Teldat Router

OSPF Protocol

Doc. *DM514-I* Rev. *8.00*

July, 1999

INDEX

Chapter 1 Introduction 1

1. The OSPF Protocol 2
2. The OSPF Routing Protocol 3
3. Configuring OSPF 4
   - 3.1. Enabling the OSPF Protocol 4
   - 3.2. Defining Backbone and Attached OSPF Areas 5
   - 3.3. Setting OSPF Interfaces 6
   - 3.4. OSPF Routing Summary 7
     - *a) Designated Router 7*
   - 3.5. Setting Non-Broadcast Network Parameters 8
   - 3.6. Enabling AS Boundary Routing 9
   - 3.7. Other Configuration Tasks 10
     - *a) Setting OSPF Router IDS 10*
     - *b) Setting Virtual Links 10*
     - *c) Configuring for Routing Protocol Comparisons 10*

Chapter 2 Configuration 12

1. Configuration Commands 13
   - 1.1. ? (HELP) 13
   - 1.2. ADD 14
     - *a) ADD NEIGHBOR 14*
     - *b) ADD RANGE <area#> <IP-address> <IP-address-mask> 14*
   - 1.3. DELETE 15
     - *a) DELETE AREA <area#> 15*
     - *b) DELETE INTERFACE <interface-IP-address> 15*
     - *c) DELETE NEIGHBOR 15*
     - *d) DELETE NON-BROADCAST <interface-IP-address> 16*
     - *e) DELETE RANGE <area#> <IP-address> 16*
     - *f) DELETE VIRTUAL-LINK 16*
   - 1.4. DISABLE 16
     - *a) AS boundary routing 16*
     - *b) OSPF routing protocol 17*
   - 1.5. ENABLE 17
     - *a) ENABLE AS boundary routing 17*
     - *b) ENABLE OSPF routing protocol 18*
   - 1.6. LIST 18
     - *a) LIST ALL 19*
     - *b) LIST AREAS 20*
     - *c) LIST INTERFACES 21*
     - *d) LIST NEIGHBORS 21*
     - *e) LIST NON-BROADCAST 22*
     - *f) LIST VIRTUAL-LINKS 22*
   - 1.7. SET 22
     - *a) SET AREA 23*
     - *b) SET COMPARISON 23*
     - *c) SET INTERFACE 24*
     - *d) SET NON-BROADCAST 24*
     - *e) SET VIRTUAL-LINK 25*
   - 1.8. EXIT 25

Chapter 3 Monitoring 26
1. Monitoring Commands 27
1.1. ? (HELP) 27
1.2. ADVERTISEMENT expansion 28
1.3. AREA summary 30
1.4. AS-EXTERNAL advertisements 31
1.5. DATABASE summary 32
1.6. DUMP routing tables 33
1.7. INTERFACE summary 34
1.8. NEIGHBOR summary 36
1.9. PING address 38
1.10. ROUTERS 38
1.11. SIZE 39
1.12. STATISTICS 40
1.13. TRACEROUTE address 41
1.14. WEIGHT 42
1.15. EXIT 42

Chapter 1
Introduction

1. The OSPF Protocol

This chapter describes how to use the Open Shortest Path First (OSPF) Protocol, which is an Interior Gateway Protocol (IGP). The **Teldat Router** supports two different IGPs for building the IP routing table, Open Shortest Path First (OSPF) Protocol, and RIP Protocol.

OSPF is based on link-state technology or the shortest-path first (SPF) algorithm. RIP is based on the Bellman-Ford or the distance-vector algorithm. The information is organized within the following sections:

- The OSPF Routing Protocol.
- Configuring the OSPF Protocol.
- OSPF Configuration Commands.
- OSPF Monitoring Commands.

Routers that use a common routing protocol form an Autonomous System (AS). This common routing protocol is called an Interior Gateway Protocol (IGP). IGPs dynamically detect network reachability and routing information within an AS and use this information to build the IP routing table. IGPs can also import external routing information into the AS.

The **Teldat Router** can simultaneously run OSPF and RIP. When doing so, OSPF routes are preferred. In general, use of the OSPF protocol is recommended due to its robustness, responsiveness, and decreased bandwidth requirements.

2. The OSPF Routing Protocol

The **Teldat Router** supports a complete implementation of the OSPF routing protocol, as specified in RFC 1247 (Version 2). This version is incompatible with bridging routers running OSPF version 1. OSPF information will not be exchanged between routers running version 1 and version 2.

OSPF is a link state dynamic routing protocol that detects and learns the best routes to (reachable) destinations. OSPF can quickly perceive changes in the topology of an AS, and after a short convergence period, calculate new routes. The OSPF protocol does not encapsulate IP packets, but forwards them based on destination address only.

OSPF is designed to provide services not available with RIP. OSPF features include the following:

- Least Cost Routing. Allows you to configure path costs based on any combination of network parameters. For example, bandwidth, delay, and cost.
- No limitations to the routing metric. While RIP restricts the routing metric to 16 hops, OSPF has no restriction.
- Multipath Routing. Allows you to use multiple paths of equal cost that connect the same points. You can then use these paths for load balancing resulting in more efficient use of network bandwidth.
- Area routing. Decreases the resources (memory and network bandwidth) consumed by the protocol and provides an additional level of routing protection.
- Variable Length Subnet Masks. Allow you to break an IP address into variable size subnets, conserving IP address space.

OSPF supports the following physical network types:

- Point-to-Point. Networks that use a communication line to join a single pair of routers. A 56 Kb serial line that connects two routers is an example of a point-to-point network.
- Broadcast. Networks that support more than two attached routers and are capable of addressing a single physical message to all attached routers. A Token Ring network is an example of a broadcast network.
- Non-Broadcast. Networks that support more than two attached routers but have no broadcast capabilities. An X.25 Public Data Network is an example of a non-broadcast network. For OSPF to function properly, this network requires extra configuration information about other OSPF routers attached to the non-broadcast network.

3. Configuring OSPF

The following steps outline the tasks required to get the OSPF protocol up and running. The sections that follow explain each step in detail, including examples.

Enable the OSPF protocol. In doing so, you must estimate the final size of the OSPF routing domain.

Define OSPF areas attached to the router. If no OSPF areas are defined, a single backbone area is assumed.

Define the router's OSPF network interfaces. The cost of sending a packet out each interface must be set, along with a collection of the OSPF operating parameters.

If the router interfaces to non-broadcast networks, you must also set the non-broadcast network parameters. This consists of a list of the other OSPF routers that are connected to the non-broadcast network.

If you want the router to import routes learned from other routing protocols (RIP or statically configured routes), you have to enable AS boundary routing. In addition, you must define whether routes are imported as Type 2 or Type 1 externals.

If you want to boot a neighboring router over an attached point-to-point interface, the neighbor's IP address must be configured. This is done by defining non-broadcast parameters for the point to point interface.

3.1. Enabling the OSPF Protocol

Before you can enable OSPF, you must supply the total number of external routes and the total number of OSPF routers. These values estimate the final size of the OSPF routing domain and should be identical in all of your OSPF routers. Each router running the OSPF protocol has a database describing a map of the routing domain. This database is identical in all participating routers. From this database the IP routing table is built through the construction of a shortest-path tree, with the router itself as root. The routing domain refers to an AS running the OSPF protocol.

To enable the OSPF protocol enter **ENABLE OSPF** and respond to the following prompts:

Example:

```
OSPF Config> ENABLE OSPF
Estimated # external routes[0]? 500
Estimated # OSPF routers[0]? 50
OSPF Config>
```

- *Estimated # external routes* - Total number of AS external routes that will be imported into the OSPF routing domain. A single destination may lead to multiple external routes when it is imported by separate AS boundary routers. For example if the OSPF routing domain has two AS boundary routers, both importing routes to the same 100 destinations, the number of AS external routes should be set to 200.
- *Estimated # OSPF routers* - Total number of OSPF routers in the routing domain.

3.2. Defining Backbone and Attached OSPF Areas

Define the OSPF areas that are directly attached to the router. If no areas are defined, the router software assumes that all the router's directly attached networks belong to the backbone area (area ID 0.0.0.0).

OSPF allows you to split the AS into regions called areas. OSPF areas are a collection of contiguous networks. The topology of any one area is hidden from that of the other areas. Hiding information significantly reduces routing traffic and protects routing within an area from outside influence.

A router has a separate database that contains the topology for each area to which it is connected. Two routers belonging to the same area have identical topologies for that area.

OSPF areas are defined as address ranges. External to the area, a single route is advertised for each address range. For example, if an OSPF area consisted of all subnets of the class B network 128.185.0.0, it would consist of a single address range. The address range would be specified as an address of 128.185.0.0 together with a mask of 255.255.0.0. Outside the area, the entire subnetted network would be advertised as a single route to network 128.185.0.0.

Every OSPF routing domain must have a backbone. The backbone is a special OSPF area having an area ID equal to 0.0.0.0. The OSPF backbone must be contiguous; however, it is possible to define areas where the backbone is not physically contiguous. When this situation exists, you must configure a virtual link to maintain the backbone's connectivity. You can configure a virtual link to maintain the backbone's connectivity. You can configure virtual links between any two backbone routers that have an interface to a common non-backbone area.

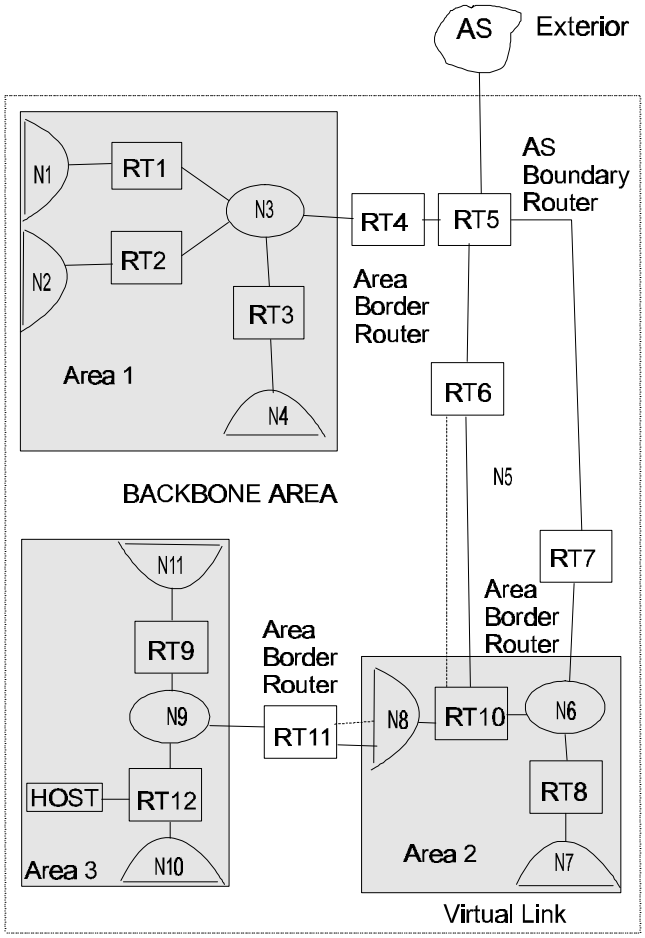

The backbone is responsible for distributing inter-area routing information. The backbone area consists of any of the following:

- Networks belonging to Area 0.0.0.0
- Routers attached to those networks
- Routers belonging to multiple areas
- Configured virtual links

To set the parameters for an OSPF area, use the **SET AREA** command and respond to the following prompts:

Example:

```
OSPF Config>SET AREA
Area number [0.0.0.0]?0.0.0.1
Authentication type [1]?1
Is this a stub area? (Yes or No): no
OSPF Config>
```

- *Area number* is the OSPF area address. An OSPF area is a contiguous group of networks that is defined by a list of address ranges, each indicated by a combination of the IP address and an address mask. A network belongs to an area if its address is in the list.
- *Authentication type* (security scheme) to be used in the area. The choices for authentication types are 1, which indicates a simple password; or 0, which indicates that no authentication is necessary exchange.
- *Stub area designation*. If you enter yes,

1. The area does not receive any AS external link advertisements, reducing the size of the area's OSPF database and decreasing memory usage for external routers in the stub area.
2. You cannot configure virtual links through a stub area.
3. You cannot configure a router within the stub area as an AS boundary routers.

NOTE: You cannot configure the backbone as a stub area. External routing in stub areas is based on a default route. Each border area router attaching to a stub area originates a default route for this purpose. The cost of this default route is also configurable in the OSPF SET AREA command.

3.3. Setting OSPF Interfaces

To set the OSPF parameters for the router's network interfaces, use the **SET INTERFACE** command, then respond to the prompts.

The cost of an OSPF interface can be dynamically changed from the router's monitoring environment. This new cost is flooded quickly throughout the OSPF routing domain, and modifies the routing immediately.

There are two special kinds of OSPF routers, area border routers and AS boundary routers.

- Area Border Routers. A router attached to multiple areas that runs multiple copies of the basic algorithm, one copy for each attached area and an additional copy for the backbone. Area border routers condense the topology information of attached areas for distribution to the backbone. The backbone then distributes this information to other areas.
- AS Boundary Routers. A router that exchanges information with routers that belong to other ASs. These routers import this information to the OSPF routing domain in AS external link advertisements.

3.4. OSPF Routing Summary

When a router is initialized, it uses the Hello Protocol to send hello packets to its neighbors, and they in turn send their packets to the router. On broadcast and point-to-point networks, the router dynamically detects its neighboring routers by sending the Hello packets to the multicast address ALLSPFRouters; on non-broadcast networks you must configure information to help the router discover its neighbors. On all multi-access networks (broadcast and non-broadcast), the Hello Protocol also elects a designated router for the network.

The router then attempts to form adjacencies with its neighbors to synchronize their topological databases. Adjacencies control the distribution (sending and receiving) of the routing protocol packets as well as the distribution of the topological database updates. On a multi-access network, the designated router determines which router becomes adjacent.

A router periodically advertises its status or link state to its adjacencies. Link state advertisements flood throughout an area ensuring that all routers have exactly the same topological database. This database is a collection of the link state advertisements received from each router belonging to an area. From the information in this database, each router can calculate a shortest path tree with itself designated as the root. Then the shortest path tree generates the routing table.

a) Designated Router

Every multi-access network has a designated router that performs two main functions for the routing protocol, it originates network link advertisements and it becomes adjacent to all other routers on the network.

When a designated router originates network link advertisements, it lists all the routers, including itself, currently attached to the network. The link ID for this advertisement is the IP interface address of the designated router. By using the subnet/network mask, the designated router obtains the IP network number.

The designated router becomes adjacent to all other routers and is tasked with synchronizing the link state databases on the broadcast network.

The Hello Protocol elects the designated router after determining the routers priority from the Rtr Pri field of the Hello packet.

When a router's interface first becomes functional, it checks to see if the network currently has designated router. If it does, it accepts that designated router regardless of that router's priority, otherwise, it declares itself the designated router at the same time another router does, the router with highest router priority (Rtr Pri) becomes the designated router. In the case that both Rtr Pris are equal, the one with the higher router ID is elected.

Once the designated router is elected. It becomes the endpoint for many adjacencies. On a broadcast network this optimizes the flooding procedure by allowing the designated route to multicast its Link State Update packets to the address ALLSPFRouters rather than sending separate packets over each adjacency.

To set the OSPF parameters for the router's network interfaces, use the **SET INTERFACE** command, then respond to the prompts.

When responding to the prompts, supply the interface's IP address for each interface in the router and answer the questions that follow. For the parameters listed below you must enter the same value for all routers attached to a common network segment.

- Hello interval
- Dead router interval
- Authentication key (if an authentication type of 1 (simple password) is used)

The code line "Attaches to area" asks for the OSPF area which the interface attaches to. In the following example, suppose that the interface address mask is 255.255.255.0, indicating that the interface attaches to a subnet (128.185.138.0) of network 128.185.0.0. All other OSPF routers attached to subnet 128.185.138.0 must also have their *hello interval* set to 10, *dead router interval* set to 40, and their interface *authentication key* set to xyz_q.

Example:

```
OSPF Config> SET INTERFACE
Interface IP address [0.0.0.0]? 192.7.1.253
Attaches to area [0.0.0.0]? 0.0.0.1
Retransmission Interval (in seconds)[5]?
Transmission Delay (in seconds)[1]?
Router Priority[1]?
Hello Interval (in seconds)[10]?
Dead Router Interval (in seconds)[40]?
Type Of Service 0 cost[1]?
Authentication Key[]?
Retype Auth. Key[]?
OSPF Config>
```

3.5. Setting Non-Broadcast Network Parameters

If the router is connected to a non-broadcast, multi-access network, such as an X.25 PDN, you have to configure the parameters below to help the router discover its OSPF neighbors. This configuration is only necessary if the router will be eligible to become designated router of the non-broadcast network.

First configure the OSPF poll interval with the following command:

Example:

```
OSPF Config> SET NON-BROADCAST
 Interface IP address [0.0.0.0]? 192.7.1.253
Poll Interval[120]?
OSPF Config>
```

Then configure the IP addresses of all other OSPF routers that will be attached to the non-broadcast network. For each router configured, you must also specify its eligibility to become the designated router.

Example:

```
OSPF Config> ADD NEIGHBOR
 Interface IP address [0.0.0.0]? 192.7.1.253
IP Address of Neighbor [0.0.0.0]? 192.7.1.251
Can that router become Designated Router on this net[Yes]? n
OSPF Config>
```

3.6. Enabling AS Boundary Routing

To import routes learned from other protocols (RIP and statically configured information) into the OSPF domain, enable AS boundary routing. You must do this even if the only route you want to import is the default route (destination 0.0.0.0).

When enabling AS boundary routing, you are asked which external routes you want to import. You can choose to import, or not to import, routes belonging to several categories. The categories are as follows:

- RIP routes
- Static routes
- Direct routes
- Default route.

For example, you can choose to import direct routes, but not RIP or static routes. All routes are imported with cost equal to their routing table cost. They are all imported as either type 1 or type 2 external routes, depending on the routing protocol comparison.

Independently of the above external categories, you can also configure whether or not to import subnet routes into the OSPF domain. This configuration item defaults to OFF (subnets not imported).

The metric type used in importing routes determines how the imported cost is viewed by the OSPF domain. When comparing two type 2 metrics, only the external cost is considered in picking the best route. When comparing two type 1 metrics, the external and internal costs of the route are combined before making the comparison.

Combinations of these options are possible. For example, you can set the router so that its default is originated only if a route to 10.0.0.0 is received from AS number 12. Setting the AS number to 0 means “from any AS”. Setting the network number to 0.0.0.0 means “any routes received”.

The syntax of the **ENABLE** command is as follows:

Example:

```
OSPF Config> ENABLE AS
Import RIP routes(Yes/No)(N)?
Import static routes(Yes/No)(N)?
Import direct routes(Yes/No)(N)?
Import subnet routes(Yes/No)(N)?
Import default route(Yes/No)(N)?
Aggregation type:
   1.- Do not aggregate
   2.- Aggregate subnets
   3.- Use aggregation routes
   4.- Aggregate subnets and use aggregation routes
Enter option: [2]?
Import routes cost[0]?
Always originate default route(Yes/No)(N)?
OSPF Config>
```

3.7. Other Configuration Tasks

a) Setting OSPF Router IDS

Every router in a OSPF routing domain must be assigned a 32-bit router-ID. The current OSPF implementation sets the OSPF router-ID to be the address of the first OSPF interface appearing in the router's configuration.

The OSPF router-ID can also be explicitly set by the **SET ROUTER-ID** command from the IP menu. The router-ID must still be one of the router's IP interface addresses.

b) Setting Virtual Links

To maintain backbone connectivity you must have all of your backbone routers interconnected either by permanent or virtual links. Virtual links may be configured between any two area border routers that share a common non-backbone and non-stub area. Virtual links are considered to be separate router interfaces connecting to the backbone area. Therefore, you are asked to also specify many of the interface parameters when configuring a virtual link.

The example below illustrates the configuration of a virtual link. Virtual links must be configured in each of the link's two endpoints. Note that OSPF router IDs are entered in the same form as IP addresses.

Example:

```
OSPF Config> SET VIRTUAL-LINK
Virtual endpoint (Router ID) [0.0.0.0]? 192.7.1.254
Link's transit area [0.0.0.1]?
Retransmission Interval (in seconds)[10]?
Transmission Delay (in seconds)[5]?
Hello Interval (in seconds)[30]?
Dead Router Interval (in seconds)[180]?
Authentication Key[]?
OSPF Config>
```

c) Configuring for Routing Protocol Comparisons

If you use a routing protocol in addition to OSPF, or when you change your routing protocol to OSPF, you must set the Routing Protocol Comparison. OSPF routing in an AS occurs on the following three levels: Intra-area, Inter-area, and exterior.

Intra-area routing occurs when a packet's source and destination address reside in the same area. For example, N1 and N2 in Area 1of. Information that is about other areas does not affect this type of routing.

Inter-area routing occurs when the packet's source and destination addresses reside in different areas of an AS, for example, N1 of Area 1 and N7 of Area 2. OSPF does inter-area routing by dividing the path into three contiguous pieces: an intra-area path from source to an area border router; a backbone path between the source and destination areas; and then another intra-area path to the destination. You can visualize this high-level of routing as a star topology with the backbone as hub and each of the areas as a spoke.

Exterior routes are paths to networks that lie outside the AS. These routes originate either from routing protocols, such as Routing Information Protocol (RIP), or from static routes entered by the network administrator. The exterior routing information provided by RIP does not interfere with the internal routing information provided by the OSPF protocol.

AS boundary routers may import exterior routes into the OSPF routing domain. OSPF represents these routes as AS external link advertisements.

OSPF imports external routes in separate levels. The first level, called type 1 routes, is used when the external metric is comparable to the OSPF metric (e.g., they might both use delay in milliseconds). The second level, called external type 2 routes, assumes that the external cost is greater than the cost of any internal OSPF (link-state) path.

Imported external routes are tagged with 32-bits of information. In a router, this 32-bit field indicates the AS number from where the route was received. This enables more intelligent RIP behavior when determining whether to re-advertise the external information to other ASs.

OSPF has a 4-level routing hierarchy as shown below. The **SET COMPARISON** command tells the router where the RIP/static routes fit in the OSPF hierarchy. The two lower levels consist of the OSPF internal routes. OSPF intra-area and inter-area routes take precedence over information obtained from any other sources, all of which are located on a single level.

To put the RIP/static routes on the same level as OSPF external type 1 routes, set the comparison to 1. To put the RIP/static routes on the same level as OSPF external type 2 routes, set the comparison to 2. The default setting is 2. For example, suppose the comparison is set to 2. In this case, when RIP routes are imported into the OSPF domain, they will be imported as type 2 externals. All OSPF external type 1 routes override received RIP routes, regardless of metric. However, if the RIP routes have a smaller cost, the RIP routes override OSPF external type 2 routes.

> ***The comparison values for all of your OSPF routers must match. If the comparison values set for the routers are inconsistent, your routing will not function properly.***

The syntax of the **SET COMPARISON** command is as follows:

Example:

```
OSPF Config>SET COMPARISON
Compare to type 1 or 2 externals [2]?
OSPF Config>
```

Chapter 2
Configuration

1. Configuration Commands

This chapter describes the OSPF configuration commands. To access to the OSPF Configuration environment you must enter the following:

```
*P 4
User Configuration
Config> PROTOCOL OSPF
Open SPF-Based Routing Protocol configuration console
OSPF Config>
```

Command	Function
?(HELP)	Displays available commands or options.
ADD	Adds to already existent OSPF information. You can add ranges to areas, and neighbors to non-broadcast networks.
DELETE	Deletes OSPF information from SRAM.
DISABLE	Disables the entire OSPF protocol, AS boundary routing capability, or IP multicast routing.
ENABLE	Enables the entire OSPF protocol, AS boundary routing capability, or IP multicast routing.
LIST	Displays OSPF configuration.
SET	Establishes or changes the configuration information concerning OSPF areas, interfaces, non-broadcast networks, or virtual links. This command also allows you to set the way in which OSPF routes are compared to information gained from other routing protocols.
EXIT	Exits the OSPF configuration process.

Letters written in **bold** type are the minimum set of characters you must enter to use a command properly.

1.1. ? (HELP)

Use the **?** (HELP) command to list the commands that are available from the current prompt level. You can also enter ? after a command to list its options.

Syntax:

```
OSPF Config> ?
```

Example:

```
OSPF Config> ?
ADD
DELETE
DISABLE
ENABLE
LIST
SET
EXIT
OSPF Config>
```

1.2. ADD

Use the **ADD** command to add more information to already existing OSPF information. With this command you can add ranges to areas as well as neighbors to non-broadcast networks.

Syntax:

```
OSPF Config> ADD ?
NEIGHBOR
RANGE
```

a) *ADD NEIGHBOR*

Adds neighbors to non-broadcast networks. If the router is connected to a non-broadcast, multi-access network, such as an X.25 PDN, you have to use this command to help the router discover its OSPF neighbors. This configuration is only necessary if the router will be eligible to become designated router of the non-broadcast network. Configure the IP addresses of all other OSPF routers that will be attached to the non-broadcast network.

For each router configured, you must also specify its eligibility to become designated router.

Example:

```
OSPF Config> ADD NEIGHBOR
 Interface IP address [0.0.0.0]? 192.7.1.253
IP Address of Neighbor [0.0.0.0]? 192.7.1.251
Can that router become Designated Router on this net[Yes]? n
 OSPF Config>
```

b) *ADD RANGE <area#> <IP-address> <IP-address-mask>*

Add ranges to OSPF areas. OSPF areas are defined in terms of address ranges. External to the area, a single route is advertised for each address range. For example, if an OSPF area were to consist of all subnets of the class B network 128.185.0.0, it would be defined as consisting of a single address range. The address range would be specified as an address of 128.185.0.0 together with a mask of 255.255.0.0. Outside of the area, the entire subnetted network would be advertised as a single route to network 128.185.0.0.

Example:

```
OSPF Config> ADD RANGE
Area number [0.0.0.0]? 0.0.0.1
IP Address [0.0.0.0]? 1.1.1.0
IP Address Mask [0.0.0.0]? 255.255.255.0
Inhibit advertisement(Yes/No)(N)? y
OSPF Config>
```

1.3. DELETE

Use the **DELETE** command to delete OSPF information from the router's configuration memory.

Syntax:

```
OSPF Config> DELETE ?
AREA
INTERFACE
NEIGHBOR
NON-BROADCAST
RANGE
VIRTUAL-LINK
```

a) DELETE AREA <area#>

Deletes OSPF areas from the current OSPF configuration.

Example:

```
OSPF Config> DELETE AREA 0.0.0.1
OSPF Config>
```

b) DELETE INTERFACE <interface-IP-address>

Deletes an interface from the current OSPF configuration.

Example:

```
OSPF Config> DELETE INTERFACE 128.185.138.19
OSPF Config>
```

c) DELETE NEIGHBOR

Deletes neighbors on non-broadcast networks from the current OSPF configuration.

Example:

```
OSPF Config> DELETE NEIGHBOR
Interface IP address [0.0.0.0]? 192.7.1.253
IP Address of Neighbor [0.0.0.0]? 192.7.1.251
OSPF Config>
```

d) DELETE NON-BROADCAST <interface-IP-address>

Deletes non-broadcast network information from the current OSPF configuration.

Example:

```
OSPF Config> DELETE NON-BROADCAST 128.185.133.21
OSPF Config>
```

e) DELETE RANGE <area#> <IP-address>

Deletes ranges from OSPF areas

Example:

```
OSPF Config> DELETE RANGE
Area number [0.0.0.0]? 0.0.0.1
IP Address [0.0.0.0]? 1.1.1.0
IP Address Mask [0.0.0.0]? 255.255.255.0
OSPF Config>
```

f) DELETE VIRTUAL-LINK

Deletes a virtual link. Virtual links can be configured between any two backbone routers that have an interface to a common non-backbone area. Virtual links are used to maintain backbone connectivity and must be configured at both endpoints.

Example:

```
OSPF Config> DELETE VIRTUAL-LINK
Virtual endpoint (Router ID) [0.0.0.0]? 192.7.1.254
Link's transit area [0.0.0.1]?
OSPF Config>
```

1.4. DISABLE

Use the **DISABLE** command to disable either the entire OSPF protocol or just the AS boundary routing capability.

Syntax:

```
OSPF Config> DISABLE?
AS boundary routing
OSPF routing protocol
```

a) AS boundary routing

Disables the AS boundary routing capability. When disabled, the router will not import external information into the OSPF domain.

Example:

```
OSPF Config> DISABLE AS
OSPF Config>
```

b) OSPF routing protocol

Disables the entire OSPF protocol.

Example:

```
OSPF Config> DISABLE OSPF
OSPF Config>
```

1.5. ENABLE

Use the **ENABLE** command to enable either the entire OSPF protocol or just the AS boundary routing capability.

Syntax:

```
OSPF Config> ENABLE?
AS boundary routing
OSPF routing protocol
```

a) ENABLE AS boundary routing

Enables the AS boundary routing capability that allows you to import routes learned from other protocols (RIP, and statically configured information) into the OSPF domain.

This also permits you to configure the type of aggregation and additional cost.

The meaning of the types of aggregation are as follows:

Do not aggregate	No aggregation is carried out of any type. Neither the aggregation routes nor the subnet aggregation routes are imported here. This is the default option.
Aggregate subnets	When in the route table a subnet route is learnt or configured, an "Sbnt" type route or a subnet aggregation route automatically appears with a destination of "subnet network" and the next hop as "none". On activating this type of aggregation, the subnet aggregation routes are imported only when the route being aggregated is one within the set of routes to import.
Use aggregation routes	The aggregation routes are not truly routes but marks that appear in the active routes table which indicate that there exists a series of routes which are being aggregated. On activating this type of aggregation, only the aggregation routes and the routes which do not belong to any aggregation are imported. This means that the aggregated routes are not imported. Also on activating this type of aggregation, the subnet aggregation routes are imported only when the route being aggregated is one within the set of routes to import.

This additional cost parameter ensures that all the imported routes have their costs increased however many units indicated by this parameter. The default value is zero.

Example:

```
OSPF Config> ENABLE AS
Import RIP routes(Yes/No)(N)? y
Import static routes(Yes/No)(N)? y
Import direct routes(Yes/No)(N)? y
Import subnet routes(Yes/No)(N)? y
Import default route(Yes/No)(N)? n
Aggregation type:
   1.- Do not aggregate
   2.- Aggregate subnets
   3.- Use aggregation routes
   4.- Aggregate subnets and use aggregation routes
Enter option: [2]? 2
Import routes cost[0]? 1
Always originate default route(Yes/No)(N)? n
OSPF Config>
```

b) ENABLE OSPF routing protocol

Enables the entire OSPF protocol. When enabling the OSPF routing protocol, you must supply the following values that will be used to estimate the size of the OSPF link state database:

- Estimated number of AS external routes that will be imported into the OSPF routing domain. A single destination may lead to multiple external routes when it is imported by separate AS boundary routers. For example, if the OSPF routing domain has two AS boundary routers, both importing routes to the same 100 destinations, the number of AS external routes should be set to 200.
- Estimated number of OSPF routers in the routing domain.

Example:

```
OSPF Config> ENABLE OSPF
Estimated # external routes[0]? 500
Estimated # OSPF routers[0]? 50
OSPF Config>
```

1.6. LIST

Use the **LIST** command to display OSPF configuration information.

Syntax:

```
OSPF Config> LIST ?
ALL
AREAS
INTERFACES
NEIGHBORS
NON-BROADCAST NETWORK DESCRIPTION
VIRTUAL-LINKS
```

a) *LIST ALL*

List all OSPF related configuration information.

Example:

```
OSPF Config> LIST ALL
        --Global configuration--
OSPF Protocol:          Enabled
# AS ext. routes:       500
Estimated # routers:    50
External comparison:    Type 2
AS boundary capability: Enabled
Import external routes: RIP STA DIR SUB
Aggregate subnets
External routes cost:   1
Orig. default route:    No (0,0.0.0.0)
Default route cost:     (1, Type 2)
Default forward. addr.: 0.0.0.0
Multicast forwarding:   Disabled

                        --Area configuration--
Area ID          AuType          Stub?  Default-cost  Import-summaries?
0.0.0.0          0=None           No     N/A           N/A
0.0.0.1          1=Simple-pass    No     N/A           N/A

                        --Interface configuration--
IP address        Area              Cost  Rtrns  TrnsDly  Pri  Hello  Dead
192.3.1.2         0.0.0.1              1      5       1     1     10    40
192.7.1.253       0.0.0.0              1      5       1     1     10    40

         --NBMA configuration--
Interface Addr      Poll Interval
                192.168.253.1       120

        --Neighbor configuration--
Neighbor Addr     Interface Address   DR eligible?
192.3.1.1         192.3.1.2           yes
OSPF Config>
```

The meaning of each field is:

OSPF protocol	Displays whether OSPF is enabled or disabled.
# AS ext. Routes	Estimated number of Autonomous System external routes. The router cannot accept more than this number of AS external routes.
Estimated # routers	Estimated number of routers found in the OSPF configuration.
External comparison	External route type used by OSPF when importing external information into the OSPF domain and when comparing OSPF external routes to RIP routes.
AS boundary capability	Displays whether the router will import external routes into the OSPF domain.

Import external routes	Displays which routes will be imported.
Aggregate subnets	Displays the type of aggregation configured.
External routes cost	Displays the configured additional cost.
Orig default route	Displays whether the router will import a default into the OSPF domain. When the value is "YES", a non-zero network number is displayed in parentheses. This indicates that the default route will originated if and only if a route to that network is available.
Default route cost	Cost and type that will be used in the imported default route.
Default forward addr	Forwarding address that will be used in the imported default route.
Multicast forwarding	Displays if the multicast routing is enabled or not.
Area-ID	Attached area ID (area summary information).
AuType	Method used for area authentication. "Simple-pass" means a simple password scheme is being used for the area's authentication.
Stub area	Displays whether or not the area being summarized is a stub area. Stub areas do not carry external route, resulting in a smaller routing database. However, stub areas cannot contain AS boundary routers, nor can they support configured virtual links.
OSPF Interfaces	For each interface, its IP address is printed, together with configured parameters. "Area" is the OSPF area to which the interface attaches. "Cost" indicates the TOS 0 cost (or metric) associated with the interface. "Rtrns" is the retransmission interval, which is the number of seconds between retransmissions of unacknowledged routing information. "TrnsDly" is the transmission delay, which is an estimate of the number of seconds it takes to transmit routing information over the interface (it must be greater than 0). "Pri" is the interface´s Router Priority, which is used when selecting the designated router. "Hello" is the number of seconds between Hello Packets sent out the interface. "Dead" is the number of seconds after Hellos cease to be heard that the router is declared down.
Virtual-link	Lists all virtual links that have been configured with this router as endpoint. "Virtual endpoint" indicated the OSPF Router ID of the other endpoint. "Transit area" indicates the non-backbone area through which the virtual link is configured. Virtual links are considered treated by the OSPF protocol similarly to point-to-point networks. The other parameters listed in the command ("Rtrns", "TrnsDly", "Hello", and "Dead") are maintained for all interfaces. See the OSPF **LIST INTERFACES** command for more information.

b) LIST AREAS

Lists all information concerning configured OSPF areas.

Example:

```
OSPF Config> LIST AREAS
Area ID          AuType          Stub? Default-cost Import-summaries?
0.0.0.0          0=None           No         N/A           N/A
0.0.0.1          1=Simple-pass    No         N/A           N/A

                        --Area ranges--
Area ID          Address          Mask             Advertise?
0.0.0.0          1.1.1.0          255.255.255.0    No
 OSPF Config>
```

Area-ID	Attached area ID (area summary information).
AuType	Method used for area authentication. "Simple-pass" means a simple password scheme is being used for the area's authentication.
Stub area	Displays whether or not the area being summarized is a stub area.

c) LIST INTERFACES

For each interface its IP address is printed, together with configured parameters.

Example:

```
OSPF Config> LIST INTERFACES
                        --Interface configuration--
IP address        Area              Cost  Rtrns  TrnsDly  Pri  Hello  Dead
192.3.1.2         0.0.0.1           1     5      1        1    10     40
192.7.1.253       0.0.0.0
OSPF Config>
```

Area	The OSPF area which the interface is connected to.
Cost	The TOS 0 (or metric) associated with the interface.
Rtrns	Retransmission interval, i.e. the number of seconds between unrecognized routing information retransmissions.
TrnsDly	This is the transmission delay which is an estimation of the number of seconds it takes to transmit the routing information through the interface (this should be a value greater than zero).
Pri	This is the Priority of the interface router which is sued when you select the Designated Router DR.
Hello	This is the number of seconds between the "Hello" packets sent by the interface.
Dead	This is the number of seconds which should pass after the "Hello" packets to consider that the router is down and not operative.

d) LIST NEIGHBORS

Lists all the information related to the neighbors.

Example:

```
OSPF Config> LIST NEIGHBORS
        --Neighbor configuration--
Neighbor Addr     Interface Address   DR eligible?
192.3.1.1         192.3.1.2           yes
OSPF Config>
```

Neighbor Addr Neighbor IP address.

Interface Address Interface IP address.

DR eligible If the designated router is eligible.

e) LIST NON-BROADCAST

List all information related to interfaces connected to non-broadcast networks. For each non-broadcast interface, as long as the router is eligible to become designated router on the attached network, the polling interval is displayed together with a list of the router's neighbors on the non-broadcast network.

Example:

```
OSPF Config> LIST NON-BROADCAST
         --NBMA configuration--
Interface Addr      Poll Interval
192.168.253.1       120
OSPF Config>
```

f) LIST VIRTUAL-LINKS

List all virtual links that have been configured with this router as endpoint. "Virtual endpoint" indicated the OSPF router ID of the other endpoint. "Transit area" indicates the non-backbone area through which the virtual link is configured. Virtual links are considered treated by the OSPF protocol similarly to point-to-point networks. The other parameters listed in the command ("Rtrns", "TrnsDly", "Hello", and "Dead") are maintained for all interfaces. See the OSPF **LIST INTERFACES** command for more information.

Example:

```
OSPF Config> LIST VIRTUAL-LINKS
              --Virtual link configuration--
Virtual endpoint     Transit area      Rtrns  TrnsDly Hello  Dead
192.7.1.153          0.0.0.1           10     5       30     180
OSPF Config>
```

1.7. SET

Use the **SET** command to display or change the configuration information concerning OSPF areas, interfaces, non-broadcast networks, or virtual links. This command also allows you to set the way in which OSPF routes are compared to information obtained from other routing protocols.

Syntax:

```
OSPF Config> SET?
AREA
COMPARISON
INTERFACE
NON-BROADCAST
VIRTUAL-LINK
```

a) SET AREA

Sets the parameters for an OSPF area. If no areas are defined, the router software assumes that all the router´s directly attached networks belong to the backbone area (area ID 0.0.0.0).

Example:

```
OSPF Config> SET AREA
Area number [0.0.0.0]? 0.0.0.2
Authentication Type[0]? 1
Is this a stub area(Yes/No)(N)? n
OSPF Config>
```

The meaning of each field is:

Area number	OSPF area address. An OSPF area is a contiguous group of networks that is defined by a list of address ranges, each indicated by a combination of the IP address mask. A network belongs to an area if its address is in the list.
Authentication type	Security scheme to be used in the area. The choices are 1, which indicates a simple password; or 0, which indicates that no authentication is necessary to pass packets.
Is this a stub area?	If you designate yes • The area does not receive any AS external link advertisements, reducing the size of your data-base and decreasing memory usage for routers in the stub area. • You cannot configure virtual links through a stub area. • You cannot configure a router within the stub area as an AS boundary routers.
External Routing in Stub Areas	You cannot configure the backbone as a stub area. External routing in stub areas is based on a default route. Each border area router attaching to a stub area originates a default route for this purpose. The cost of this default route is also configurable with the **SET AREA** command.

b) SET COMPARISON

Tells the router where the RIP/static routes fit in the OSPF hierarchy. The two lower levels consist of the OSPF internal routes. OSPF internal routes take precedence over information gained from any other sources, all of which are located on a single level.

Example:

```
OSPF Config> SET COMPARISON
Compare to type 1 or 2 externals[2]?
OSPF Config>
```

c) SET INTERFACE

Sets the OSPF parameters for the router's network interfaces.

Example:

```
OSPF Config> SET INTERFACE
Interface IP address [0.0.0.0]? 192.7.1.253
Attaches to area [0.0.0.0]?
Retransmission Interval (in seconds)[5]?
Transmission Delay (in seconds)[1]?
Router Priority[1]?
Hello Interval (in seconds)[10]?
Dead Router Interval (in seconds)[40]?
Type Of Service 0 cost[1]?
Authentication Key[]?
Retype Auth. Key[]?
OSPF Config>
```

When responding to the prompts, supply the IP address for each interface in the router and answer the questions that follow. For the parameters listed below you must enter the same value for all routers attached to a common network.

- Hello interval
- Dead router interval
- Authentication key (if an authentication of 1 is used)

The code line "Attaches to area" prompt asks for the OSPF area which the interface attaches to. For example, suppose that the interface address mask is 255.255.255.0, indicating that the interface attaches to a subnet (128.185.138.0) of network 128.185.0.0. All other OSPF routers attached to subnet 128.185.138.0 must also have their hello interval set to 10, dead router interval set to 40, and their interface authentication key set to xyz_q.

d) SET NON-BROADCAST

Helps the router discover its OSPF neighbors. This configuration is only necessary if the router will be eligible to become designated router of the non-broadcast network. After using this command you must then configure the IP addresses of all other OSPF routers that will be attached to the non-broadcast network. See the **ADD NEIGHBOR** command for more information.

Example:

```
OSPF Config> SET NON-BROADCAST
 Interface IP address [0.0.0.0]? 192.7.1.253
Poll Interval[120]?
OSPF Config>
```

e) SET VIRTUAL-LINK

Configures virtual links between any two area border routers. To maintain backbone connectivity you must have all of your backbone routers interconnected either by permanent or virtual links. Virtual links are considered to be separate router interfaces connecting to the backbone area. Therefore, you are asked to also specify many of the interface parameters when configuring a virtual link.

Example:

```
OSPF Config> SET VIRTUAL-LINK
Virtual endpoint (Router ID) [0.0.0.0]? 192.7.1.153
Link's transit area [0.0.0.1]?
Retransmission Interval (in seconds)[10]?
Transmission Delay (in seconds)[5]?
Hello Interval (in seconds)[30]?
Dead Router Interval (in seconds)[180]?
Authentication Key[]?
OSPF Config>
```

1.8. EXIT

Use the **EXIT** command to return to the previous prompt level.

Syntax:

```
OSPF Config> EXIT
```

Example:

```
OSPF Config> EXIT
Config>
```

Chapter 3
Monitoring

1. Monitoring Commands

This section describes the OSPF monitoring commands. To access to the OSPF Configuration environment you must enter the following:

```
*P 3
Console Operator
+PROTOCOL OSPF
Open SPF-Based Routing Protocol console
OSPF>
```

Command	Function
?(HELP)	Displays available commands or options.
ADVERTISEMENT expansion	Displays a link state advertisement belonging to the OSPF database.
AREA SUMMARY	Displays OSPF area statistics and parameters.
AS-EXTERNAL advertisement	Lists the AS external advertisements belonging to the OSPF link state database.
DATABASE summary	Displays the advertisements belonging to an OSPF area's link state database.
DUMP routing tables	Displays the OSPF routes contained in the routing table.
INTERFACE summary	Displays OSPF interface statistics and parameters.
NEIGHBOR summary	Displays OSPF neighbor statistics and parameters.
PING address	Used to test for network reachability and fault isolation.
ROUTERS	Displays the reachable OSPF area-border routers and AS-boundary routers.
SIZE	Displays the number of LSAs currently in the link state database, categorized by type.
STATISTICS	Displays OSPF statistics detailing memory and network usage.
TRACEROUTE address	Displays the route taken to a host or network gateway.
WEIGHT	Dynamically changes the cost of an OSPF interface.
EXIT	Exits the OSPF monitor process.

Letters written in **bold** type are the minimum set of characters you must enter to use a command properly.

1.1. ? (HELP)

Use the **?** (HELP) command to list the commands that are available from the current prompt level. You can also enter ? after a command to list its options.

Syntax:

```
OSPF> ?
```

Example:

```
OSPF> ?
ADVERTISEMENT expansion
AREA summary
AS-EXTERNAL advertisements
DATABASE summary
DUMP routing tables
INTERFACE summary
NEIGHBOR summary
PING address
ROUTERS
SIZE
STATISTICS
TRACEROUTE address
WEIGHT
EXIT
OSPF>
```

1.2. ADVERTISEMENT expansion

Use the **ADVERTISEMENT expansion** command to print the contents of a link state advertisement contained in the OSPF database. For a summary of the router's advertisements use the **DATABASE summary** command.

A link state advertisement is defined by its link state type, link state ID and its advertising router. There is a separate link state database for each OSPF area. Providing an area-id on the command line tells the software which database you want to search.

> ***NOTE: Link State IDs, advertising routers (specified by their router IDs), and area IDs take the same format as IP addresses. For example, the backbone area can be entered as 0.0.0.0.***

The different kinds of advertisements, which depend on the value given for link-state-type, are

- Router links - Contain descriptions of a single router's interface.
- Network links - Contain the list of routers attached to a particular interface.
- Summary nets - Contain descriptions of a single inter-area route.
- Summary AS boundary routers - Contain descriptions of the route to an AS boundary router in another area.
- AS external nets - Contain descriptions of a single route.

The example below shows an expansion of a router links advertisement. The router's ID is 128.185.184.11. It is an AS boundary router and has three interfaces to the backbone area (all of cost 1). Detailed field descriptions are provided with the example shown below.

This command has also been enhanced in two ways. First of all, when displaying router-LSAs and network-LSAs, the reverse cost of each router-to-router link and router-to-transit-network link is displayed, as well as the previously displayed forward cost. This is done because routing of multicast datagrams whose source lies in different areas/ASs is based on reverse cost instead of forward cost. In

those cases where there is no reverse link (which means that the link will never be used by the Dijkstra), the reverse cost is shown as "1-way".

In addition, the LSA's OSPF options are displayed in the same manner as they were displayed in the detailed OSPF **NEIGHBOR summary** command.

New group-membership-LSAs can also be displayed. The "LS destination" of each group-membership-LSA is a group address. A router originates a group-membership-LSA for each group with members on one or more of the router's attached networks.

Syntax:

```
OSPF> ADVERTISEMENT
Link state type[1]?
Link state ID (destination) [0.0.0.0]?
For which area [0.0.0.0]?
```

Example:

```
OSPF> ADVERTISEMENT
Link state type[1]?
Link state ID (destination) [0.0.0.0]? 192.7.1.253
For which area [0.0.0.0]?

        LS age:     1693
        LS options:  E
        LS type:     1
        LS destination (ID): 192.7.1.253
        LS originator:   192.7.1.253
        LS sequence no:  0x80000002
        LS checksum:     0xCF63
        LS length:       36
        Router type:  ABR,ASBR
        # router ifcs:   1
                Link ID:          192.7.1.254
                Link Data:        192.7.1.253
                Interface type:   2
                        No. of metrics: 0
                        TOS 0 metric:   1 (0)

OSPF>
```

LS age	Age of the advertisement in seconds
LS options	Optional OSPF capabilities supported by the piece of the routing domain described by the advertisement. These capabilities are denoted by E (processes type 5 externals; when this is not set to the area to which the advertisement belongs has been configured as a stub), T (can route based on TOS).
LS type	Classifies the advertisement and dictates its contents: 1 (router links advertisement), 2 (network link advertisement), 3 (summary link advertisement), 4 (summary ASBR advertisement), 5 (AS external link) and 6 (group-membership advertisement).
LS destination	Identifies what is being described by the advertisement. Depends on the advertisement type. For router links and ASBR summaries, it is the OSPF router ID. For network links, it is the IP address of the network's designated router. For summary links and AS external links, it is a network/subnet number.
LS originator	OSPF router ID of the originating router.

LS sequence number	Used to distinguish separate instances of the same advertisement. Should be looked at as a signed 32-bit integer. Starts at 0x80000001, and increments by one each time the advertisement is updated.
LS checksum	A checksum of advertisement contents, used to detect data corruption.
LS length	The size of the advertisement in bytes.
Router type	Level of functionality of the router. ASBR means that the router is an AS boundary router, ABR that the router is an area border router, and W that the router is a wildcard multicast receiver.
# Router ifcs	Router interface described in the advertisement.
Link ID	Indicates what the interface connects to. Depends on the interface type. For interfaces to routers (i.e., point-to-point links), the Link ID is the neighbor's router ID. For interfaces to transit networks, it is the IP address of the network designated router. For interfaces to stub networks, it is the network's network/subnet number.
Link Data	4 bytes of extra information concerning the link, it is either the IP address of the interface (for interfaces to point-to-point networks and transit networks), or the subnet mask (for interfaces to stub networks).
Interface type	One of the following: 1 (point-to-point connection to another router), 2 (connection to transit network), 3 (connection to stub network) or 4 (virtual link).
No. of metrics	The number of non-zero TOS values for which metrics are provided for this interface.
TOS 0 metric	The cost of the interface. In parenthesis the reverse cost of the link is given (derived from another advertisement). If there is no reverse link, "1-way" is displayed.

The *LS age*, *LS options*, *LS type*, *LS destination*, *LS originator*, *LS sequence no*, *LS checksum* and *LS length* fields are common to all advertisements. The *Router Type* and *# Router ifcs* are seen only in router links advertisements. Each link in the router advertisement is described by the *Link ID*, *Link Data*, and *Interface type* fields.

Each link can also be assigned a separate cost for each IP Type of Service (TOS); this is described by the *No. of metrics* and *TOS 0 metric* fields (the router currently does not router based on TOS, and looks at the TOS 0 cost only).

1.3. AREA summary

Use the **AREA summary** command to display the statistics and parameters for all OSPF areas attached to the router.

In the example below, the router attaches to a single area (the backbone area). A simple password scheme is being used for the area's authentication. The router has three interfaces attaching to the area, and has found 4 transit networks, 7 routers and no area border routers when doing the SPF tree calculation for the backbone.

Syntax:

```
OSPF> AREA
```

Example:

```
OSPF> AREA
Area ID          Authentication  #ifcs  #nets  #rtrs  #brdrs
0.0.0.0           None              1      1      2      1
0.0.0.1           None              1      0      1      1
OSPF>
```

 # ifcs Router interfaces attached to the particular area. These interfaces are not necessarily functional.

 # nets Transit networks found while doing the SPF tree calculation for this area.

 # rtrs Routers found when doing the SPF tree calculation for this area.

 # brdrs Area border routers found when doing the SPF tree calculation for this area.

1.4. AS-EXTERNAL advertisements

Use the **AS-EXTERNAL advertisements** command to list the AS external advertisements belonging to the OSPF routing domain. On line is printed for each advertisement. Each advertisement is defined by the following three parameters: its link state type (always 5 for AS external advertisements), its link state ID (called the LS destination), and the advertising router (called the LS originator).

Syntax:

```
OSPF> AS-EXTERNAL advertisements
```

Example:

```
OSPF> AS-EXTERNAL advertisements
Type  Ls destination   LS originator   Seqno       Age  Xsum
5     0.0.0.0          128.185.123.22  0x80000084  430  0x41C7
5     128.185.131.0    128.185.123.22  0x80000080  450  0x71DC
5     128.185.132.0    128.185.123.22  0x80000080  450  0x66E6
5     128.185.144.0    128.185.123.22  0x80000002  329  0xF2CA
5     128.185.178.0    128.185.123.22  0x80000081  450  0x72AA
5     128.185.178.0    128.185.129.40  0x80000080  382  0xDD28
5     129.9.0.0        128.185.123.22  0x80000082  451  0x4F30
5     129.9.0.0        128.185.126.24  0x80000080  676  0x324A
5     134.216.0.0      128.185.123.22  0x80000082  451  0x505A
5     134.216.0.0      128.185.126.24  0x80000080  676  0x3374
5     192.9.3          128.185.123.22  0x80000082  451  0xF745
5     192.9.3          128.185.126.24  0x80000080  677  0xDA5F
5     192.9.12         128.185.123.22  0x80000082  452  0x949F
5     192.9.12         128.185.128.41  0x80000080  679  0x31B2
5     192.26.100.0     128.185.123.22  0x80000081  452  0xFDCD
5     192.26.100.0     128.185.126.24  0x80000080  21   0xDEE8
etc.
```

```
# advertisements:  133
Checksum total:    0x43CC41
OSPF>
```

Type	Always 5 for AS external advertisements.
LS destination	IP network/subnet number. These network numbers belong to other Autonomous Systems.
LS originator	Advertising router.
Seqno, Age, Xsum	It is possible for several instances of an advertisement to be present in the OSPF routing domain at any one time. However, only the most recent instance is kept in the OSPF link state database (and printed by this command). The LS sequence number (Seqno), LS age (Age) and LS checksum fields (Xsum) are compared to see which instance is most recent. The LS age field is expressed in seconds. Its maximum value is 3600.

At the end of the display, the total number of AS external advertisements is printed, along with a checksum total over all of their contents. The checksum total is simply the 32-bit sum (carries discarded) of the individual advertisement's LS checksum fields. This information can be used to quickly determine whether two OSPF routers have synchronized databases.

1.5. DATABASE summary

Use the **DATABASE summary** command to display a description of the contents of a particular OSPF area's link state database. AS external advertisements are omitted from the display. A single line is printed for each advertisement. Each advertisement is defined by the following three parameters: its link state type (called Type), its link state ID (called the LS destination) and the advertising router (called the LS originator).

Syntax:

```
OSPF> DATABASE summary <area-id>
```

Example:

```
OSPF> DATABASE summary
For which area [0.0.0.0]?0.0.0.0
Type  LS destination   LS originator  Seqno       Age  Xsum
1*    10.1.2.7        10.1.2.7       0x80000025  390  0xB13C
1*    10.1.26.9       10.1.26.9      0x80000016  393  0x987D
1*    10.1.26.41      10.1.26.41     0x80000018  122  0x533D
1*    10.1.40.40      10.1.40.40     0x80000015  192  0x317C
1*    10.1.50.16      10.1.50.16     0x80000031  394  0x7A74
2*    10.1.25.40      10.1.40.40     0x80000006  193  0xCB35
2*    10.1.26.16      10.1.50.16     0x80000007  401  0x9669
3*    10.2.50.9       10.1.26.9      0x80000010  397  0xA430
3*    10.5.0.0        10.1.26.41     0x8000000F  133  0x4E9E
3*    10.5.50.41      10.1.26.9      0x80000006  394  0x5D5D
3*    128.185.214.0   10.1.40.40     0x8000000E  740  0x3CA2
6     224.185.0.0     10.1.50.16     0x8000000F  469  0x9B7A
6     225.0.1.36      10.1.2.7       0x80000006  405  0x5CC8
6     225.0.1.36      10.1.26.9      0x8000000F  404  0x8265
6     225.0.1.36      10.1.26.41     0x8000000F  133  0x3A4
6     225.0.1.36      10.1.40.40     0x8000000E  755  0x1D71
6     225.0.1.100     10.1.50.16     0x80000006  476  0x5E14
# advertisements:  17
Checksum total:   0x73121
OSPF>
```

Type	Separate LS types are numerically displayed: type 1 (router links advertisements), type 2 (network links advertisements), type 3 (network summaries), type 4 (AS boundary router summaries), and type 6 (group-membership-LSAs).
LS destination	Indicates what is being described by the advertisement.
LS originator	Advertising router.
Seqno, Age, Xsum	It is possible for several instances of an advertisement to be present in the OSPF routing domain at any one time. However, only the most recent instance is kept in the OSPF link state database (and printed by this command). The LS sequence number (Seqno), LS age (Age) and LS checksum fields (Xsum) are compared to see which instance is most recent. The LS age field is expressed in seconds. Its maximum value is 3,600.

At the end of the display, the total number of advertisements in the area database is printed, along with a checksum total over all of their contents. The checksum total is simply the 32-bit sum (carries discarded) of the individual advertisement's LS checksum fields. This information can be used to quickly determine whether two OSPF routers have synchronized databases.

1.6. DUMP routing tables

Displays the routes that have been calculated by OSPF and are now present in the routing table. Its output is similar to the monitoring **DUMP routing tables** command.

Syntax:

```
OSPF> DUMP
```

Example:

```
OSPF> DUMP
Type       Dest net          Mask      Cost Age  Next hop(s)

Sbnt(0)    1.0.0.0           FF000000  1    0    None
Aggr(0)A   1.1.0.0           FFFF0000  1    0    None
Stat(1)a   1.1.1.0           FFFFFF00  1    0    192.7.1.1
Sbnt(0)    5.0.0.0           FF000000  1    0    None
SPE2(0)    5.4.3.2           FFFFFFFF  0    1    192.7.1.1
Dir(1)     10.0.0.0          FF000000  1    0    FR/0
Dir(1)     192.3.1.0         FFFFFF00  1    0    FR/0
SPF(0)     192.3.1.2         FFFFFFFF  0    0    FR/0
Stat(1)    192.6.1.0         FFFFFF00  1    0    192.7.1.1
SPF(1)     192.7.1.0         FFFFFF00  1    1    Eth/0
OSPF>
```

Type	Route type and how the route was derived. SPF indicates an intra-area route; Rnge indicates an active area address range which is not used in forwarding packets. SPE1 and SPE2 indicate OSPF inter-area routes type 1 and 2, respectively.
Dest net	Destination host or network.
Mask	Entry's subnet mask.

Cost Route cost.

Age This field is applicable when running RIP. For OSPF routes, this field equals the cost of the route.

Next hop(s) Address of the next router on the path toward the destination host. A number in parentheses at the end of the column indicates the number of equal-cost routes to the destination. The first hops belonging to these routes can be displayed with the IP monitoring route.

1.7. INTERFACE summary

Use the **INTERFACE summary** command to display statistics and parameters related to OSPF. If no arguments are given, a single line is printed summarizing each interface. If an interface's IP address is given, detailed statistics for that interface will be displayed.

Syntax:

```
OSPF> INTERFACE < Interface-IP-address >
```

Example 1:

```
OSPF> INTERFACE
Ifc Address       Phys    assoc. Area       Type     State   #nbrs  #adjs
192.7.1.253       Eth/0   0.0.0.0           Brdcst   32      1      1
192.3.1.2         FR/0    0.0.0.1           P-2-MP   8       1      0
OSPF>
```

Ifc Address Interface IP address.

Assoc Area Attached area ID

Type Can be either Brdcst (broadcast, e.g., an Ethernet interface), P-P (a point-to-point network, e. g. a synchronous serial line), Multi (non-broadcast multi-access, e.g., an X.25 connection) and VLink (an OSPF virtual link).

State Can be one of the following: 1 (down), 2 (looped back), 4 (waiting), 8 (point-to-point), 16 (DR other), 32 (backup DR) or 64 (designated router).

#nbrs Number of neighbors. This is the number of routers whose hellos have been received, plus those that have been configured.

#adjs Number of adjacencies. This is the number of neighbors with whom the router has synchronized or is in the process of synchronization.

Example 2:

```
OSPF> INTERFACE 192.7.1.253
Interface address:      192.7.1.253
Attached area:          0.0.0.0
Physical interface:     Eth/0
Interface mask:         255.255.255.0
Interface type:         Brdcst
State:                  32
Designated Router:      192.7.1.254
Backup DR:              192.7.1.253

DR Priority:       1  Hello interval:   10  Rxmt interval:     5
Dead interval:    40  TX delay:          1  Poll interval:     0
Max pkt size:   1500  TOS 0 cost:        1

# Neighbors:       1  # Adjacencies:     1  # Full adjs.:      1
# Mcast floods:    5  # Mcast acks:      4

MC forwarding:   Off  DL unicast:      Off  IGMP monitor:    Off
# MC data in:      0  # MC data acc:     0  # MC data out:     0
IGMP polls snt:    0  IGMP polls rcv:    0  Unexp polls:       0
IGMP reports:      0
OSPF>
```

Interface address	Interface IP address.
Attached Area	Attached area ID.
Physical interface	Displays physical interface type and number.
Interface Mask	Interface subnet mask.
Interface type	Can be either Brdcst (broadcast, e.g., an Ethernet interface), P-P (a point-to-point network, e. g. a synchronous serial line), Multi (non-broadcast multi-access, e.g., an X.25 connection) and VLink (an OSPF virtual link).
State	Can be one of the following: 1 (down), 2 (looped back), 4 (waiting), 8 (point-to-point), 16 (DR other), 32 (backup DR) or 64 (designated router).
Designated Router	IP address of the designated router.
Backup DR	IP address of the backup designated router.
DR Priority	Priority assigned to designated router.
Hello interval	Current hello interval value.
Rxmt interval	Current retransmission interval value.
Dead interval	Current dead interval value.
TX delay	Current transmission delay value.
Poll interval	Current poll interval value.
Max pkt size	Maximum size for an OSPF packet sent out this interface.
TOS 0 cost	Interface's TOS 0 cost.
# Neighbors	Routers whose hellos have been received, plus those that have been configured.
# Adjacencies	Neighbors in state Exchange or greater.
# Full adjs.	Full adjacencies is the number or neighbors whose state is Full (and therefore, with which the router has synchronized databases).
# Mcast floods	Link state updates flooded out the interface (not counting retransmissions).
# Mcast acks	Link state acknowledgments flooded out the interface (not counting retransmissions).
MC forwarding	Displays whether multicast forwarding has been enabled for the interface.

DL unicast	Displays whether multicast datagrams are to be forwarded as data-link unicasts.
IGMP monitor	Displays whether IGMP is enabled on the interface.
# MC data in	Multicast datagrams that have been received on this interface and then successfully forwarded.
# MC data acc	Multicast datagrams that have been successfully forwarded.
# MC data out	Datagrams that have been forwarded out the interface (either as data-link multicasts or data-link unicast).
IGMP polls snt	IGMP Host Membership Queries that have been sent out the interface.
IGMP polls rcv	IGMP Host Membership Queries that have been received on the interface.
Unexp polls	IGMP Host Membership Queries that have been received on the interface that were unexpected (i.e. received when the router itself was sending them).
IGMP reports	IGMP Host Membership Reports received on the interface.

1.8. NEIGHBOR summary

Use the **NEIGHBOR summary** command to display statistics and parameters related to OSPF neighbors. If no arguments are given, a single line is printed summarizing each neighbor. If a neighbor's IP address is given, detailed statistics for that neighbor will be displayed.

Syntax:

```
OSPF> NEIGHBOR <neighbor-ip-address>
```

Example 1:

```
OSPF> NEIGHBOR
Neighbor addr    Neighbor ID      State  LSrxl DBsum LSreq  Ifc
192.7.1.254      192.7.1.254      128    0     0     0      Eth/0
192.3.1.1        0.0.0.0          1      0     0     0      FR/0
OSPF>
```

Neighbor addr	Displays the neighbor address.
Neighbor ID	Displays the neighbor's OSPF router ID.
Neighbor State	Can be one of the following: 1 (Down), 2 (Attempt), 4 (Init), 8 (2-Way), 16 (ExStart), 32 (Exchange), 64 (Loading) or 128(Full).
LSrxl	Size of the current link state retransmission list for this neighbor.
DBsum	Size of the database summary list waiting to be sent to the neighbor.
LSreq	Number of more recent advertisements that are being requested from the neighbor.
Ifc	Interface shared by the router and the neighbor.

Example 2:

```
OSPF> NEIGHBOR 128.185.184.34
Neighbor IP address:    192.7.1.254
OSPF Router ID:         192.7.1.254
Neighbor State:         128
Physical interface:     Eth/0
DR choice:              192.7.1.254
Backup choice:          192.7.1.253
DR Priority:            1
Nbr options:            E

DB summ qlen:      0  LS rxmt qlen:      0  LS req qlen:       0
Last hello:        2

# LS rxmits:       0  # Direct acks:     0  # Dup LS rcvd:     0
# Old LS rcvd:     0  # Dup acks rcv:    1  # Nbr losses:      0
# Adj. resets:     0
OSPF>
```

The meaning of each field is:

Neighbor IP address	Neighbor IP address.
OSPF router ID	Neighbor's OSPF router ID.
Neighbor State	Can be one of the following: 1(Down), 2 (Attempt), 4(Init), 8 (2-Way), 16 (ExStart), 32 (Exchange), 64 (Loading) or 128 (Full).
Physical interface	Displays physical interface type and number of the router and neighbor's common network.
DR choice	Indicate the value seen in the last hello received from the neighbor.
Backup choice	Indicate the value seen in the last hello received from the neighbor.
DR Priority	Indicate the value seen in the last hello received from the neighbor.
Nbr options	Indicates the optional OSPF capabilities supported by the neighbor. These capabilities are denoted by E (processes type 5 externals; when this is not set the area to which the common network belongs has been configured as a stub), T (can route based on TOS). This field is valid only for those neighbors in state Exchange or greater.
DB summ qlen	Indicates the number of advertisements waiting to be summarized in Database Description packets. It should be zero except when the neighbor is in state Exchange.
LS rxmt qlen	Indicates the number of advertisements that have been flooded to the neighbor, but not yet acknowledged.
LS req qlen	Indicates the number of advertisements that are being requested from the neighbor in state Loading.
Last hello	Indicates the number of seconds since a hello has been received from the neighbor.
# LS rxmits	Indicates the number of retransmissions that have occurred during flooding.
# Direct acks	Indicates responses to duplicate link state advertisements.
# Dup LS rcvd	Indicates the number of duplicate retransmissions that have occurred during flooding.
# Old LS rcvd	Indicates the number of old advertisements received during flooding.
# Dup acks rcvd	Indicates the number of duplicate acknowledgments received.
# Nbr losses	Indicates the number of times the neighbor has transitioned to Down state.
# Adj. Resets	Counts entries to state ExStart.

1.9. PING address

Ping utilizes the ICMP protocol's mandatory ECHO_REQUEST datagram to elicit an ICMP ECHO_RESPONSE from the specified host or network gateway. PING is generally used to test for reachability between network nodes.

Syntax:

```
OSPF> PING <address>
```

Example:

```
OSPF> PING 10.1.155.29
PING 192.7.1.254: 56 data bytes
64 bytes from 192.7.1.254: icmp_seq=0. time=33. ms
64 bytes from 192.7.1.254: icmp_seq=1. time=8. ms
64 bytes from 192.7.1.254: icmp_seq=2. time=22. ms
64 bytes from 192.7.1.254: icmp_seq=3. time=22. ms
64 bytes from 192.7.1.254: icmp_seq=4. time=8. ms
64 bytes from 192.7.1.254: icmp_seq=5. time=8. ms
64 bytes from 192.7.1.254: icmp_seq=6. time=8. ms
64 bytes from 192.7.1.254: icmp_seq=7. time=19. ms

----192.7.1.254 PING Statistics----
8 packets transmitted, 8 packets received, 0% packet loss
round-trip (ms)  min/avg/max = 8/16/33
OSPF>
```

1.10. ROUTERS

Use the **ROUTERS** command to display routes that have been calculated by OSPF and are now present in the routing table.

> ***NOTE: The ROUTERS command does not show all known (discovered) routers. The only routers listed by the command are the border routers, used for calculating inter-area routes, and boundary routers, used for calculating external routes.***

Syntax:

```
OSPF> ROUTERS
```

Example:

```
OSPF> ROUTERS
DType  RType  Destination       Area      Cost  Next hop(s)
ASBR   SPF    128.185.142.9     0.0.0.0   1     128.185.142.9
Fadd   SPF    128.185.142.98    0.0.0.0   1     0.0.0.0
Fadd   SPF    128.185.142.7     0.0.0.0   1     0.0.0.0
Fadd   SPF    128.185.142.48    0.0.0.0   1     0.0.0.0
Fadd   SPF    128.185.142.111   0.0.0.0   1     0.0.0.0
Fadd   SPF    128.185.142.38    0.0.0.0   1     0.0.0.0
Fadd   SPF    128.185.142.11    0.0.0.0   1     0.0.0.0
BR     SPF    128.185.142.9     0.0.0.0   1     128.185.142.9
BR     SPF    128.185.142.9     0.0.0.0   2     128.185.184.114
Fadd   SPF    128.185.142.48    0.0.0.0   1     0.0.0.0
OSPF>
```

DType Indicates destination type. “Net” indicates that the destination is a network, “ASBR” indicates that the destination is an AS boundary router, and “ABR” indicates that the destination is an area border router, and “Fadd” indicates a forwarding address (for external routes).

RType Indicates route type and how the route was derived. “SPF” indicates that the route is an intra-area route (comes from the Dijkstra calculation), “SPIA” indicates that it is an inter-area route (comes from considering summary link advertisements).

Destination Destination router’s OSPF ID. For Type D entries, one of the router’s IP addresses is displayed (which corresponds to a router in another AS).

Area Area which it belongs to.

Cost Displays the route cost.

Next hop(s) Address of the next router on the path toward the destination host. A number in parentheses at the end of the column indicates the number of equal-cost routes to the destination.

1.11. SIZE

Use the **SIZE** command to display the number of LSAs currently in the link state database, categorized by type.

Syntax:

```
OSPF> SIZE
```

Example:

```
OSPF> SIZE
# Router-LSAs:            7
# Network-LSAs:           6
# Summary LSAs:           14
# Summary Router-LSAs:    2
# AS External-LSAs:       44
# Group-membership-LSAs:  21
OSPF>
```

1.12. STATISTICS

Use the **STATISTICS** command to display statistics generated by the OSPF routing protocol. The statistics indicate how well the implementation is performing, including its memory and network utilization. Many of the fields displayed are confirmation of the OSPF configuration.

Syntax:

```
OSPF> STATISTICS
```

Example:

```
OSPF> STATISTICS
S/W version:              2.1
OSPF Router ID:           192.7.1.253
External comparison:      Type 2
AS boundary capability: yes
Import external routes: RIP STA SUB
Do not aggregate
External routes cost:     5
Orig. default route:      No (0,0.0.0.0)
Default route cost:       (1, Type 2)
Default forward. addr:    0.0.0.0

Attached areas:                    2  Estimated # external routes:    500
Estimated # OSPF routers:         50  Estimated heap usage:         66800
OSPF packets rcvd:               484  OSPF packets rcvd w/ errs:        0
Transit nodes allocated:          20  Transit nodes freed:             16
LS adv. allocated:                26  LS adv. freed:                   20
Queue headers alloc:              32  Queue headers avail:             32

# Dijkstra runs:                  4  Incremental summ. updates:        0
Incremental VL updates:            0  Buffer alloc failures:            0
Multicast pkts sent:             481  Unicast pkts sent:                5
LS adv. aged out:                  0  LS adv. flushed:                  0
Incremental ext. updates:          9
OSPF>
```

The meaning of each field is:

S/W version	Displays the OSPF software revision level.
OSPF Router ID	Displays the router's OSPF ID.
External comparison	Displays the external route type used by the router when importing external routes.
AS boundary capability	Displays whether external routes will be imported.
Import external routes	Displays which external routes will be imported.
Aggregation Type	Aggregation type.
External routes cost	Additional cost for the imported routes.
Orig default route	Displays whether the router will advertise an OSPF default route. If the value is "Yes" and a non-zero number is displayed in parentheses then a default route will be advertised only when a route to the network exists.

Default route cost	Displays the cost and type of the default route (if advertised).
Default forward addr	Displays the forwarding address specified in the default route (if advertised).
Attached areas	Indicates the number of areas that the router has active interfaces to.
Estimated # external routes	Estimated number of external routes.
Estimated heap usage	Link state database size estimate in bytes.
OSPF packets rcvd	OSPF packets received.
OSPF packets rcvd w/ errs	OSPF packets received with errors.
Transit nodes	Allocated to store router links and network links advertisements.
LS adv.	Allocated to store summary link and AS external link advertisements.
Queue headers	Form lists of link state advertisements used in the flooding and database exchange processes. If the number of queue headers allocated is not equal to the number freed, database synchronization with some neighbor is in progress.
# Dijkstra runs	Indicates how many times the OSPF routing table has been calculated from scratch.
Incremental summ. updates	Indicates that the new summary link advertisements have caused the routing table to be partially rebuilt.
Incremental VL updates	Indicates that the new summary link advertisements have caused the routing table to be partially rebuilt.
Buffer alloc failures	Indicates buffer allocation failures. The OSPF system will recover from temporary lack of packet buffers.
Multicast pkts sent	Covers OSPF hello packets and packets sent during the flooding procedure.
Unicast pkts sent	Covers OSPF packet retransmissions and the Database Exchange Procedure.
LS adv. aged out	Counts the number of advertisements that have hit 60 minutes. Link state advertisements are aged out after 60 minutes. Usually they will be refreshed before this time.
LS adv. flushed	Indicates number of advertisements removed (and not replaced) from the link state database.
Incremental ext. updates.	Displays number of changes to external destinations that are incrementally installed in the routing table.

1.13. TRACEROUTE address

Traceroute works similarly to PING, and is used to test for connectivity and isolate faults in the network. Unlike PING, Traceroute records and displays the route taken to host or network gateway.

Syntax:

```
OSPF> TRACEROUTE <address>
```

Example:

```
OSPF> TRACEROUTE 10.1.151.29: 56 data bytes
1 10.1.151.29 2 ms 0 ms 0 ms
OSPF>
```

1.14. WEIGHT

Use the **WEIGHT** command to change the cost of one of the routers OSPF interfaces. This new cost is immediately flooded throughout the OSPF routing domain, causing routes to be updated accordingly.

The cost of the interface reverts to its configured cost whenever you restart or reload the router. To make the cost change permanent, reconfigure the appropriate OSPF interface after invoking the **WEIGHT** command. This command causes a new router links advertisement to be originated, unless the cost of the interface does not change.

Syntax:

```
OSPF> WEIGHT <IP-interface-address> <new-cost>
```

Example:

```
OSPF> WEIGHT 128.185.124.22 2
Interface address[]?
TOS 0 new cost [1]?
OSPF>
```

1.15. EXIT

Use the **EXIT** command to return to the previous prompt level.

Syntax:

```
OSPF> EXIT
```

Example:

```
OSPF> EXIT
Config>
```