

Table of Contents

<u>Guide to OSPF Application on the CSS 11000</u>	1
<u>Introduction</u>	1
<u>Before You Begin</u>	1
<u>Conventions</u>	2
<u>Prerequisites</u>	2
<u>Components Used</u>	2
<u>Description</u>	2
<u>OSPF Configuration Task List</u>	2
<u>Configuration</u>	2
<u>Global OSPF Commands</u>	3
<u>OSPF Interface Commands</u>	4
<u>OSPF Show Commands</u>	4
<u>Related Information</u>	7

Guide to OSPF Application on the CSS 11000

Introduction

Before You Begin

Conventions

Prerequisites

Components Used

Description

OSPF Configuration Task List

Configuration

Global OSPF Commands

OSPF Interface Commands

OSPF Show Commands

Related Information

Introduction

Open Shortest Path First (OSPF) is a link state routing protocol that maintains a local view of each area at each router, and to which a router may have an attached interface. When an OSPF router comes up, it exchanges hello messages to discover its neighbors and (in the case of a Local Area Network (LAN)) elects Designated and Backup Designated Routers (DR and BDR). At this stage, it records its state in the neighbor structures. Then, it proceeds to build its local view of the area.

First, the router exchanges a database summary message with its immediate neighbors. These messages are used to determine which Link State Advertisements (LSAs) need to be requested from the neighbors. The replies to the Link State Requests (LSRs) are the Link State Updates (LSUs) that are sent until the neighbor acknowledges in a link state acknowledgment. The process of achieving synchronization among all routers in an area is known as routing convergence. In the case of a LAN, the database synchronization occurs between the routers and the DR and BDR separately. There is no router-to-router exchange other than with the DR or BDR, hence the number of messages is considerably reduced. OSPF supports the notion of hierarchical routing. For example, an Autonomous System (AS) is organized into areas containing no more than 50 routers, and a backbone area (area 0). Each area must contain at least one router with an interface in the backbone area. In addition, the backbone area must be connected. In other words, the routers in the backbone area must be connected either directly by links in the backbone area or by a "virtual link" that crosses a transit area.

OSPF is intended for use where customers are currently running OSPF as their routing protocol and need the Content Services Switch (CSS) 11000 content services switch to participate in the learning and advertising of OSPF routes.

The following are two examples of when customers would run OSPF on the CSS:

1. When the CSS is used in a transparent or proxy cache environment where it is placed in the middle of the network and needs to learn routes back to clients.
2. In a firewall load balancing implementation where the firewall routes need to be redistributed into the OSPF domain downstream from the CSS.

Before You Begin

Conventions

For more information on document conventions, see the Cisco Technical Tips Conventions.

Prerequisites

There are no specific prerequisites for this document.

Components Used

This document is not restricted to specific software and hardware versions.

The information presented in this document was created from devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If you are working in a live network, ensure that you understand the potential impact of any command before using it.

Description

The CSS 11000 implementation of OSPF supports the following:

1. The ability to route in a single area between other OSPF routers (inter-area routes support).
2. The ability to route in multiple areas between OSPF routers (inter-area route support).
3. Hierarchical routing across multiple areas.
4. Route summarization between areas.
5. The AS boundary router support.
6. The stub-area support.
7. Routing Information Protocol (RIP) route leakage.
8. Redistribution of local, RIP, static, and firewall route into OSPF domain.
9. Simple authentication.
10. Management Information Base (MIB) per Request for Comments (RFC) 1850.

OSPF Configuration Task List

Perform the steps below to configure OSPF.

1. Configure an OSPF router ID. It is recommended that the IP address of the first OSPF interface be used.
2. Enable OSPF.
3. Configure an OSPF area. OSPF backbone area 0.0.0.0 is created by default.
4. Configure OSPF on an IP interface. The interface is added into the backbone area by default.
5. Enable OSPF on that interface.
6. Configure the advertisement of Versatile Interface Processors (VIPs) if needed (issue the **ospf advertise** command). This will advertise that network/host out all OSPF interfaces.
7. Configure the route redistribution into the OSPF domain, if needed.
8. Configure the OSPF area summarization, if needed.

Configuration

Global OSPF Commands

- **advertise** – Advertises a route as OSPF AS external through all OSPF interfaces. The default type is type2. Primarily used to advertise a VIP or a range of VIPs into an OSPF domain. The command syntax is shown below.

```
beta-rules(config)# ospf advertise 200.200.200.200 /32 optional sub commands
```

Sub commands of the advertise command include the following:

- ◆ **metric** – the metric to advertise.
- ◆ **tag** – 32-bit tag to advertise.
- ◆ **type1** – Advertise as ASE type 1 (comparable cost to OSPF metric).
- **metric** – Ranges from 1 to 15 and indicates the relative cost of this route. The larger the cost, the less preferable the route. The default is 1.
- **tag** – A 32-bit field attached to each external route. This is not used by the OSPF protocol itself. It may be used to communicate information between AS boundary routers.
- **type1** – Expressed in the same units as OSPF interface cost (that is, in terms of the link state metric). Type 2 external metrics are an order of magnitude larger; any Type 2 metric is considered greater than the cost of any path internal to the AS. This configuration parameter can be used to have an OSPF domain prefer type1 VIPs over type2.

Note: The CSS must be configured as an Autonomous System Boundary (ASB) router before issuing the **type1** command.

- **area** – Configures an OSPF area. By default, area 0.0.0.0 is already configured. You can also specify an area as being a stub area, as shown below.

```
beta-rules(config)# ospf area 2.2.2.2 stub ?
```

- ◆ **default-metric** – Metric for the default route advertised into the stub area.
- ◆ **send-summaries** – Propagates summary LSAs into this stub area.
- ◆ **as-boundary** – Configures the CSS as an ASB router.

An ASB is a router that exchanges routing information with routers belonging to other ASs, such as RIP domains. Issue this command to advertise VIPs, local, firewall, and RIP learned routes into an OSPF domain.

- **default** – Advertises a default route as ASE through OSPF. Options include **metric**, **tag**, and **type1** (type2 is default).
- **equal-cost** – Number of equal cost routes OSPF can use. The range is 1 through 15.
- **enable** – Enables OSPF globally.
- **range** – Configures route summarization between OSPF areas.

```
beta-rules(config)# ospf range 0.0.0.0 10.10.0.0 255.255.0.0
```

The OSPF area 0.0.0.0 contains the contiguous networks that you would like to advertise to other areas. You also have the ability to block the advertisement of a range. An example is provided below.

```
beta-rules(config)# ospf range 0.0.0.0 10.10.0.0 255.255.0.0 block
```

- **redistribute** – Advertises routes from other protocols through OSPF. Options include the following:
 - ◆ **firewall** – Advertises firewall routes through OSPF.
 - ◆ **local** – Advertises local routes through OSPF.

- ◆ *rip* – Advertises RIP routes through OSPF.
- ◆ *static* – Advertises static routes through OSPF. Sub options are *metric* , *tag* , and *type1* .
- **router-id** – Configures the OSPF router ID. It is recommended that you use the IP address of the first OSPF interface configured.

OSPF Interface Commands

The command syntax is shown below.

```
beta-rules(config-circuit-ip[VLAN2-20.20.1.2])# ospf ?
```

The command options are shown below.

- **area** – Configures the OSPF area to which this interface belongs. By default, an OSPF interface is already a member of the 0.0.0.0 area.
- **cost** – Sets the cost of sending a packet on this interface. The default cost is 10.
- **dead** – Sets the dead router interval (in seconds) for this interface. It is the number of seconds before the CSS's neighbors will declare it down, when they stop hearing the CSS's hello packets. The default is 40.
- **enable** – Enables OSPF on this interface.
- **hello** – Sets the hello interval (in seconds) for this interface. It is the length of time, in seconds, between the hello packets that the CSS sends on the interface. The default is ten.
- **password** – Sets the simple password (a maximum of eight characters) for this interface. Simple password authentication guards against routers inadvertently joining the routing domain; each router must first be configured with its attached networks' passwords before it can participate in routing. The password is in clear text.
- **poll** – Sets the poll interval (in seconds) for this interface. If a neighboring router has become inactive (hello packets have not been seen for RouterDeadInterval seconds), then it may still be necessary to send hello packets to the dead neighbor. These hello packets are sent at the reduced rate PollInterval, which should be much larger than HelloInterval. The default is ??.
- **priority** – Sets the router priority. When two routers attached to a network both attempt to become the DR, the one with the highest router priority takes precedence. If there is still a tie, the router with the highest router ID takes precedence. A router whose router priority is set to 0 is ineligible to become DR on the attached network. The default is 1.
- **retransmit** – Sets the retransmit interval (in seconds) for this interface. It is the number of seconds between LSA retransmissions, for adjacencies belonging to this interface. It is also used when retransmitting database description and link state request packets. This should be well over the expected round-trip delay between any two routers on the attached network. The setting of this value should be conservative, or needless retransmissions will result. The default is five.
- **retransmit** – Sets the retransmit interval (in seconds) for this interface. It is the number of seconds between LSA retransmissions, for adjacencies belonging to this interface. It is also used when retransmitting database description and link state request packets. This should be well over the expected round-trip delay between any two routers on the attached network. The setting of this value should be conservative, or needless retransmissions will result. The default is 5.

OSPF Show Commands

The list below contains sample output from various **show ospf** commands.

1. show ospf advertise

```
beta-rules# show ospf advertise
OSPF Advertise Routes Entries:
```

```

Advertise Routes Prefix :    200.200.200.200
Advertise Routes Prefix Length :    32
Advertise Routes Metric :    1
Advertise Routes Type :    aseType2
Advertise Routes Tag :    0

```

Note: In the above **show** command screen, a VIP with a 32-bit mask is advertised. Defaults are used for the other parameters.

2. show ospf areas

```

beta-rules# show ospf areas

```

Area ID	Type	SPF Runs	Routers	Routers	LSAs	Summaries
0.0.0.0	Transit	46	0	1	3	N/A
2.2.2.2	Stub	5	0	1	1	Yes

3. show ospf ase

```

beta-rules# show ospf ase

```

Link State ID	Router ID	Age	T	Tag	Metric	Forwarding Address
0.0.0.0	192.168.151.1	1	2	00000000	1	0.0.0.0
200.200.200.200	192.168.151.1	593	2	00000000	1	0.0.0.0

Note: Data traffic for the advertised destination will be forwarded to the forwarding address. If the forwarding address is set to 0.0.0.0, data traffic will be forwarded instead to the LSA's originator (that is, the responsible ASB router).

4. show ospf global

```

beta-rules# show ospf global
OSPF Global Summary:

Router ID:                192.168.151.1
Admin Status:              enabled
Area Border Router:        FALSE
AS Boundary Router:        TRUE
External LSAs :            2
LSA Sent :                  8
LSA Received :             5

```

5. show ospf interfaces

```

beta-rules# show ospf interfaces
OSPF Interface Summary:

IP Address:                192.168.151.1
Admin State:               enabled
Area:                      0.0.0.0   Type:                broadcast
State:                     BDR       Priority:              1
DR:                        192.168.151.2 BDR:                  192.168.151.1
Hello:                     10        Dead:                  40
Transmit Delay:             1        Retransmit:           5
Cost:                      10

```

6. show ospf lsd

```

beta-rules# show ospf lsd
OSPF LSDB Summary:

Area:                      0.0.0.0   Type:                Router
Link State ID:             192.168.151.1 ADV Router:           192.168.151.1
Age:                       699

```

```

Sequence:                0x80000003
Checksum:                0xdf5d

Area:                    0.0.0.0   Type:                Router
Link State ID:          192.168.151.2  ADV Router:          192.168.151.2
Age:                    706
Sequence:                0x80000004
Checksum:                0xd565

Area:                    0.0.0.0   Type:                Network
Link State ID:          192.168.151.2  ADV Router:          192.168.151.2
Age:                    706
Sequence:                0x80000001
Checksum:                0xbd93

Area:                    0.0.0.0   Type:                ASE
Link State ID:          0.0.0.0     ADV Router:          192.168.151.1
Age:                    114
Sequence:                0x80000001
Checksum:                0xb51a

Area:                    0.0.0.0   Type:                ASE
Link State ID:          200.200.200.200  ADV Router:          192.168.151.1
Age:                    706
Sequence:                0x80000001
Checksum:                0xa10b

```

7. show ospf neighbors

```

beta-rules# show ospf neighbors
  Address      Neighbor ID    Prio    State    Type    Rxmt_Q
  -----
192.168.151.2  192.168.151.2  1      Full    Dynamic  0

```

8. show ospf range

```

beta-rules# show ospf range
Area ID      LsdbType      Addr Range    Mask Range    Effect
-----
2.2.2.2      summaryLink    150.0.0.0     255.0.0.0     advertise

```

9. show ospf redistribute

```

beta-rules# show ospf redistribute
Redistribution via OSPF Summary:

Static Routes Redistribution :      disabled

RIP Routes Redistribution :        disabled

Local Routes Redistribution :      disabled

Firewall Routes Redistribution :    disabled

```

10. show ip routes ospf

```

beta-rules# show ip routes ospf
prefix/length  next hop      if  type  proto  age  metric
-----
20.20.20.0/24  150.150.150.2  1021 remote ospf  5  1

```

Related Information

- **OSPF Technical Support**
 - **OSPF Design Guide**
 - **Technical Support – Cisco Systems**
-

All contents are Copyright © 1992–2003 Cisco Systems, Inc. All rights reserved. Important Notices and Privacy Statement.