Release Supplement: Web OS Switch Software



Release 7.1

Alteon WebSystems
Web Speed for e-Business

50 Great Oaks Boulevard
San Jose, California 95119
408-360-5500 Main
408-360-5501 Fax
www.alteonwebsystems.com

Part Number: 050109, Revision D, February 2001

Copyright 2001 Alteon WebSystems, Inc., 50 Great Oaks Boulevard, San Jose, California 95119, USA. All rights reserved. Part Number: 050109, Revision D.

This document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this document may be reproduced in any form by any means without prior written authorization of Alteon WebSystems, Inc. Documentation is provided “as is” without warranty of any kind, either express or implied, including any kind of implied or express warranty of non-infringement or the implied warranties of merchantability or fitness for a particular purpose.

U.S. Government End Users: This document is provided with a “commercial item” as defined by FAR 2.101 (Oct 1995) and contains “commercial technical data” and “commercial software documentation” as those terms are used in FAR 12.211-12.212 (Oct 1995). Government End Users are authorized to use this documentation only in accordance with those rights and restrictions set forth herein, consistent with FAR 12.211- 12.212 (Oct 1995), DFARS 227.7202 (JUN 1995) and DFARS 252.227-7015 (Nov 1995).

Alteon WebSystems, Inc. reserves the right to change any products described herein at any time, and without notice. Alteon WebSystems, Inc. assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by Alteon WebSystems, Inc. The use and purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of Alteon WebSystems, Inc.

Web OS and Alteon 708 are trademarks of Alteon WebSystems, Inc. in the United States and certain other countries. Any other trademarks appearing in this manual are owned by their respective companies.

Contents

Preface 5
What is in This Document? 5
Related Manuals on CD-ROM 6
News and Support 6

Chapter 1: Web OS 7.1 Release Notes 7
New in This Release 7
Additional Hardware Installation Instructions 7
Activating Optional Software Keys 7
1000Base-T Support 8
New Syslog Menu 8
Configurable Frame Buffer Capacity 9
Limitations 10
Known Issues 11

Chapter 2: OSPF Application Guide 13
OSPF Overview 13
- Types of OSPF Areas 14
- Types of OSPF Routing Devices 15
- Neighbors and Adjacencies 16
- The Link-State Database 16
- The Shortest Path First Tree 17
- Authentication 17
- Internal versus External Routing 17

Web OS 7.1 OSPF Implementation 18
Configurable Parameters 18
Defining Areas 19
Interface Cost 21
Electing the Designated Router and Backup 21
Summarizing Routes 21
Default Routes 22
Virtual Links 23
The Router ID 23
Authentication 24
Host Routes for Load Balancing 24
OSPF Features Not Supported in This Release 25
OSPF Configuration Examples 25
Example 1: Simple OSPF Domain 26
Example 2: Virtual Links 28
Example 3: Summarizing Routes 32
Example 4: Host Routes 34
Verifying OSPF Support 40

Chapter 3: Web OS 7.1 Command Reference 41
Information Menu 42
OSPF Information Menu 43
OSPF Area Information 43
OSPF Interface Information 44
OSPF Neighbor Information 44
Management Processor Statistics Menu 45
OSPF Statistics Menu 45
System Configuration Menu 46
Syslog Configuration Menu 47
IP Menu Configuration 48
OSPF Menu Configuration 49
OSPF Area Configuration 51
OSPF Summary Range Configuration 52
OSPF Interface Configuration 54
OSPF Virtual Link Configuration 56
OSPF Host Entry Configuration 58

Preface

What is in This Document?

This *Release Supplement* modifies and extends your Web OS 7.0 documentation for use with Web OS 7.1.17 and higher. It contains the following material:

Chapter 1, "Web OS 7.1 Release Notes," provides the latest information about new features and limitations as well as any errata for the Web OS 7.0 documentation.

Chapter 2, "OSPF Application Guide," provides an extensive overview of the primary new feature in Web OS 7.1, including detailed configuration examples.

Chapter 3, "Web OS 7.1 Command Reference," provides details for all the new and modified commands.

Please keep this information with your Alteon WebSystems product manuals.

Related Manuals on CD-ROM

Product manuals are now supplied on CD-ROM for your convenience. The CD-ROM included with your Web switch contains the following documentation for installing and configuring the Alteon 708 Web switch:

- *Alteon 700 Series Hardware Installation Guide* (part number 050055, Revision B)
- *Web OS 7.0 Command Reference* (part number 050066, Revision B)
- *Web OS 7.0 Application Guide* (part number 050108, Revision A)
- *Web OS BBI Quick Guide* (part number 050034, Revision C)

The manuals are included as PDF files which can be read and printed using the free Acrobat Reader software available from Adobe Systems Incorporated (http://www.adobe.com).

To access the manuals, open the `doclist.pdf` file located in the `manuals` directory of the CD-ROM. The `doclist.pdf` file provides a master list of the available manuals. To open the manual for your specific needs, just click on the appropriate list item.

Once the manual opens, you can navigate through the book by selecting the predefined bookmarks on the right side of the window and by scrolling through the pages.

To return to the `doclist.pdf` file from any manual, click on the “Doc List” bookmark at the top of the bookmark list.

NOTE – To obtain a set of manuals in hardcopy format, contact your Alteon WebSystems sales representative and order part number 050134.

News and Support

Alteon WebSystems online access is available at the following URL:

http://www.alteonwebsystems.com

The Web site includes product information, software updates, release notes, and white papers. The Web site also includes access to Alteon WebSystems Customer Support for accounts under warranty or that are covered by a maintenance contract.

CHAPTER 1

Web OS 7.1 Release Notes

These release notes provide the latest information regarding your Alteon WebSystems switch with Web OS 7.1 software.

New in This Release

- The Open Shortest Path First (OSPF) routing protocol is now supported (see Chapter 2, "OSPF Application Guide").
- 1000Base-T network interface modules are supported (see page 8)
- System logging commands have been consolidated under the new `/cfg/sys/syslog` menu (see page 8).
- To relieve possible congestion at the switch, a new command has been added to extend the switch frame buffer size (see page 9).
- Port trunk group restrictions have been removed: up to four ports of any type or position can now be included in the same trunk group.
- The number of static routes supported has been increased from 128 to 2048.

Additional Hardware Installation Instructions

When installing modules, use care to prevent physical interference between the top of the module being inserted and the handles of any module immediately above. In some conditions, the handles can catch on the module being inserted. Careful insertion will prevent this problem.

Activating Optional Software Keys

Before using the `/oper/swkey` command to activate optional Web OS software features, you must first purge any demonstration keys by rebooting the switch. The switch can be rebooted using the `/boot/reset` command.

1000Base-T Support

The switch supports 1000Base-T (Gigabit Ethernet over copper) using the following link option:

/cfg/port <*slot-number port-number*>**/link 1000base-t**

New Syslog Menu

A new menu has been created which consolidates the commands involved with logging system events. The new Syslog menu appears under the System Configuration menu and can be accessed with the following command:

/cfg/sys/syslog

The Syslog menu includes the following options:

```
[Syslog Menu]
      host    - Set IP address of first syslog host
      host2   - Set IP address of second syslog host
      facil   - Set facility of first syslog host
      facil2  - Set facility of second syslog host
      console - Enable/disable console output of syslog messages
      current - Display current syslog settings
```

The Syslog menu includes a number of commands which have been moved from other locations in the menu system:

Old Web OS 7.0 Menu Command	Equivalent Web OS 7.1 Menu Command
/cfg/ip/log	/cfg/sys/syslog/host
/cfg/ip/log2	/cfg/sys/syslog/host2
/cfg/ip/fac	/cfg/sys/syslog/facil
/cfg/ip/fac2	/cfg/sys/syslog/facil2

For more information, see “Syslog Configuration Menu” on page 47.

Configurable Frame Buffer Capacity

In some environments, traffic congestion can occur at the Web switch under the following conditions:

- The Web switch is placed between a server on a Gigabit port and a client on a Fast Ethernet port, *and*...
- The server sends traffic bursts of *greater* than 16K (two 8K blocks or one 16K block).

Transfers may fail or may experience unacceptable delays when congestion occurs.

To extend the frame buffer size and relieve congestion on the Web switch, the following command has been added:

```
/cfg/sys/extbuf enable|disable
```

By default, this option is disabled. Enabling the extended buffer mode increases the Web switch frame buffer to 64K, accomodating bursts of up to two 32K blocks, four 16K blocks, or eight 8K blocks.

NOTE – Enabling the extended buffer mode also reduces the layer-4/layer-7 session capacity from 256K sessions to 128K sessions. This reduces the number of simultaneous sessions that can be processed by the Web switch for virtual servers, filtering, and content intelligent switching features.

To change the frame buffer mode, perform the following steps:

1. **Use the `/cfg/sys/extbuf` command to enable or disable the extended buffer mode.**
2. **Confirm the command (`y` or `n`) when prompted.**
3. **Use the `/boot/reset` command to reboot the Web switch.**

   The Web switch will reboot using the new buffer setting.

Limitations

The following functions are not supported in this release:

- OSPF support has the following limitations:
  - Up to eight areas are supported.
  - The Web switch will not emulate an Autonomous System Border Router (ASBR). External route summarization and redistributing external routes are not supported, though external routes can be distributed into a Not-So-Stubby Area (NSSA).
  - Filtering OSPF routes is not supported
  - Equal cost route load balancing is not supported
  - Using OSPF to forward multicast routes is not supported
- Half-duplex Ethernet (including connected hubs) is not supported.
- The `pbind` mode of operation is not supported.
- Jumbo Frames are not supported.
- Flow control is not supported.
- Mirroring is not supported.
- Class of Service (COS) is not supported.
- Network Address Translation (NAT) is not supported in a firewall load balancing environment.
- Direct Access Mode (DAM) is not supported in a firewall load balancing environment.
- Scripted health checking is not supported.
- Hot swapping is limited:
  - Hot swapping of power supplies and fans is supported.
  - Hot swapping of any other module may result in unstable operation of the switch or could cause the switch to stop processing. Removal or insertion of modules with the power on is not supported, but it will not damage the switch equipment.

Known Issues

The following issues are reported with this release:

- During initialization, some or all Gigabit SX ports will generate line errors, which may be detected by some test equipment. (This condition is common for Gigabit Ethernet equipment.).
- When displaying routing information in some OSPF configurations, a false IP address of 0.0.0.0 may be shown as the next hop or destination for some routes. This false address only appears for summary routes or host routes being injected into a non-backbone area through a virtual link. Although IP address 0.0.0.0 is normally considered invalid, it is merely a display problem in this case and does not affect OSPF routing or performance.

CHAPTER 2

OSPF Application Guide

Web OS 7.1 supports the Open Shortest Path First (OSPF) routing protocol. Alteon WebSystems' implementation conforms to the OSPF version 2 specifications detailed in Internet RFC 1583. The following sections discuss OSPF support for the Alteon 708 Web switch:

- "OSPF Overview" on page 13
- "Web OS 7.1 OSPF Implementation" on page 18
- "OSPF Configuration Examples" on page 25

OSPF Overview

OSPF is designed for routing traffic within a single IP domain called an Autonomous System (AS). The AS can be divided into smaller logical units known as *areas*.

All routing devices maintain link information in their own Link State Database (LSDB). The LSDB for all routing devices within an area is identical but is not exchanged between different areas. Only routing updates are exchanged between areas, thereby significantly reducing the overhead for maintaining routing information on a large, dynamic network.

The following sections describe key OSPF concepts.

Types of OSPF Areas

An AS can be broken into logical units known as *areas*. In any AS with multiple areas, one area must be designated as area 0, known as the *backbone*. The backbone acts as the central OSPF area. All other areas in the AS must be connected to the backbone. Areas inject summary routing information into the backbone, which then distributes it to other areas as needed.

As shown in Figure 1, OSPF defines the following types of areas:

- Stub Area—an area that is connected to only one other area. External route information is not distributed into stub areas.
- Not-So-Stubby-Area (NSSA)—similar to a stub area with additional capabilities. Routes originating from within the NSSA can be propagated to adjacent transit and backbone areas. External routes from outside the AS can be advertised within the NSSA but are not distributed into other areas.
- Transit Area—an area that allows area summary information to be exchanged between routing devices. The backbone (area 0), any area that contains a virtual link to connect two areas, and any area that is not a stub area or an NSSA are considered transit areas.

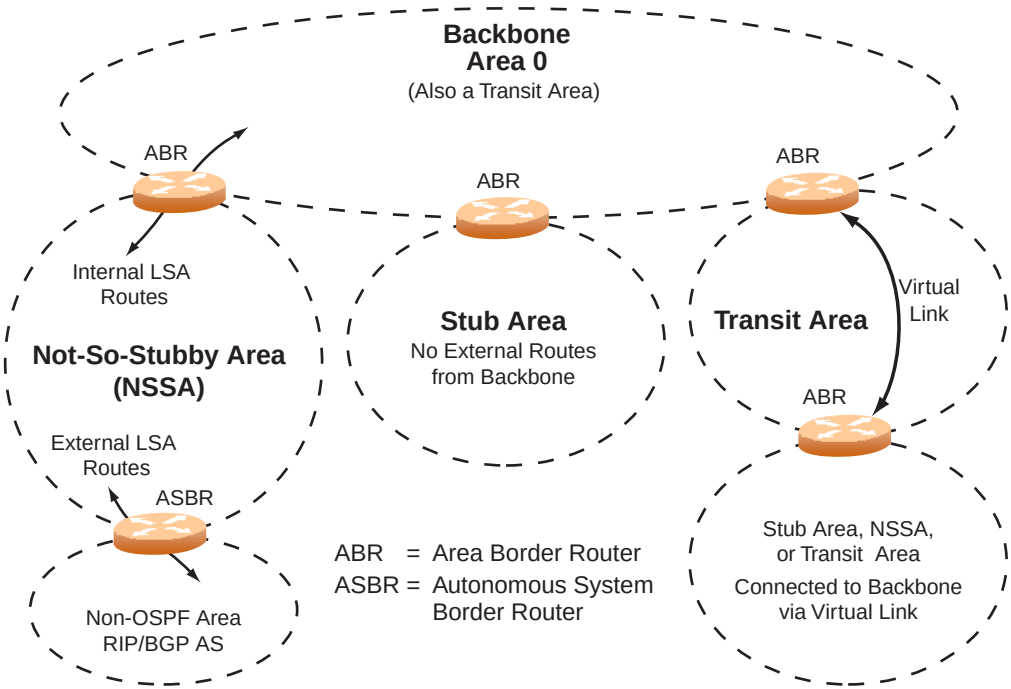


Figure 1 OSPF Area Types

Types of OSPF Routing Devices

As shown in Figure 2, OSPF uses the following types of routing devices:

- Internal Router (IR)—a router that has all of its interfaces within the same area. IRs maintain LSDBs identical to those of other routing devices within the local area.
- Area Border Router (ABR)—a router that has interfaces in multiple areas. ABRs maintain one LSDB for each connected area and disseminate routing information between areas.
- Autonomous System Border Router (ASBR)—a router that acts as a gateway between the OSPF domain and non-OSPF domains, such as IGRP, EIGRP, IS-IS, RIP, BGP, and static routes.

NOTE – At this time, an Alteon 708 Web switch can *not* be used as an ASBR.

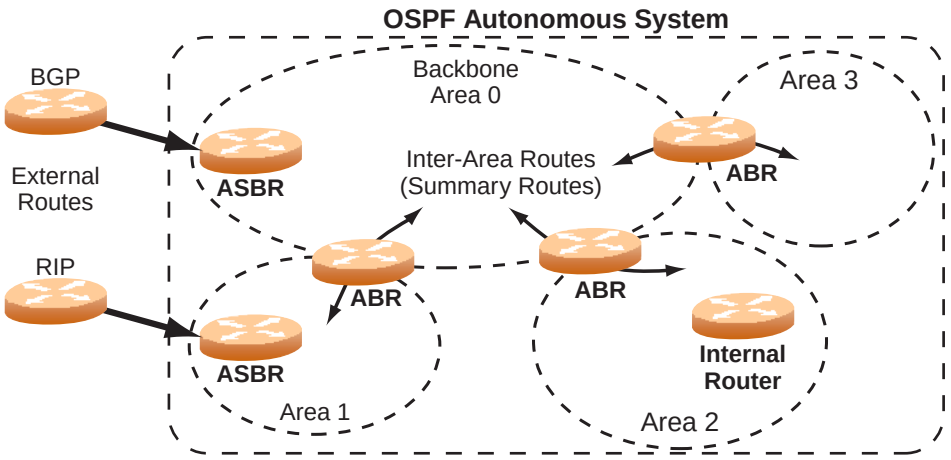


Figure 2 OSPF Domain and an Autonomous System

Neighbors and Adjacencies

In areas with more than two routing devices, *neighbors* and *adjacencies* are formed.

Neighbors are routing devices that maintain information about each others' health. To establish neighbor relationships, routing devices periodically send hello packets on each of their interfaces. All routing devices that share a common network segment, appear in the same area, and have the same health parameters (`hello` and `dead` intervals) and authentication parameters respond to each others hello packets and become neighbors. Neighbors continue to send periodic hello packets to advertise their health to neighbors. In turn, they listen to hello packets to determine the health of their neighbors and to establish contact with new neighbors.

Adjacencies are neighbors that exchange OSPF database information. In order to limit the number of database exchanges, not all neighbors in an area become adjacent to each other. Instead, the hello process is used for electing one of the neighbors as the area's Designated Router (DR) and one as the area's Backup Designated Router (BDR).

The DR is adjacent to all other neighbors and acts as the central contact for database exchanges. Each neighbor sends its database information to the DR, which relays the information to the other neighbors.

Because of the overhead required for establishing a new DR in case of failure, the hello process also elects a Backup Designated Router (BDR). The BDR is adjacent to all other neighbors (including the DR). Each neighbor sends its database information to the BDR just as with the DR, but the BDR merely stores this data and does not distribute it. If the DR fails, the BDR will take over the task of distributing database information to the other neighbors.

The Link-State Database

OSPF is a link-state routing protocol. A *link* represents an interface (or routable path) from the routing device. By establishing an adjacency with the DR, each routing device in an OSPF area maintains an identical Link-State Database (LSDB) describing the network topology for its area.

Each routing device transmits a Link-State Advertisement (LSA) on each of its interfaces. LSAs are entered into the LSDB of each routing device. OSPF uses *flooding* to distribute LSAs between routing devices.

When LSAs result in changes to the routing device's LSDB, the routing device forwards the changes to the adjacent neighbors (the DR and BDR) for distribution to the other neighbors.

OSPF routing updates occur only when changes occur, instead of periodically. For each new route, if an adjacency is interested in that route (for example, if configured to receive static routes and the new route is indeed static), an update message containing the new route is sent to the adjacency. For each route removed from the route table, if the route has already been sent to a adjacency, an update message containing the route to withdraw is sent.

The Shortest Path First Tree

The routing devices use a link-state algorithm (Dijkstra's algorithm) to calculate the shortest path to all known destinations, based on the cumulative *cost* required to reach the destination.

The cost of an individual interface in OSPF is an indication of the overhead required to send packets across it. The cost is inversely proportional to the bandwidth of the interface. A lower cost indicates a higher bandwidth.

Authentication

OSPF also allows packet authentication and uses IP multicast when sending and receiving packets. This ensures less processing on routing devices that are not listening to OSPF packets.

Internal versus External Routing

To ensure effective processing of network traffic, every routing device on your network needs to know how to send a packet (directly or indirectly) to any other location/destination in your network. This is referred to as *internal routing* and can be done with static routes or using active internal routing protocols such as OSPF, RIP, or RIPv2.

It is also useful to tell routers outside your network (upstream providers or *peers*) about the routes you have access to in your network. Sharing of routing information between autonomous systems is known as *external routing*.

Typically, an AS will have one or more border routers (peer routers that exchange routes with other OSPF networks) as well as an internal routing system enabling every router in that AS to reach every other router and destination within that AS.

When a routing device *advertises* routes to border routers on other autonomous systems, it is effectively committing to carry data to the IP space represented in the route being advertised. For example, if the routing device advertises 192.204.4.0/24, it is declaring that if another router sends data destined for any address in the 192.204.4.0/24 range, it will carry that data to its destination.

Web OS 7.1 OSPF Implementation

The following sections describe issues specific to the OSPF implementation in Web OS 7.1:

- "Configurable Parameters" on page 18
- "Defining Areas" on page 19
- "Interface Cost" on page 21
- "Electing the Designated Router and Backup" on page 21
- "Summarizing Routes" on page 21
- "Default Routes" on page 22
- "Virtual Links" on page 23
- "The Router ID" on page 23
- "Authentication" on page 24
- "Host Routes for Load Balancing" on page 24
- "OSPF Features Not Supported in This Release" on page 25

Configurable Parameters

In Web OS 7.1, OSPF parameters can be configured through the Command Line Interface (CLI). The CLI supports the following parameters: interface output cost, interface priority, dead and hello intervals, retransmission interval, and interface transmit delay.

In addition to the above parameters, you can also specify the following:

- Shortest Path First (SPF) interval—Time interval between successive calculations of the shortest path tree using the Dijkstra's algorithm.
- Stub area metric—A stub area can be configured to send a numeric metric value such that all routes received via that stub area carry the configured metric to potentially influence routing decisions.
- Default routes—Default routes with weight metrics can be manually injected into transit areas. This helps establish a preferred route when multiple routing devices exist between two areas. It also helps route traffic to external networks.

NOTE – At this time, Web OS 7.1 does not support configuration of OSPF parameters through the Web OS Browser-Based Interface (BBI) or through SNMP.

Defining Areas

If you are configuring multiple areas in your OSPF domain, one of the areas must be designated as area 0, known as the *backbone*. The backbone is the central OSPF area and is usually physically connected to all other areas. The areas inject routing information into the backbone which, in turn, disseminates the information into other areas.

Since the backbone connects the areas in your network, it must be a contiguous area. If the backbone is partitioned (possibly as a result of joining separate OSPF networks), parts of the AS will be unreachable, and you will need to configure *virtual links* to reconnect the partitioned areas (see "Virtual Links" on page 23).

Up to eight OSPF areas can be connected to an Alteon 708 with Web OS 7.1 software. To configure an area, the OSPF number must be defined and then attached to a network interface on the Web switch. The full process is explained in the following sections.

An OSPF area is defined by assigning ***two*** pieces of information—an *area index* and an *area ID*. The command to define an OSPF area is as follows:

`/cfg/ip/ospf/aindex` <*area index*>**`/areaid`** <*n.n.n.n*>

NOTE – The `aindex` option above is an arbitrary index used only on the switch and does not represent the actual OSPF area number. The actual OSPF area number is defined in the `areaid` portion of the command as will be explained below.

Using the Area ID to Assign the OSPF Area Number

The OSPF area number is defined in the `areaid` <*n.n.n.n*> option. The octet format is used in order to be compatible with two different systems of notation used by other OSPF network vendors. There are two valid ways to designate an area ID:

- Placing the area number in the last octet (0.0.0.*n*)

  Most common OSPF vendors express the area ID number as a single number. For example, the Cisco IOS-based router command "`network 1.1.1.0 0.0.0.255 area 1`" defines the area number simply as "`area 1`". On an Alteon 708 Web switch, using the last octet in the area ID, "`area 1`" is equivalent to "`areaid 0.0.0.1`".

- Multi-octet (*n.n.n.n*)

  Some OSPF vendors express the area ID number in multi-octet format. For example, "`area 2.2.2.2`" represents OSPF area 2 and can be specified directly on an Alteon 708 Web switch as "`areaid 2.2.2.2`".

NOTE – Although both types of area ID formats are supported, be sure that the area IDs are in the same format throughout an area.

Assigning the Area Index

The `aindex` <*area index*> option is actually just an arbitrary index (0-7) used only by the Web switch. This index does not necessarily represent the OSPF area number, though for configuration simplicity, it should where possible.

For example, both of the following sets of commands define OSPF area 0 (the backbone) and area 1 because that information is held in the area ID portion of the command. However, the first set of commands is easier to configure and maintain because the arbitrary area indexes agree with the area IDs:

- Area index and area ID agree

  `/cfg/ip/ospf/aindex 0/areaid 0.0.0.0` *(Use index 0 to set area 0 in ID octet format)*
  `/cfg/ip/ospf/aindex 1/areaid 0.0.0.1` *(Use index 1 to set area 1 in ID octet format)*

- Also valid, the area index can be set to an arbitrary value

  `/cfg/ip/ospf/aindex 4/areaid 0.0.0.0` *(Use index 4 to set area 0 in ID octet format)*
  `/cfg/ip/ospf/aindex 5/areaid 0.0.0.1` *(Use index 5 to set area 1 in ID octet format)*

Attaching an Area to a Network

Once an OSPF area has been defined, it must be associated with a network. To attach the area to a network, you must assign the OSPF area index to an IP interface that participates in the area. The format for the command is as follows:

`/cfg/ip/ospf/if` <*interface number*>**`/aindex`** <*area index*>

For example, the following commands could be used to configure IP interface 14 for a presence on the 10.10.10.1/24 network, to define OSPF area 1, and to attach the area to the network:

```
>> # /cfg/ip/if 14                          (Select menu for IP interface 14)
>> IP Interface 14# addr 10.10.10.1         (Define IP address on backbone network)
>> IP Interface 14# enable                  (Enable IP interface 14)
>> IP Interface 14# ../ospf/aindex 1        (Select menu for area index 1)
>> OSPF Area (index) 1 # areaid 0.0.0.1     (Define area ID as OSPF area 1)
>> OSPF Area (index) 1 # enable             (Enable area index 1)
>> OSPF Area (index) 1 # ../if 14           (Select OSPF menu for interface 14)
>> OSPF Interface 14# aindex 1              (Attach area to network on interface 14)
>> OSPF Interface 14# enable                (Enable interface 14 for area index 1)
```

NOTE – Attaching the Web switch to an area requires only one IP interface. If multiple network ranges exist within an area, you do not need to configure an IP interface for each range.

Interface Cost

The OSPF link-state algorithm (Dijkstra's algorithm) places each routing device at the root of a tree and determines the cumulative *cost* required to reach each destination. Usually, the cost is inversely proportional to the bandwidth of the interface. Low cost indicates high bandwidth. You can manually enter the cost for the output route with the following command:

`/cfg/ip/ospf/if` *<OSPF interface number>* **`/cost`** *<cost value (1-65535)>*

Electing the Designated Router and Backup

In any area with more than two routing devices, a Designated Router (DR) is elected as the central contact for database exchanges among neighbors, and a Backup Designated Router (BDR) is elected in case the DR fails.

DR and BDR elections are made through the hello process. The election can be influenced by assigning a priority value to the Web switch's OSPF interfaces. The command is as follows:

`/cfg/ip/ospf/if` *<OSPF interface number>* **`/prio`** *<priority value (0-127)>*

A priority value of 127 is the highest, and 1 is the lowest. A priority value of 0 specifies that the interface cannot be used as a DR or BDR. In case of a tie, the routing device with the lowest router ID wins.

Summarizing Routes

Route summarization condense routing information. Without summarization, each routing device in an OSPF network would retain a route to every subnet in the network. With summarization, routing devices can reduce some sets of routes to a single advertisement, reducing both the load on the routing device and the perceived complexity of the network. The importance of route summarization increases with network size.

Summary routes can be defined for up to 64 IP address ranges using the following command:

`/cfg/ip/ospf/range` *<range number>* **`/addr`** *<x.x.x.x>* **`/mask`** *<z.z.z.z>*

where *<range number>* is a number between 1 and 64, *<x.x.x.x>* is the base IP address for the range, and *<z.z.z.z>* is the IP address mask for the range. For a detailed configuration example, see "Example 3: Summarizing Routes" on page 32.

NOTE – Because an Alteon 708 Web switch does not emulate an ASBR, it does not support OSPF external summary routes.

Default Routes

When an OSPF routing device encounters traffic for a destination address it does not recognize, it forwards that traffic along the *default route*. Typically, the default route leads upstream toward the backbone until it reaches the intended area or an external router.

Each Alteon 708 Web switch acting as an ABR automatically inserts a default route into each attached area. In simple OSPF stub areas or NSSAs with only one ABR leading upstream (see area 1 in Figure 3), any traffic for IP address destinations outside the area will be forwarded to the switch's IP interface, and then into the connected transit area (usually the backbone). Since this is automatic, no further configuration is required for such areas.

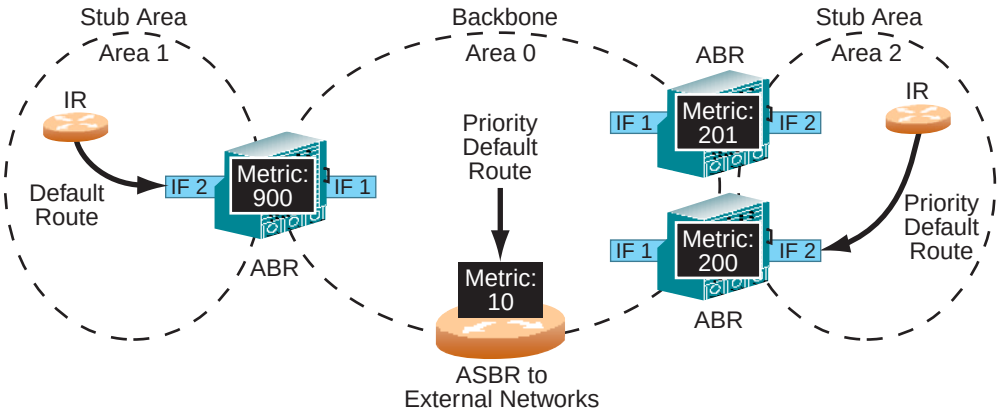


Figure 3 Injecting Default Routes

In more complex OSPF areas with multiple ABRs or ASBRs (such as area 0 and area 2 in Figure 3), there are multiple routes leading from the area. In such areas, traffic for unrecognized destinations cannot tell which route leads upstream without further configuration.

To resolve the situation and select one default route among multiple choices in an area, you can manually configure a metric value on each ABR. The metric assigns a priority to the ABR for its selection as the priority default route in an area. The following command is used for setting the metric value:

`/cfg/ip/ospf/default` *<metric value> <metric type (1 or 2)>*

where *<metric value>* is any number between 1 (highest priority for default route) and 16,777,215 (lowest priority).

To clear a default route metric from the switch, use the following command:

`/cfg/ip/ospf/default none`

Virtual Links

Usually, all areas in an OSPF AS are physically connected to the backbone. In some cases where this is not possible, you can use a *virtual link.* Virtual links are created to connect one area to the backbone through another non-backbone area (see Figure 1 on page 14).

The area which contains a virtual link must be a transit area and have full routing information. Virtual links cannot be configured inside a stub area or NSSA. The area type must be defined as `transit` using the following command:

`/cfg/ip/ospf/aindex` *<area index>***`/type transit`**

The virtual link must be configured on the routing devices at each endpoint of the virtual link, though they may traverse multiple routing devices. To configure an Alteon 708 Web switch as one endpoint of a virtual link, use the following command:

`/cfg/ip/ospf/virt` *<link number>***`/aindex`** *<area index>***`/nbr`** *<x.x.x.x>*

where *<link number>* is a value between 1 and 8, *<area index>* is the OSPF area index of the transit area, and *<x.x.x.x>* is the router ID of the virtual neighbor (the routing device at the target endpoint).

To provide the Alteon 708 Web switch with a router ID, which will be needed when configuring a virtual link in the other direction, the following command is used:

`/cfg/ip/rtrid` *<x.x.x.x>*

For a detailed configuration example, see "Example 2: Virtual Links" on page 28.

The Router ID

Routing devices in OSPF areas are identified by a router ID. The router ID is expressed in IP address format. The IP address of the router ID is not required to be included in any IP interface range or in any OSPF area.

The router ID can be configured in one of the following two ways:

- Statically—Use the following command to manually configure the router ID:

  `/cfg/ip/rtrid` *<x.x.x.x>*

- Dynamically—OSPF protocol configures the lowset active IP interface IP address as the router ID. This is the default. To use a dynamic router ID after having set it statically, set the router ID to 0.0.0.0 and reboot the Web switch.

Authentication

OSPF protocol exchanges are authenticated so that only trusted routing devices can participate. Web OS 7.1 supports only simple authentication (type 1 plain text passwords) among neighboring routing devices within an area.

OSPF passwords are enabled or disabled on the Web switch using the following commands:

`/cfg/ip/ospf/auth password` *(Turn on OSPF password authentication)*
`/cfg/ip/ospf/auth none` *(Turn off OSPF password authentication)*

Plain text passwords up to eight characters long can be configured individually for each OSPF IP interface and virtual link using the following commands:

`/cfg/ip/ospf/if` *<OSPF interface number>* **`/key`** *<password>*
`/cfg/ip/ospf/virt` *<link number>* **`/key`** *<password>*

Host Routes for Load Balancing

The Web OS 7.1 implementation of OSPF includes host routes. Host routes are used for advertising network device IP addresses to external networks, accomplishing the following goals:

- Server Load Balancing (SLB) within OSPF

  Host routes advertise virtual server IP addresses to external networks. This allows for standard SLB between the Web switch and the server pools in an OSPF environment. For more information on SLB, see your Web OS *Application Guide* and *Command Reference*.

- ABR Load Sharing

  As a second form of load balancing, host routes can be used for dividing OSPF traffic among multiple ABRs. To accomplish this, each Web switch provides identical services but advertises a host route for a different virtual server IP address to the external network. If each virtual server IP address serves a different and equal portion of the external world, incoming traffic from the upstream router should be split evenly among ABRs.

- ABR Failover

  Complementing ABR load sharing, identical host routes are configured on each ABR. These are all given different costs so that a different ABR is selected as the preferred route for each virtual server and the others are available as backups for failover purposes.

For a configuring example, see "Example 4: Host Routes" on page 34.

NOTE – If redundant routes via multiple routing processes (such as OSPF, RIP, BGP, or static routes) exist on your network, the Web switch will default to the OSPF-derived route.

OSPF Features Not Supported in This Release

The following OSPF features are not supported in this release:

- External route summarization
- Redistributing external routes (you can distribute external routes into an NSSA, but the Web switch will not emulate an ASBR)
- Filtering OSPF routes
- Configuring equal cost route load balancing
- Using OSPF to forward multicast routes

OSPF Configuration Examples

As covered in detail in the following example sections, here are the basic steps for using the CLI to configure OSPF support on the Web switch:

1. **Configure IP interfaces.**

   One IP interface is required for each desired network (range of IP addresses) being assigned to an OSPF area on the Web switch.

2. **(Optional) Configure the router ID.**

   The router ID is required only when configuring virtual links on the Web switch.

3. **Enable OSPF on the switch.**
4. **Define the OSPF areas.**
5. **Configure OSPF interface parameters.**

   IP interfaces are used for attaching networks to the various areas.

6. **(Optional) Configure route summarization between OSPF areas.**
7. **(Optional) Configure route summarization when redistributing routes into OSPF.**
8. **(Optional) Configure virtual links.**
9. **(Optional) Configure host routes.**

NOTE – At this time, Web OS 7.1 does not support configuration of OSPF parameters through the Web OS Browser-Based Interface (BBI) or through SNMP.

Example 1: Simple OSPF Domain

In this example, two OSPF areas are defined—one area is the backbone and the other is a stub area. A stub area does not allow advertisements of external routes, thus reducing the size of the database. Instead, a default summary route of IP address 0.0.0.0 is automatically inserted into the stub area. Any traffic for IP address destinations outside the stub area will be forwarded to the stub area's IP interface, and then into the backbone.

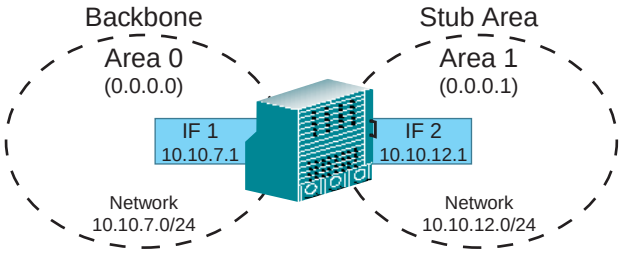


Figure 4 A Simple OSPF Domain

Follow this procedure to configure OSPF support as shown in Figure 4:

1. **Configure IP interfaces on each network that will be attached to OSPF areas.**

   In this example, two IP interfaces are needed: one for the backbone network on 10.10.7.0/24 and one for the stub area network on 10.10.12.0/24.

```
>> # /cfg/ip/if 1                    (Select menu for IP Interface 1)
>> IP Interface 1 # addr 10.10.7.1   (Set IP address on backbone network)
>> IP Interface 1 # enable           (Enable IP interface 1)
>> IP Interface 1 # ../if 2          (Select menu for IP Interface 2)
>> IP Interface 2 # addr 10.10.12.1  (Set IP address on stub area network)
>> IP Interface 2 # enable           (Enable IP interface 2)
```

2. **Enable OSPF.**

```
>> IP Interface 2 # /cfg/ip/ospf/on   (Enable OSPF on the Web switch)
```

3. **Define the backbone.**

   The backbone is always configured as a transit area using `areaid 0.0.0.0`.

```
>> Open Shortest Path First # aindex 0          (Select menu for area index 0)
>> OSPF Area (index) 0 # areaid 0.0.0.0        (Set the ID for backbone area 0)
>> OSPF Area (index) 0 # type transit          (Define backbone as transit type)
>> OSPF Area (index) 0 # enable                (Enable the area)
```

4. **Define the stub area.**

```
>> OSPF Area (index) 0 # ../aindex 1          (Select menu for area index 1)
>> OSPF Area (index) 1 # areaid 0.0.0.1       (Set the area ID for OSPF area 1)
>> OSPF Area (index) 1 # type stub            (Define area as stub type)
>> OSPF Area (index) 1 # enable               (Enable the area)
```

5. **Attach the network interface to the backbone.**

```
>> OSPF Area 1 # ../if 1                      (Select OSPF menu for IP interface 1)
>> OSPF Interface 1 # aindex 0                (Attach network to backbone index)
>> OSPF Interface 1 # enable                  (Enable the backbone interface)
```

6. **Attach the network interface to the stub area.**

```
>> OSPF Interface 1 # ../if 2                 (Select OSPF menu for IP interface 2)
>> OSPF Interface 2 # aindex 1                (Attach network to stub area index)
>> OSPF Interface 2 # enable                  (Enable the stub area interface)
```

7. **Apply and save the configuration changes.**

```
>> OSPF Interface 2 # apply                   (Global command to apply all changes)
>> OSPF Interface 2 # save                    (Global command to save all changes)
```

Example 2: Virtual Links

In the example shown in Figure 5, area 2 is not physically connected to the backbone as is usually required. Instead, area 2 will be connected to the backbone via a virtual link through area 1. The virtual link must be configured at each endpoint.

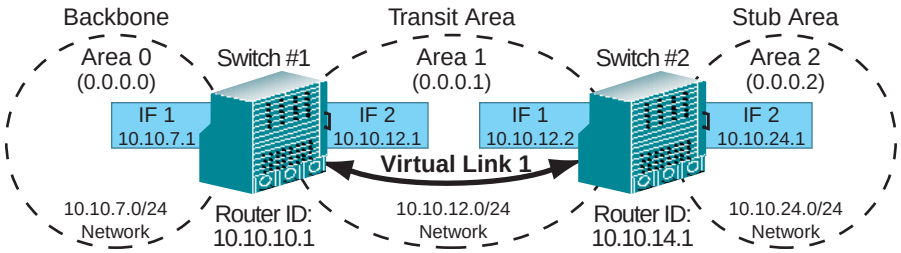


Figure 5 Configuring a Virtual Link

Configuring OSPF for a Virtual Link on Switch #1

1. **Configure IP interfaces on each network that will be attached to the switch.**

   In this example, two IP interfaces are needed on Switch #1: one for the backbone network on 10.10.7.0/24 and one for the transit area network on 10.10.12.0/24.

```
>> # /cfg/ip/if 1                    (Select menu for IP Interface 1)
>> IP Interface 1 # addr 10.10.7.1   (Set IP address on backbone network)
>> IP Interface 1 # enable           (Enable IP interface 1)
>> IP Interface 1 # ../if 2          (Select menu for IP Interface 2)
>> IP Interface 2 # addr 10.10.12.0  (Set IP address on transit area network)
>> IP Interface 2 # enable           (Enable interface 2)
```

2. **Configure the router ID.**

   A router ID is required when configuring virtual links. Later, when configuring the other end of the virtual link on Switch #2, the router ID specified here will be used as the `nbr` (target virtual neighbor) address.

```
>> IP Interface 2 # /cfg/ip/rtrid 10.10.10.1  (Set static router ID on Web switch #1)
```

3. **Enable OSPF.**

```
>> IP # /cfg/ip/ospf/on                (Enable OSPF on Web switch #1)
```

4. **Define the backbone.**

```
>> Open Shortest Path First # aindex 0          (Select menu for area index 0)
>> OSPF Area (index) 0 # areaid 0.0.0.0        (Set the area ID for backbone area 0)
>> OSPF Area (index) 0 # type transit          (Define backbone as transit type)
>> OSPF Area (index) 0 # enable                (Enable the area)
```

5. **Define the transit area.**

The area that contains the virtual link must be configured as a transit area.

```
>> OSPF Area (index) 0 # ../aindex 1           (Select menu for area index 1)
>> OSPF Area (index) 1 # areaid 0.0.0.1        (Set the area ID for OSPF area 1)
>> OSPF Area (index) 1 # type transit          (Define area as transit type)
>> OSPF Area (index) 1 # enable                (Enable the area)
```

6. **Attach the network interface to the backbone.**

```
>> OSPF Area (index) 1 # ../if 1               (Select OSPF menu for IP interface 1)
>> OSPF Interface 1 # aindex 0                 (Attach network to backbone index)
>> OSPF Interface 1 # enable                   (Enable the backbone interface)
```

7. **Attach the network interface to the transit area.**

```
>> OSPF Interface 1 # ../if 2                  (Select OSPF menu for IP interface 2)
>> OSPF Interface 2 # aindex 1                 (Attach network to transit area index)
>> OSPF Interface 2 # enable                   (Enable the transit area interface)
```

8. **Configure the virtual link.**

```
>> OSPF Interface 2 # ../virt 1                (Specify a virtual link number)
>> OSPF Virtual Link 1 # aindex 1              (Specify the transit area for the virtual link)
>> OSPF Virtual Link 1 # nbr 10.10.14.1        (Specify the router ID of the recipient)
>> OSPF Virtual Link 1 # enable                (Enable the virtual link)
```

The `nbr` router ID configured in this step must be the same as the router ID that will be configured for Switch #2 in Step 2 on page 30.

9. **Apply and save the configuration changes.**

```
>> OSPF Interface 2 # apply                    (Global command to apply all changes)
>> OSPF Interface 2 # save                     (Global command to save all changes)
```

Configuring OSPF for a Virtual Link on Switch #2

1. **Configure IP interfaces on each network that will be attached to OSPF areas.**

   Two IP interfaces are needed on Switch #2: one for the transit area network on 10.10.12.0/24 and one for the stub area network on 10.10.24.0/24.

```
>> # /cfg/ip/if 1                          (Select menu for IP Interface 1)
>> IP Interface 1 # addr 10.10.12.2        (Set IP address on transit area network)
>> IP Interface 1 # enable                 (Enable IP interface 1)
>> IP Interface 1 # ../if 2                (Select menu for IP Interface 2)
>> IP Interface 2 # addr 10.10.24.1        (Set IP address on stub area network)
>> IP Interface 2 # enable                 (Enable IP interface 2)
```

2. **Configure the router ID.**

   A router ID is required when configuring virtual links. This router ID should be same one specified as the `nbr` (target virtual neighbor) on Web switch #1 in Step 8 on page 29.

```
>> IP Interface 2 # /cfg/ip/rtrid 10.10.14.1  (Set static router ID on Web switch #2)
```

3. **Enable OSPF.**

```
>> IP# /cfg/ip/ospf/on                    (Enable OSPF on Web switch #2)
```

4. **Define the backbone.**

   This version of Web OS 7.1 requires that a backbone index be configured on the non-backbone end of the virtual link as follows:

```
>> Open Shortest Path First # aindex 0     (Select the menu for area index 0)
>> OSPF Area (index) 0 # areaid 0.0.0.0    (Set the area ID for OSPF area 0)
>> OSPF Area (index) 0 # enable            (Enable the area)
```

5. **Define the transit area.**

```
>> OSPF Area (index) 0 # ../aindex 1       (Select menu for area index 1)
>> OSPF Area (index) 1 # areaid 0.0.0.1    (Set the area ID for OSPF area 1)
>> OSPF Area (index) 1 # type transit      (Define area as transit type)
>> OSPF Area (index) 1 # enable            (Enable the area)
```

6. **Define the stub area.**

```
>> OSPF Area (index) 1 # ../aindex 2          (Select the menu for area index 2)
>> OSPF Area (index) 2 # areaid 0.0.0.2       (Set the area ID for OSPF area 2)
>> OSPF Area (index) 2 # type stub            (Define area as stub type)
>> OSPF Area (index) 2 # enable               (Enable the area)
```

7. **Attach the network interface to the backbone.**

```
>> OSPF Area (index) 2 # ../if 1              (Select OSPF menu for IP interface 1)
>> OSPF Interface 1 # aindex 1                (Attach network to transit area index)
>> OSPF Interface 1 # enable                  (Enable the transit area interface)
```

8. **Attach the network interface to the transit area.**

```
>> OSPF Interface 1 # ../if 2                 (Select OSPF menu for IP interface 2)
>> OSPF Interface 2 # aindex 2                (Attach network to stub area index)
>> OSPF Interface 2 # enable                  (Enable the stub area interface)
```

9. **Configure the virtual link.**

The `nbr` router ID configured in this step must be the same as the router ID that was configured for Web switch #1 in Step 2 on page 28.

```
>> OSPF Interface 2 # ../virt 1               (Specify a virtual link number)
>> OSPF Virtual Link 1 # aindex 1             (Specify the transit area for the virtual link)
>> OSPF Virtual Link 1 # nbr 10.10.10.1       (Specify the router ID of the recipient)
>> OSPF Virtual Link 1 # enable               (Enable the virtual link)
```

10. **Apply and save the configuration changes.**

```
>> OSPF Interface 2 # apply                   (Global command to apply all changes)
>> OSPF Interface 2 # save                    (Global command to save all changes)
```

Other Virtual Link Options

- You can use redundant paths by configuring multiple virtual links.
- Only the endpoints of the virtual link are configured. The virtual link path may traverse multiple routers in an area as long as there is a routable path between the endpoints.

Example 3: Summarizing Routes

By default, ABRs advertise all the network addresses from one area into another area. Route summarization can be used for consolidating advertised addresses and reducing the perceived complexity of the network.

If the network IP addresses in an area are assigned to a contiguous subnet range, you can configure the ABR to advertise a single summary route that includes all the individual IP addresses within the area.

The following example shows one summary route from area 1 (stub area) injected into area 0 (the backbone). The summary route consists of all IP addresses from 36.128.192.0 through 36.128.254.255.

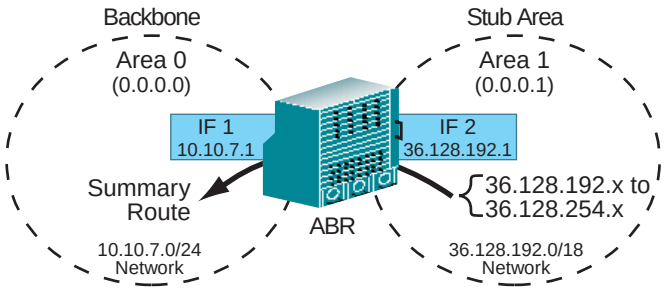


Figure 6 Summarizing Routes

NOTE – You can also specify a range of addresses to *prevent* advertising by using the `hide` option (see page 53).

Follow this procedure to configure OSPF support as shown in Figure 6:

1. **Configure IP interfaces for each network which will be attached to OSPF areas.**

```
>> # /cfg/ip/if 1                    (Select menu for IP Interface 1)
>> IP Interface 1 # addr 10.10.7.1   (Set IP address on backbone network)
>> IP Interface 1 # enable           (Enable IP interface 1)
>> IP Interface 1 # ../if 2          (Select menu for IP Interface 2)
>> IP Interface 2 # addr 36.128.192.1 (Set IP address on stub area network)
>> IP Interface 2 # enable           (Enable IP interface 2)
```

2. **Enable OSPF.**

```
>> IP Interface 2 # /cfg/ip/ospf/on   (Enable OSPF on the Web switch)
```

3. **Define the backbone.**

```
>> Open Shortest Path First # aindex 0          (Select menu for area index 0)
>> OSPF Area (index) 0 # areaid 0.0.0.0        (Set the ID for backbone area 0)
>> OSPF Area (index) 0 # type transit          (Define backbone as transit type)
>> OSPF Area (index) 0 # enable                (Enable the area)
```

4. **Define the stub area.**

```
>> OSPF Area (index) 0 # ../aindex 1           (Select menu for area index 1)
>> OSPF Area (index) 1 # areaid 0.0.0.1        (Set the area ID for OSPF area 1)
>> OSPF Area (index) 1 # type stub             (Define area as stub type)
>> OSPF Area (index) 1 # enable                (Enable the area)
```

5. **Attach the network interface to the backbone.**

```
>> OSPF Area (index) 1 # ../if 1               (Select OSPF menu for IP interface 1)
>> OSPF Interface 1 # aindex 0                 (Attach network to backbone index)
>> OSPF Interface 1 # enable                   (Enable the backbone interface)
```

6. **Attach the network interface to the stub area.**

```
>> OSPF Interface 1 # ../if 2                  (Select OSPF menu for IP interface 2)
>> OSPF Interface 2 # aindex 1                 (Attach network to stub area index)
>> OSPF Interface 2 # enable                   (Enable the stub area interface)
```

7. **Configure route summarization by specifying the starting address and mask of the range of addresses to be summarized**

```
>> OSPF Interface 2 # ../range 1               (Select menu for summary range)
>> OSPF Summary Range 1 # addr 36.128.192.0    (Set base IP address of summary range)
>> OSPF Summary Range 1 # mask 255.255.192.0   (Set mask address for summary range)
>> OSPF Summary Range 1 # aindex 0             (Inject summary route into backbone area)
>> OSPF Summary Range 1 # enable               (Enable summary range)
```

8. **Apply and save the configuration changes.**

```
>> OSPF Summary Range 1 # apply                (Global command to apply all changes)
>> OSPF Summary Range 1 # save                 (Global command to save all changes)
```

Example 4: Host Routes

The Web OS 7.1 implementation of OSPF includes host routes. Host routes are used for advertising network device IP addresses to external networks. This makes Server Load Balancing (SLB) possible within OSPF. It also make ABR load sharing and failover possible.

Consider the example network in Figure 7. Both Web switches have access to servers with identical content and are configured with the same virtual server IP addresses: 10.10.10.1 and 10.10.10.2. Web switch #1 is given a host route with a low cost for virtual server 10.10.10.1 and another host route with a high cost for virtual server 10.10.10.2. Web switch #2 is configured with the same hosts but with the costs reversed; one host route has a high cost for virtual server 10.10.10.1 and another has a low cost for virtual server 10.10.10.2.

All four host routes are injected into the upstream router and advertised externally. Traffic comes in for both virtual server IP addresses (10.10.10.1 and 10.10.10.2). The upstream router sees that both addresses exist on both Web switches and uses the host route with the lowest cost for each. Traffic for 10.10.10.1 goes to Web switch #1 because its host route has the lowest cost for that address. Traffic for 10.10.10.2 goes to Web switch #2 because its host route has the lowest cost. This effectively shares the load among ABRs. Both Web switches then use standard Server Load Balancing (SLB) to distribute traffic among available real servers.

In addition, if one of the Web switches were to fail, the upstream routing device would forward the traffic to the ABR whose host route has the next lowest cost. In this example, the remaining Web switch would assume the entire load for both virtual servers.

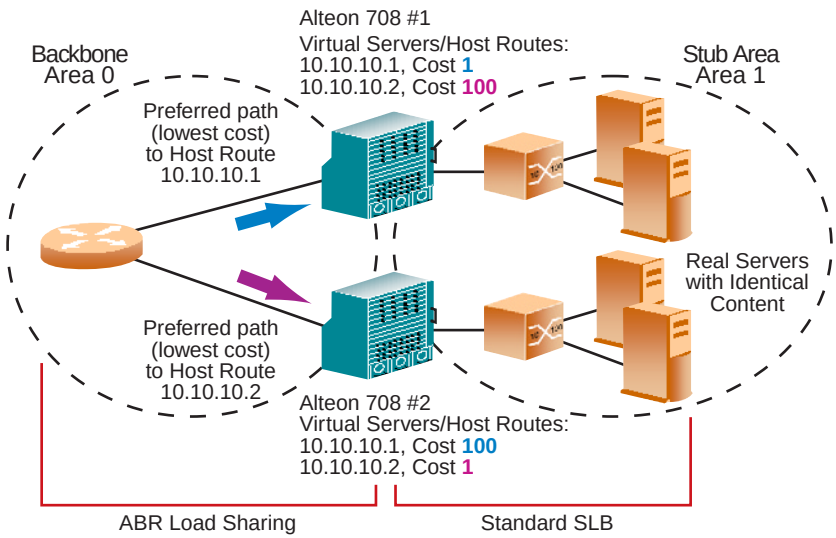


Figure 7 Configuring OSPF Host Routes

Configuring OSPF for Host Routes on Web Switch #1

1. **Configure basic SLB parameters.**

   Web switch #1 is connected to two real servers. Each real server is given an IP address and is placed in the same real server group.

```
>> # /cfg/slb/real 1                          (Select menu for real server 1)
>> Real server 1 # rip 100.100.100.25         (Set the IP address for real server 1)
>> Real server 1 # enable                     (Enable the real server)
>> Real server 1 # ../real 2                  (Select menu for real server 2)
>> Real server 2 # rip 100.100.10.26          (Set the IP address for real server 2)
>> Real server 2 # enable                     (Enable the real server)
>> Real server 2 # ../group 1                 (Select menu for real server group 1)
>> Real server group 1 # add 1                (Add real server 1 to group)
>> Real server group 1 # add 2                (Add real server 2 to group)
>> Real server group 1 # enable               (Enable the group)
>> Real server group 1 # ../on                (Turn SLB on)
```

2. **Configure the primary virtual server.**

   Web switch #1 will be preferred for virtual server 10.10.10.1

```
>> Layer 4 # virt 1                           (Select menu for virtual server 1)
>> Virtual server 1 # vip 10.10.10.1          (Set the IP address for virtual server 1)
>> Virtual server 1 # enable                  (Enable the virtual server)
>> Virtual server 1 # service http            (Select menu for service on virtual server)
>> Virtual server 1 http service # group 1    (Use real server group 1 for http service)
```

3. **Configure the backup virtual server.**

   Web switch #1 will act as a backup for virtual server 10.10.10.2. Both virtual servers in this example are configured with the same real server group and provide identical services.

```
>> Virtual server 2 http service # /cfg/slb/virt 2  (Select menu for virtual server 2)
>> Virtual server 1 # vip 10.10.10.2          (Set the IP address for virtual server 2)
>> Virtual server 1 # enable                  (Enable the virtual server)
>> Virtual server 1 # service http            (Select menu for service on virtual server)
>> Virtual server 1 # group 1                 (Use real server group 1 for http service)
```

4. **Configure IP interfaces for each network that will be attached to OSPF areas.**

```
>> Virtual server 1 # /cfg/ip/if 1          (Select menu for IP Interface 1)
>> IP Interface 1 # addr 10.10.7.1          (Set IP address on backbone network)
>> IP Interface 1 # enable                  (Enable IP interface 1)
>> IP Interface 1 # ../if 2                 (Select menu for IP Interface 2)
>> IP Interface 2 # addr 10.10.10.40        (Set IP address on stub area network)
>> IP Interface 2 # enable                  (Enable IP interface 2)
```

5. **Enable OSPF on Web switch #1.**

```
>> IP Interface 2 # ../ospf/on              (Enable OSPF on Web switch #1)
```

6. **Define the backbone.**

```
>> Open Shortest Path First # aindex 0      (Select menu for area index 0)
>> OSPF Area (index) 0 # areaid 0.0.0.0     (Set the ID for backbone area 0)
>> OSPF Area (index) 0 # type transit       (Define backbone as transit type)
>> OSPF Area (index) 0 # enable             (Enable the area)
```

7. **Define the stub area.**

```
>> OSPF Area (index) 0 # ../aindex 1        (Select menu for area index 1)
>> OSPF Area (index) 1 # areaid 0.0.0.1     (Set the ID for stub area 1)
>> OSPF Area (index) 1 # type stub          (Define area as stub type)
>> OSPF Area (index) 1 # enable             (Enable the area)
```

8. **Attach the network interface to the backbone.**

```
>> OSPF Area (index) 1 # ../if 1            (Select OSPF menu for IP interface 1)
>> OSPF Interface 1 # aindex 0              (Attach network to backbone index)
>> OSPF Interface 1 # enable                (Enable the backbone interface)
```

9. **Attach the network interface to the stub area.**

```
>> OSPF Interface 1 # ../if 2               (Select OSPF menu for IP interface 2)
>> OSPF Interface 2 # aindex 1              (Attach network to stub area index)
>> OSPF Interface 2 # enable                (Enable the stub area interface)
```

10. Configure host routes.

One host route is needed for each virtual server on Web switch #1. Since virtual server 10.10.10.1 is preferred for Web switch #1, its host route has a low cost. Because virtual server 10.10.10.2 is used as a backup in case Web switch #2 fails, its host route has a high cost.

```
>> OSPF Interface 2 # ../host 1                 (Select menu for host route 1)
>> OSPF Host Entry 1 # addr 10.10.10.1          (Set IP address same as virtual server 1)
>> OSPF Host Entry 1 # aindex 0                 (Inject host route into backbone area)
>> OSPF Host Entry 1 # cost 1                   (Set low cost for preferred path)
>> OSPF Host Entry 1 # enable                   (Enable the host route)
>> OSPF Host Entry 1 # ../host 2                (Select menu for host route 2)
>> OSPF Host Entry 2 # addr 10.10.10.2          (Set IP address same as virtual server 2)
>> OSPF Host Entry 2 # aindex 0                 (Inject host route into backbone area)
>> OSPF Host Entry 2 # cost 100                 (Set high cost for use as backup path)
>> OSPF Host Entry 2 # enable                   (Enable the host route)
```

11. Apply and save the configuration changes.

```
>> OSPF Host Entry 2 # apply                    (Global command to apply all changes)
>> OSPF Host Entry 2 # save                     (Global command to save all changes)
```

Configuring OSPF for Host Routes on Web Switch #2

1. Configure basic SLB parameters.

Web switch #2 is connected to two real servers. Each real server is given an IP address and is placed in the same real server group.

```
>> # /cfg/slb/real 1                           (Select menu for real server 1)
>> Real server 1 # rip 100.100.100.27           (Set the IP address for real server 1)
>> Real server 1 # enable                       (Enable the real server)
>> Real server 1 # ../real 2                    (Select menu for real server 2)
>> Real server 2 # rip 100.100.10.28            (Set the IP address for real server 2)
>> Real server 2 # enable                       (Enable the real server)
>> Real server 2 # ../group 1                   (Select menu for real server group 1)
>> Real server group 1 # add 1                  (Add real server 1 to group)
>> Real server group 1 # add 2                  (Add real server 2 to group)
>> Real server group 1 # enable                 (Enable the group)
>> Real server group 1 # ../on                  (Turn SLB on)
```

2. **Configure the virtual server parameters.**

   The same virtual servers are configured as on Web switch #1.

```
>> Layer 4 # virt 1                          (Select menu for virtual server 1)
>> Virtual server 1 # vip 10.10.10.1         (Set the IP address for virtual server 1)
>> Virtual server 1 # enable                 (Enable the virtual server)
>> Virtual server 1 # service http           (Select menu for service on virtual server)
>> Virtual server 1 http service # group 1   (Use real server group 1 for http service)
>> Virtual server 2 http service # /cfg/slb/virt 2  (Select menu for virtual server 2)
>> Virtual server 1 # vip 10.10.10.2         (Set the IP address for virtual server 2)
>> Virtual server 1 # enable                 (Enable the virtual server)
>> Virtual server 1 # service http           (Select menu for service on virtual server)
>> Virtual server 1 # group 1                (Use real server group 1 for http service)
```

3. **Configure IP interfaces for each network that will be attached to OSPF areas.**

```
>> Virtual server 1 # /cfg/ip/if 1           (Select menu for IP Interface 1)
>> IP Interface 1 # addr 10.10.7.2           (Set IP address on backbone network)
>> IP Interface 1 # enable                   (Enable IP interface 1)
>> IP Interface 1 # ../if 2                  (Select menu for IP Interface 2)
>> IP Interface 2 # addr 10.10.10.41         (Set IP address on stub area network)
>> IP Interface 2 # enable                   (Enable IP interface 2)
```

4. **Enable OSPF on Web switch #2.**

```
>> IP Interface 2 # ../ospf/on               (Enable OSPF on Web switch #2)
```

5. **Define the backbone.**

```
>> Open Shortest Path First # aindex 0       (Select menu for area index 0)
>> OSPF Area (index) 0 # areaid 0.0.0.0      (Set the ID for backbone area 0)
>> OSPF Area (index) 0 # type transit        (Define backbone as transit type)
>> OSPF Area (index) 0 # enable              (Enable the area)
```

6. **Define the stub area.**

```
>> OSPF Area (index) 0 # ../aindex 1         (Select menu for area index 1)
>> OSPF Area (index) 1 # areaid 0.0.0.1      (Set the ID for stub area 1)
>> OSPF Area (index) 1 # type stub           (Define area as stub type)
>> OSPF Area (index) 1 # enable              (Enable the area)
```

7. Attach the network interface to the backbone.

```
>> OSPF Area (index) 1 # ../if 1          (Select OSPF menu for IP interface 1)
>> OSPF Interface 1 # aindex 0            (Attach network to backbone index)
>> OSPF Interface 1 # enable              (Enable the backbone interface)
```

8. Attach the network interface to the stub area.

```
>> OSPF Interface 1 # ../if 2             (Select OSPF menu for IP interface 2)
>> OSPF Interface 2 # aindex 1            (Attach network to stub area index)
>> OSPF Interface 2 # enable              (Enable the stub area interface)
```

9. Configure host routes.

Host routes are configured just like those on Web switch #1, except their costs are *reversed*. Since virtual server 10.10.10.2 is preferred for Web switch #2, its host route has been given a low cost. Because virtual server 10.10.10.1 is used as a backup in case Web switch #1 fails, its host route has been given a high cost.

```
>> OSPF Interface 2 # ../host 1           (Select menu for host route 1)
>> OSPF Host Entry 1 # addr 10.10.10.1    (Set IP address same as virtual server 1)
>> OSPF Host Entry 1 # aindex 0           (Inject host route into backbone area)
>> OSPF Host Entry 1 # cost 100           (Set high cost for use as backup path)
>> OSPF Host Entry 1 # enable             (Enable the host route)
>> OSPF Host Entry 1 # ../host 2          (Select menu for host route 2)
>> OSPF Host Entry 2 # addr 10.10.10.2    (Set IP address same as virtual server 2)
>> OSPF Host Entry 2 # aindex 0           (Inject host route into backbone area)
>> OSPF Host Entry 2 # cost 1             (Set low cost for primary path)
>> OSPF Host Entry 2 # enable             (Enable the host route)
```

10. Apply and save the configuration changes.

```
>> OSPF Host Entry 2 # apply              (Global command to apply all changes)
>> OSPF Host Entry 2 # save               (Global command to save all changes)
```

Verifying OSPF Support

Use the following commands to verify the OSPF configuration on your switch:

- **`/info/ospf/general`**
- **`/info/ospf/route`**
- **`/info/ospf/nbr`**

CHAPTER 3

Web OS 7.1 Command Reference

This section describes commands and options new to the Web OS 7.1 Command Line Interface (CLI). The CLI can be accessed via a local terminal or remote Telnet session as described in your complete Web OS 7.0 documentation.

The following menu items have been updated:

- Information menu (`/info`) on page 42
  - OSPF Information menu (`/info/ospf`) on page 43
- Management Processor Statistics menu (`/stats/mp`) on page 45
  - OSPF Statistics menu (`/stats/mp/ospf`) on page 45
- System Configuration menu (`/cfg/sys`) on page 46
  - Syslog Configuration menu (`/cfg/sys/syslog`) on page 47
  - Extended Buffer option (`/cfg/sys/extbuf`) on page 46
- IP Configuration menu (`/cfg/ip`) on page 48
  - Router ID option (`/cfg/ip/rtrid`) on page 48
  - OSPF Configuration menu (`/cfg/ip/ospf`) on page 49

    OSPF Area Configuration menu (`/cfg/ip/ospf/aindex`) on page 51
    OSPF Summary Range Configuration menu (`/cfg/ip/ospf/range`) on page 52
    OSPF Interface Configuration menu (`/cfg/ip/ospf/if`) on page 54
    OSPF Virtual Link Configuration menu (`/cfg/ip/ospf/virt`) on page 56
    OSPF Host Entry Configuration menu (`/cfg/ip/ospf/host`) on page 58

In the following sections, new and updated menus items appear in **bold** type. The syntax and usage of updated options follow each menu.

/info
Information Menu

```
[Information Menu]
      sys      - Show system information
      env      - Show temperature and fan status
      log      - Show last 10 syslog messages
      link     - Show link status
      port     - Show port status
      slot     - Show slot status
      vlan     - Show VLAN information
      stp      - Show STP information
      trunk    - Show Trunk Group information
      arp      - ARP information menu
      fdb      - Forwarding Database information menu
      ip       - Show IP information
      route    - IP routing information menu
      ospf     - OSPF routing information menu
      slb      - Layer 4 information menu
      vrrp     - Show VRRP information
      swkey    - Show enabled software features
      dump     - Dump all information
```

Table 1 OSPF IP Menu Options (/cfg/ip/ospf)

Command Syntax and Usage
ospf Displays the OSPF Information menu, used for gathering information about OSPF configuration and performance. For menu options, see page 43.

/info/ospf

OSPF Information Menu

```
[OSPF Information Menu]
      general - Show general information
      aindex  - Show area(s) information
      if      - Show interface(s) information
      virtual - Show details of virtual links
      nbr     - Show neighbor(s) information
      dbase   - Database Menu
      reqlist - Show request list of this neighbor
      retlist - Show retransmission list of this neighbor
      sumaddr - Show summary address list
      nsumadd - Show NSAA summary address list
      routes  - Show OSPF routes
```

The information provided in the area, interface, and neighbor menus are briefly described in the following sections, with examples of each information screen.

/info/ospf/aindex *<index>*

OSPF Area Information

```
Area (index) 1 with ID 0.0.0.1
   Number of interfaces in this area is 2 of which 2 are UP
   Area authentication: none
   LS types accepted by this area 63
   SPF algorithm executed 2 times
   Number of LSDB entries of this area 10
   Count of ASBR local to this area 1
   Count of ABR local to this area 2
   Total neighbors are 2 of which
               2 are >=INIT state,
               2 are >=EXCH state,
               2 are =FULL state
```

/info/ospf/if *<interface number>*

OSPF Interface Information

```
Ip Address 10.10.10.1, Area 0.0.0.1, Admin Status UP
   Router ID 10.10.12.1, State Backup DR, Priority 5
   Designated Router (ID) 10.10.13.2, Ip Address 10.10.10.2
   Backup Designated Router (ID) 10.10.12.1, Ip Address 10.10.10.1
   Timer intervals, Hello 10, Dead 40, Wait 0, Retransmit 5,
   Poll interval 0, Transit delay 1
   Neighbor count is 1   If Event 8, Authentication type None
```

/info/ospf/nbr

OSPF Neighbor Information

```
Intf  NeighborID       Prio  State         Address
----  ----------       ----  -----         -------
   1  10.10.13.2          1  Full           10.10.10.2
   2  10.10.9.2           1  Full           10.10.9.2
```

/stats/mp
Management Processor Statistics Menu

```
[MP Statistics Menu]
      if      - IP interface ("if") statistics
      ip      - IP statistics
      icmp    - ICMP statistics
      tcp     - TCP statistics
      udp     - UDP statistics
      snmp    - SNMP statistics
      arp     - ARP statistics
      route   - Route statistics
      dns     - DNS server statistics
      ospf    - OSPF Statistics Menu
      vrrp    - VRRP statistics
      maint   - MP Maintenance Statistics Menu
      clear   - Clear statistics
```

Table 2 OSPF IP Menu Options (/cfg/ip/ospf)

Command Syntax and Usage
ospf Displays the OSPF Statistics menu.

/stats/mp/ospf
OSPF Statistics Menu

```
[OSPF Statistics Menu]
      general - Show global stats
      aindex  - Show area(s) stats
      if      - Show interface(s) stats
```

/cfg/sys
System Configuration Menu

```
[System Menu]
      syslog  - Syslog Menu
      date    - Set system date
      time    - Set system time
      idle    - Set timeout for idle CLI sessions
      bootp   - Enable/disable use of BOOTP
      telnet  - Enable/disable Telnet access
      snmp    - Enable/disable SNMP management access
      web     - Enable/disable Web management access
      wport   - Set Web server port number
      mnet    - Set management network
      mmask   - Set management netmask
      banner  - Set login banner
      usrpw   - Set user password
      admpw   - Set administrator password
      l4apw   - Set L4 administrator password
      l4opw   - Set L4 operator password
      extbuf  - Enable/disable Extended Buffer Mode
      current - Display current system-wide configuration
```

Table 3 System Configuration Menu Options (/cfg/sys)

Command Syntax and Usage
syslog Displays the Syslog Configuration menu. This menu controls the level and destination of system log messages. For menu options, see page 47.
extbuf enable|disable In some environments, transfers may fail or experience unacceptable delays as a result of buffer congestion on the Web switch. If such congestion occurs, enable the extended buffer mode. This increases the Web switch frame buffer size from 16K to 64K, allowing bursts with larger block sizes. However, enabling the extended buffer mode reduces the layer-4/layer-7 session capacity from 256K sessions to 128K sessions, limiting the number of simultaneous sessions that can be processed for virtual servers, filtering, and content intelligent switching. By default, the extended buffer mode is disabled. When changing the mode setting, you must reboot the switch before the changes will take effect.

/cfg/sys/syslog

Syslog Configuration Menu

```
[Syslog Menu]
      host    - Set IP address of first syslog host
      host2   - Set IP address of second syslog host
      facil   - Set facility of first syslog host
      facil2  - Set facility of second syslog host
      console - Enable/disable console output of syslog messages
      current - Display current syslog settings
```

Table 4 System Configuration Menu Options (/cfg/sys)

Command Syntax and Usage
host *<IP address>* Set the IP address, in dotted-decimal notation, of the first host where system log (syslog) messages are to be sent. This command was formerly `/cfg/ip/log`.
host2 *<IP address>* Set the IP address, in dotted-decimal notation, of the second host where syslog messages are to be sent. This command was formerly `/cfg/ip/log2`.
facil *<level (0-7)>* Set the facility level of syslog messages sent to the first syslog host. This command was formerly `/cfg/ip/logfac`.
facil2 *<level (0-7)>* Set the facility level of syslog messages sent to the second syslog host. This command was formerly `/cfg/ip/logfac2`.
console enable|disable Enable/disable sending syslog messages to the console port.
current Display the current syslog configuration parameters.

/cfg/ip
IP Menu Configuration

```
[IP Menu]
      if      - Interface menu
      gw      - Default gateway menu
      route   - Static route menu
      frwd    - Forwarding menu
      rip     - Routing Information Protocol (RIP) menu
      ospf    - Open Shortest Path First (OSPF) menu
      port    - Port menu
      dns     - Domain Name System menu
      bootp   - Bootstrap Protocol Relay Menu
      rtrid   - Set router ID
      metric  - Set default gateway metric
      rearp   - Set re-ARP period in minutes
      current - Display current IP configuration
```

Table 5 OSPF IP Menu Options (/cfg/ip/ospf)

Command Syntax and Usage
ospf Displays the OSPF Configuration menu. For menu options, see page 49.
rtrid *<router ID>* Sets a static router ID for this switch. The router ID is expressed in dotted decimal IP address format (10.10.10.1 for example). OSPF, when enabled, uses the router ID to identify the routing device. If no router ID is specified (or if the router ID is set to 0.0.0.0 and the switch is rebooted), the switch dynamically selects one of the active IP interfaces on the switch as the router ID. For more information, see "The Router ID" on page 23. When using OSPF virtual links, you are required to set the router ID (see "Virtual Links" on page 23).

NOTE – The following syslog related commands have been moved from the `/cfg/ip` menu to the new `/cfg/sys/syslog` menu: `log`, `log2`, `fac`, and `fac2`. These commands have been renamed `host`, `host2`, `facil`, and `facil2` respectively. For more information, see "Syslog Configuration Menu" on page 47.

/cfg/ip/ospf

OSPF Menu Configuration

```
[Open Shortest Path First Menu]
      aindex  - OSPF Area (index) menu
      range   - OSPF Summary Range menu
      if      - OSPF Interface menu
      virt    - OSPF Virtual Links menu
      host    - OSPF Host Entry menu
      default - Originate default route information
      on      - Globally turn OSPF ON
      off     - Globally turn OSPF OFF
      current - Display current OSPF configuration
```

Table 6 OSPF IP Menu Options (/cfg/ip/ospf)

Command Syntax and Usage
aindex *<area index (0-7)>* Displays the OSPF Area menu. This menu is used for defining OSPF area numbers and parameters. For menu options, see page 51. **Note—**The area index specified in this menu option does not represent the actual OSPF area number. It is an arbitrary index used only on the switch. For more information, see "Defining Areas" on page 19.
range *<range number (1-64)>* Displays the OSPF Summary Range menu. This menu is used for defining OSPF summary routes, useful for condensing OSPF routing information. For menu options, see page 52. For more information, see "Summarizing Routes" on page 21.
if *<IP interface number (1-1024)>* Displays the OSPF Interface menu. This menu is used for attaching IP interface networks to OSPF areas. For menu options, see page 54. For more information, see "Attaching an Area to a Network" on page 20.
virt *<virtual link number (1-8)>* Displays the OSPF Virtual Links menu. This menu is used for connecting partitioned areas together. For menu options, see page 56. For more information, see "Virtual Links" on page 23.
host *<host number (1-1024)>* Displays the Host Entry menu. This menu is used for configuring host routes for server load balancing or ABR load sharing and failover. For menu options, see page 58. For more information, see "Host Routes for Load Balancing" on page 24.

Table 6 OSPF IP Menu Options (/cfg/ip/ospf)

Command Syntax and Usage
default *<metric value (1-16777215)> <metric type (1 or 2)>* **default none** The metric value sets the priority for choosing this switch for default routes where multiple ABRs or ASBRs exist in an area. For more information, see "Default Routes" on page 22. Metric type determines the method for influencing routing decisions for external routes. Type 1 external routes use the OSPF internal link state metrics. Type 2 external routes use the cost of the interface. The default is type 1.
on Globally turns on OSPF.
off Globally turns off OSPF.
current Displays the current OSPF configuration.

/cfg/ip/ospf/aindex <area index>

OSPF Area Configuration

```
[OSPF Area (index) 1 Menu]
      areaid  - Set area ID
      type    - Set area type
      metric  - Set stub area metric
      auth    - Set authentication type
      spf     - Set time interval between two SPF calculations
      enable  - Enable area
      disable - Disable area
      delete  - Delete area
      current - Display current OSPF area configuration
```

Table 7 OSPF Area Menu Options (/cfg/ip/ospf/area)

Command Syntax and Usage
areaid *<area ID (such as 0.0.0.0)>* Sets the OSPF area number in dotted decimal notation. The area number can be set using the last octet format (0.0.0.1 for area 1) or using multi-octet format (1.1.1.1), though the same format should be used throughout an area. For more information, see "Using the Area ID to Assign the OSPF Area Number" on page 19.
type transit|stub|nssa Sets the area type: ■ `transit` for the backbone or any area that contains a virtual link. ■ `stub` for any area that contains no external routes. ■ `nssa` for any area that can process external routes but does not advertise external routes originating from outside its area. For more information, see "Types of OSPF Areas" on page 14.
metric *<value (1-65535)>* Sets the stub area metric. Other routing devices add this value to the cost of routing to this stub area when building their SPF tree.
auth none|password Sets the authentication type for this area: ■ `none` turns off OSPF authentication. ■ `password` turns on type 1 (plain text) password authentication. The passwords are set using the `key` options for OSPF interface (see page 55) and OSPF virtual links (see page 57).
spf *<time interval (0-255)>* Sets the time interval, in seconds, between each calculation of the shortest path tree. The default is 10.

Table 7 OSPF Area Menu Options (/cfg/ip/ospf/area)

Command Syntax and Usage
enable Enables this area.
disable Disables this area.
delete Deletes this area.
current Displays the current OSPF area configuration.

/cfg/ip/ospf/range <range number>

OSPF Summary Range Configuration

```
[OSPF Summary Range 1 Menu]
      addr    - Set IP address
      mask    - Set IP mask
      aindex  - Set area index
      hide    - Enable/disable hide range
      enable  - Enable range
      disable - Disable range
      delete  - Delete range
      current - Display current OSPF summary range configuration
```

This menu is used for defining OSPF summary routes. Without summarization, each routing device in an OSPF network would retain a route to every subnet in the network. With summarization, routing devices can reduce some sets of routes to a single advertisement, reducing both the load on the routing device and the perceived complexity of the network. The importance of route summarization increases with network size.

Table 8 OSPF Summary Range Menu Options (/cfg/ip/ospf/range)

Command Syntax and Usage
addr *<IP address (for example, 192.4.17.1)>* Sets the base IP address for the summary range, using dotted decimal notation.
mask *<IP mask (for example, 255.255.255.0)>* Sets the IP mask for the summary range, using dotted decimal notation.

Table 8 OSPF Summary Range Menu Options (/cfg/ip/ospf/range)

Command Syntax and Usage
aindex *<area index (0-7)>* Sets the area index number into which the summary range is to be injected.
hide enable|disable When enabled, this forces the address range to be removed from any other summary ranges being injected into the defined area by this switch. This is useful for removing sections from large summary ranges which are not fully contiguous or contain gaps. This option is disabled by default.
enable Enables this range.
disable Disables this range.
delete Deletes this range.
current Displays the current OSPF summary range configuration.

/cfg/ip/ospf/if <interface number>

OSPF Interface Configuration

```
[OSPF Interface 1 Menu]
      aindex  - Set area index
      prio    - Set interface router priority
      cost    - Set interface cost
      hello   - Set hello interval in seconds
      dead    - Set dead interval in seconds
      trans   - Set transit delay in seconds
      retra   - Set retransmit interval in seconds
      key     - Set authentication key
      enable  - Enable interface
      disable - Disable interface
      delete  - Delete interface
      current - Display current OSPF interface configuration
```

NOTE – The hello interval (`hello`), dead interval (`dead`), transit interval (`trans`) and retransmit interval (`retra`) must be the same on *all* OSPF routing devices within an area. Using incompatible values could keep adjacencies from forming and could stop (or loop) routing updates.

Table 9 OSPF Interface Menu Options (/cfg/ip/ospf/if)

Command Syntax and Usage
aindex *<area index (0-7)>* Sets the OSPF area index to attach to the network for the current IP interface.
prio *<value (0-127)>* Sets the IP interface (IF) priority that will used when electing a Designated Router (DR) and Backup Designated Router (BDR) for the area. The default is 1 (lowest priority). A value of 0 specifies that the interface to be elected as `DROTHER` and cannot be used as a DR or BDR. For more information, see "Neighbors and Adjacencies" on page 16.
cost *<value (1-65535)>* Sets the cost of output routes on this interface. Cost is used in calculating the shortest path tree throughout the AS. Cost is based on bandwidth. Low cost indicates high bandwidth. The default is 1.
hello *<value (1-65535)>* Sets the hello interval, in seconds. The switch sends hello messages to inform neighbors that the link is up. The default is 10 seconds. This value must be the same on all routing devices within the area.

Table 9 OSPF Interface Menu Options (/cfg/ip/ospf/if)

Command Syntax and Usage
dead *<value (1-65535)>* Sets the router dead interval, in seconds. If the switch does not receive a hello on the IP interface within the deal interval, the switch will declare the interface to be down. Typically, the dead value is four times the value for Hello. The default is 40 seconds. This value must be the same on all routing devices within the area.
trans *<value (1-65535)>* Sets the transit delay, in seconds. This is the estimated time required to transmit an LSA to adjacencies on this interface, taking into account transmission and propagation delays. The default is 1 second. This value must be the same on all routing devices within the area.
retra *<value (1-65535)>* Sets the time interval, in seconds, between each transmission of LSAs to adjacencies on this interface. The default is 5 seconds. This value must be the same on all routing devices within the area.
key *<type 1 password (eight characters maximum)>* This option is used in conjunction with the `auth` option (see page 51). When `auth` is set to `none`, the key option is ignored. When `auth` is set to `password`, the key option sets the password to be used for OSPF authentication. Specify a type 1 (plain text) password of up to eight characters. To clear the key, specify **none** as the value.
enable Enables this interface.
disable Disables this interface.
delete Deletes this interface.
current Displays the current OSPF interface configuration.

/cfg/ip/ospf/virt <link number>

OSPF Virtual Link Configuration

```
[OSPF Virtual Links Menu]
      aindex  - Set area index
      hello   - Set hello interval in seconds
      dead    - Set dead interval in seconds
      trans   - Set transit delay in seconds
      retra   - Set retransmit interval in seconds
      nbr     - Set router ID of virtual neighbor
      key     - Set authentication key
      enable  - Enable interface
      disable - Disable interface
      delete  - Delete interface
      current - Display current OSPF interface configuration
```

Virtual links are typically created to connect one area to the backbone through another non-backbone area. The virtual link must be configured at each endpoint of the virtual link, though they may traverse multiple routing devices.

The minimum requirements for configuring a virtual link are the `aindex` and `nbr` options in this menu and the `rtrid` option in the IP Configuration menu (see page 48). For more information, see "Virtual Links" on page 23.

NOTE – The hello interval (`hello`), dead interval (`dead`), transit interval (`trans`) and retransmit interval (`retra`) must be the same on *all* OSPF routing devices within an area. Using incompatible values could keep adjacencies from forming, and may stop or loop routing updates.

Table 10 OSPF Virtual Link Options (/cfg/ip/ospf/virt)

Command Syntax and Usage
aindex *<area number (0-7)>* Sets the OSPF area index through which the virtual link passes.
hello *<value (1-65535)>* Sets the hello interval, in seconds. The switch sends hello messages to inform other network devices that the virtual link is up. The default is 10 seconds. This value must be the same on all routing devices within the area.

Table 10 OSPF Virtual Link Options (/cfg/ip/ospf/virt)

Command Syntax and Usage
dead *<value (1-65535)>* Sets the dead interval, in seconds. If the switch does not receive a hello on the IP interface within the deal interval, the switch will declare the virtual link to be down. Typically, the dead value is four times the hello value. The default is 40 seconds. This value must be the same on all routing devices within the area.
trans *<value (1-65535)>* Sets the transit delay, in seconds. This is the estimated time required to transmit an LSA to adjacencies, taking into account transmission and propagation delays. The default is 1 second. This value must be the same on all routing devices within the area.
retra *<value (1-65535)>* Sets the time interval, in seconds, between each transmission of LSAs to adjacencies. The default is 5 seconds. This value must be the same on all routing devices within the area.
nbr *<router ID (e.g., 192.4.17.101)>* Set router ID of the recipient neighbor (endpoint of the virtual link). The neighbor router ID is specified in dotted decimal format.
key *<type 1 password (eight characters maximum)>* This option is used in conjunction with the `auth` option (see page 51). When `auth` is set to `none`, the key option is ignored. When `auth` is set to `password`, the key option sets the password to be used for OSPF authentication. Specify a type 1 (plain text) password of up to eight characters. To clear the key, specify **none** as the value.
enable Enables this virtual link.
disable Disables this virtual link.
delete Deletes this virtual link.
current Displays the current OSPF virtual link configuration.

/cfg/ip/ospf/host <range number>

OSPF Host Entry Configuration

```
[OSPF Host Entry 1 Menu]
      addr    - Set host entry IP address
      aindex  - Set area index
      cost    - Set cost of this host entry
      enable  - Enable host entry
      disable - Disable host entry
      delete  - Delete host entry
      current - Display current OSPF host entry configuration
```

This menu is used for configuring Host routes. Host routes are used for advertising network device IP addresses to external networks. This makes Server Load Balancing (SLB) possible within OSPF. It also makes ABR load sharing and failover possible. For more information, see "Host Routes for Load Balancing" on page 24.

Table 11 OSPF Host Entry Menu Options (/cfg/ip/ospf/host)

Command Syntax and Usage
addr *<IP address (for example, 192.4.17.1)>* Specify the host IP address and mask that you want to advertise to the network.
aindex *<area index (0-7)>* Sets the area index into which the host entry is to be announced.
cost *<value (1-65535)>* Sets the cost of output routes on this host. Cost is used in calculating the shortest path tree throughout the AS. Cost is usually based on bandwidth (low cost indicates high bandwidth) but can be used to set route preferences. The default is 1.
enable Enables this host entry.
disable Disables this host entry.
delete Deletes this host entry.
current Displays the current OSPF host entry configuration.

Notes