the previously determined cost of the route based on other criteria.
- The **exportlimit** statement specifies that ASE link state advertisements are sent out in groups of **150** routes.
- The **exportinterval** statement specifies that ASE link state advertisements are broadcast every **10** seconds.
- The **area** statement specifies the router as being in Area **1**. Since interface 1.1.1.1 is configured into Area 1, it has a priority of **10**, and requires simple authentication with a key of **test**.
- This router has been configured with a virtual link to the router with an ID of 2.2.2.2. (A router identification is specified in the **gated.conf** file using the definition statement. See Chapter 4 for specific details.) The **transitarea** (i.e., the configured area that the virtual link passes through) is area **1**.
- The **hellointerval** statement specifies that Hello packets are sent on this link every **10** seconds.
- The **routerdeadinterval** statement specifies that after **15** seconds, if no packets are received from a neighbor, then the neighbor is declared down.

Application Example

The following diagram illustrates a simple AS using OSPF:



In this network, Switch 1 and Switch 2 are in Area 1, while Switch 3 and Switch 4 are in Area 2. Switch 2 and Switch 3 create the backbone for the areas. A **gated.conf** file for each switch is listed below.

Switch 1

```
router-id 208.8.1.1;
ospf yes {
  area 1 {
    interface 208.8.1.1 {
      priority 1;
    };
  };
};
```

- Switch 1 has a router ID of **208.8.1.1** (this is the same as the interface IP address). The router ID is a way for OSPF to identify individual switches in the network.
- OSPF is activated with the **ospf yes** command.
- This switch is assigned to Area 1 with the **area** command.
- Interface **208.8.1.1** is defined.
- Interface **208.8.1.1** is given a priority of **1**. This means that when a new designated router must be elected, most likely this router will become the designated router (lower values are considered to be better candidates for designated router).

Switch 2

```
router-id 208.8.2.2;
ospf yes {
  backbone {
    interface 208.8.2.2 {
      priority 5;
    };
  };
  area 1{
    interface 208.8.2.1 {
      priority 5;
    };
  };
};
```

- Switch 2 has a router ID of **208.8.2.2** (this is the same as the backbone interface IP address). The router ID is a way for OSPF to identify individual switches in the network.
- OSPF is activated with the **ospf yes** command.
- This switch is assigned as the backbone, using the **backbone** command. Interface 208.8.2.2 is the backbone interface.
- Interface **208.8.2.2** is defined.
- Interface **208.8.2.2** is given a priority of **5**.
- This switch is also assigned to Area 1 using the **area** command. Interface **208.8.2.1** is the Area 1 interface.
- Interface **208.8.2.1** is given a priority of **5**.

Switch 3

```
router-id 208.8.3.2;
ospf yes {
  backbone {
    interface 208.8.3.2 {
    };
  };
  area 2 {
    interface 208.8.3.1
  };
};
```

- Switch 3 has a router ID of **208.8.3.2** (this is the same as the backbone interface IP address). The router ID is a way for OSPF to identify individual switches in the network.
- OSPF is activated with the **ospf yes** command.
- This switch is assigned as the backbone, using the **backbone** command. Interface **208.8.3.2** is the backbone interface.
- Interface **208.8.3.2** is defined.
- Interface **208.8.3.2** is given a priority of **5**.
- This switch is also assigned to Area 2 using the **area** command. Interface **208.8.3.1** is the Area 1 interface.
- Interface **208.8.3.1** is given a priority of **5**.

Switch 4

```
router-id 208.8.4.1;
ospf yes {
  area 2 {
    interface 208.8.4.1 {
      priority 4;
    };
  };
};
```

- Switch 4 has a router ID of **208.8.4.1** (this is the same as the interface IP address). The router ID is a way for OSPF to identify individual switches in the network.
- OSPF is activated with the **ospf yes** command.
- This switch is assigned to Area 2 with the **area** command.
- Interface **208.8.4.1** is defined
- Interface **208.8.4.1** is given a priority of **4**.

