ERICSSON

Troubleshooting OSPF

SmartEdge OS

FAULT TRACING DIRECT

13/154 51-CRA 119 1170/1 Uen D

Disclaimer

The contents of this document are subject to revision without notice due to continued progress in methodology, design and manufacturing. Ericsson shall have no liability for any error or damage of any kind resulting from the use of this document.

Trademark List

SmartEdge is a registered trademark of Telefonaktiebolaget LM Ericsson.

Contents

1 Troubleshooting General OSPF Issues

1.1 Sample OSPF Topology

Use the following sample Open Shortest Path First (OSPF) topology and configuration as a guide to troubleshooting general OSPF issues. The configuration and sample output in this section match the sample topology. For specific OSPF issues, see Section 2 on page 43.

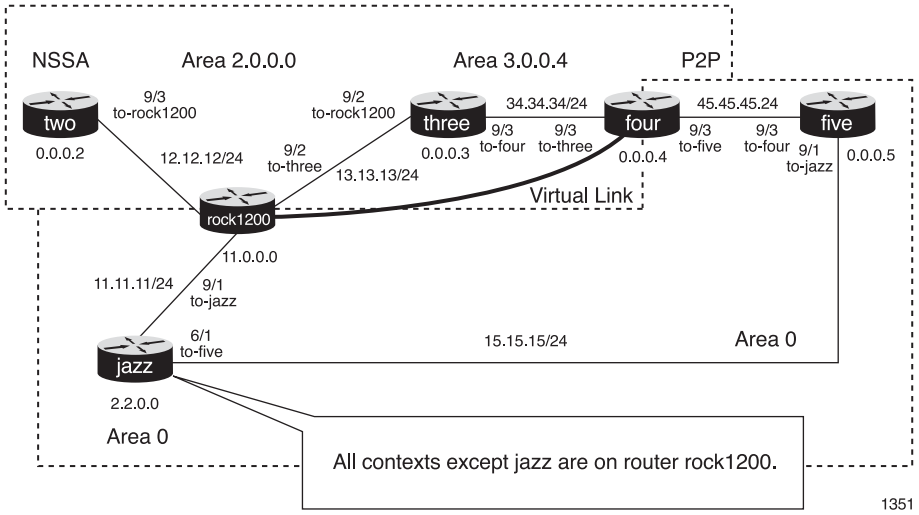

Figure 1 Sample OSPF Topology

Note: Troubleshooting OPSFv3 is beyond the scope of this document.

1.2 Sample OSPF Configuration

The following are sample OSPF configurations on the `jazz` and `rock1200` routers, which match the sample topology in Section 1.1 on page 1.

The more complex configuration is on router `rock1200`. The other routers simply have a context with an OSPF instance configured for area 0. This setup produces all possible LSAs.

Routers two, three, four, and five are contexts on router `rock1200`, not standalone SmartEdge® routers.

1.2.1 jazz Router Configuration

The SmartEdge OS supports multiple OSPF processes on the SmartEdge router. The OSPF configuration on the `jazz` router shows that the OSPF process running on that router is process 22. The SmartEdge OS also supports multiple OSPF processes on the same router, which can help in the redistribution of routes between OSPF processes and keep routes from different OSPF processes separate.

```
Building configuration...

Current configuration:
!
context jazz
!
 no ip domain-lookup
!
  router ospf 22
  fast-convergence
  router-id 22.0.0.0
  area 0.0.0.0
   interface to-rock1200
   interface to-five
!

!
!
!
!
!
End

[local]Redback#show config
Building configuration...

Current configuration:
!
!  Configuration last changed by user 'test' at Tue Aug 31 15:17:58 2010
!
!
!
!
!
!
service multiple-contexts
!
!
!
!
!
!
!
```

```
!
!
!
!
!
!
dpi access-list list1
!
!
!
!
!
context local
!
 no ip domain-lookup
!
  interface mgmt
  ip address 10.18.17.102/24
!
 interface pim1
 logging console
!
!
!
 administrator test encrypted 1 $1$........$kvQfdsjs0ACFMeDHQ7n/o.
   privilege start 15
!
!
 ip route 10.0.0.0/11 10.18.17.254
 ip route 155.53.0.0/16 10.18.17.254
!
!
!
context jazz
!
 no ip domain-lookup
!
 interface to-five
  ip address 15.15.15.1/24
 !
 interface to-rock1200
  ip address 11.11.11.2/24
 no logging console
!
 router ospf 22
  fast-convergence
  router-id 22.0.0.0
  area 0.0.0.0
   interface to-rock1200
   interface to-five
!
!
!
!
!
!
!
card ge-5-port 6
!
port ethernet 6/1
 no shutdown
 encapsulation dot1q
 dot1q pvc 1
  bind interface to-rock1200 jazz
 dot1q pvc 2
  bind interface to-five jazz
!
!
port ethernet 7/1
! XCRP management ports on slot 7 and 8 are configured through 7/1
 no shutdown
 bind interface mgmt local
!
card ch-ds3-3-port 10
!
card ge-2-ort 14
```

```
!
port ethernet 14/1
 no shutdown
 encapsulation dot1q
!
 timeout session idle 999
!
no service console-break
!
service crash-dump-dram
!
no service auto-system-recovery
!
end
```

1.2.2 rock1200 Router Configuration

```
context rock1200
!
 no ip domain-lookup
!
 interface to-jazz
  ip address 11.11.11.1/24
  !
 interface to-three
  ip address 13.13.13.1/24
 !
 interface to-two
  ip address 12.12.12.1/24

 no logging console
!
 router ospf 1
  fast-convergence
  router-id 11.0.0.0
  area 0.0.0.0
   virtual-link 3.0.0.4 0.0.0.4
   interface to-jazz
  area 2.0.0.0
   area-type nssa
   interface to-two
  area 3.0.0.4
   interface to-three
!
 !
!
!
!
```

```
context two
!
 no ip domain-lookup
!
 interface to-rock1200
  ip address 12.12.12.2/24
  no logging console
!
 router ospf 2
  fast-convergence
  router-id 0.0.0.2
  area 2.0.0.0
   area-type nssa
   interface to-rock1200
  redistribute static
!
 ip route 88.88.88.88/32 null0
!
!
!
!
!
```

```
context three
!
 no ip domain-lookup
!
 interface to-four
  ip address 34.34.34.1/24
!
 interface to-rock1200
  ip address 13.13.13.2/24
 no logging console
!
  ip route 99.99.99.99/32 null0
!
!
!
!
!
context four
!
 no ip domain-lookup
!
 interface to-five p2p
  ip address 45.45.45.1/24
 !
 interface to-three
  ip address 34.34.34.2/24
  no logging console
!
 router ospf 4
  fast-convergence
  router-id 0.0.0.4
  area 0.0.0.0
   interface to-five
   virtual-link 3.0.0.4 11.0.0.0
  area 3.0.0.4
   interface to-three
!
 !
!
!
context five
!
 no ip domain-lookup
!
 interface to-four p2p
  ip address 45.45.45.2/24
!
 interface to-jazz
  ip address 15.15.15.2/24
 no logging console
!
 router ospf 5
  fast-convergence
  router-id 0.0.0.5
  area 0.0.0.0
   interface to-four
   interface to-jazz
!

!
card ge-20-port 9
!
port ethernet 9/1
 no shutdown
 encapsulation dot1q
 dot1q pvc 1
  bind interface to-jazz rock1200
 dot1q pvc 2
  bind interface to-jazz five
!
port ethernet 9/2
 no shutdown
 encapsulation dot1q
 dot1q pvc 1
  bind interface to-two rock1200
 dot1q pvc 2
```

```
  bind interface to-three rock1200
 dot1q pvc 3
  bind interface to-four three
 dot1q pvc 4
  bind interface to-five four
!
port ethernet 9/3
 no shutdown
 encapsulation dot1q
 dot1q pvc 1
  bind interface to-rock1200 two
 dot1q pvc 2
  bind interface to-rock1200 three
 dot1q pvc 3
  bind interface to-three four
 dot1q pvc 4
  bind interface to-four five
!
End
```

1.3 Tasks for Troubleshooting General OSPF Issues

Use the following table as a guide to troubleshooting general OSPF issues. More information about each step is provided in subsequent sections.

Table 1 Tasks to Troubleshoot OSPF

Task	Command	Notes	Checked?
Step 1: Navigate to the Context	**show context all** **show ospf *instance*-id**	• Display all the contexts on the router and then navigate to the context you want to troubleshoot.	
	show process ospf	• Display high-level information for all OSPF instances, or optionally, for a specific instance. • Check for process restarts and uptime.	
Step 2: Verify Port Status	**show port** **show port counters**	• Make sure your ports are enabled and the circuit configuration matches the configuration (encapsulation and circuit number) at each end of the circuit .	
	ping	• Make sure the port counters are incrementing. • Verify that links between routers are operational.	
Step 3: Verify Interfaces	**show ip interface brief**	• Make sure your interfaces are up.	
	show ospf interface	• Verify OSPF interfaces.	
	show ospf interface detail	• Display detailed information about OSPF interfaces.	
Step 4: Verify Connectivity	**ping** **traceroute**	Verify that links between routers are operational.	
Step 5: Verify OSPF Configuration	**show configuration ospf** **show ospf global**	• Make sure both endpoint types match; for example, point-to multipoint, point-to-point (P2P). Otherwise you cannot establish an adjacency.	
		• Display how many OSPF instances are configured on the SmartEdge router in the local context.	
Step 6: Check OSPF Adjacency	**show ospf neighbor**	• Verify that every router establishes OSPF neighborship.	
	show ospf neighbor interface	• Display OSPF adjacency information for a specific interface.	
	traceroute		
Step 7: Verify LSAs in OSPF Database	**show ospf database**	This database includes information about the network topology for this area. All the routers in this area should have the same database. Make sure SPF has the required LSAs.	
Step 8: Verify OSPF Routes	**show ip route ospf**	• View OSPF route entries in the RIB.	
	show ospf route	• View OSPF route entries in the OSPF route table.	
	show ospf route *ip-address*	• View a specific route and hops in the OSPF route table.	

Table 1 Tasks to Troubleshoot OSPF

Task	Command	Notes	Checked?
Step 9: Verify IP Routes	**show ip route**	• Verify the active (best routes) in the RIB.	
	show ip route all	• View all routes stored in the RIB.	
Step 10: Check OSPF Statistics	**show ospf statistics [*instance-id*| Interface] [detail]**	Verify the OSPF traffic information.	
Step 11: Verify OSPF SPF Log	**show ospf spf**	Display a history of the SPF calculation results.	
Step 12: Check OSPF in Logs	**show log | grep ospf**	Filter the log for entries relating to OSPF.	
		You must enable the **log-neighbor-up-down** command to view OSPF logs.	
Step 13: Monitor OSPF Events	**monitor ospf interface**	• Display continuously updated information about OSPF interfaces.	
	monitor ospf neighbor	• Display continuously updated information about OSPF neighbors.	
	monitor ospf spf last	• Display continuously updated information about the most recent OSPF SPF calculation.	
	monitor ospf statistics	• Display continuously updated information about OSPF statistics.	
Step 14: Debug OSPF	**debug ospf packet errors**	Check for MTU, area ID, authentication, and interface issues.	
		Note: Risk of performance loss. Enabling the generation of debug messages can severely affect system performance. To reduce the risk, exercise caution when enabling the generation of debug messages on a production system.	

1.4 Step 1: Navigate to the Correct Context

Run the **`show context all`** command to display all the contexts on the router, and then navigate to the context you want to troubleshoot—in this case, `rock1200`.

```
[local]rock1200#show context all
Context Name               Context ID        VPN-RD               Description

--------------------------------------------------------------------------------
local                      0x40080001

rock1200                   0x40080082

two                        0x40080083

three                      0x40080084

four                       0x40080085

five                       0x40080086

[local]rock1200#context rock1200
[rock1200]rock1200#
```

Run the **`show ospf [`*`instance-id`*`]`** command to display high-level information for all OSPF instances or for a specific instance.

```
[rock1200]rock1200#show ospf

  --- OSPF Instance 1/Router ID 11.0.0.0 ---

Intra-Distance  : 110              Inter-Distance  : 110
Ext-Distance    : 110              Type of Service : TOS-Type0
Area Border Rtr : Yes              AS Boundary Rtr : Yes
Auto-Cost       : Yes              Flood Queued    : 0
SPF Delay       : 5                SPF Holdtime    : 10
Full SPF Count  : 12               Incr SPF Count  : 0
Full SPF Vers   : 12               Incr SPF Vers   : 0
SPF LastCompute : 01:42:27         Nbrs Adjacent   : 3
Nbrs Exchanging : 0                Global Exchg Max: 300
Redist Metric   : Unspecified      Redist Queued   : 0
Redist Count    : 0                Redist Quantum  : 2000
Stub Rtr Config : None             Stub Rtr Delay  : 0
Stub Router     : No               BGP Converged   : No
MPLS Traffic Eng: No               MPLS Shortcuts  : No
Demand DC Clear : 0                Demand Indicate : 0
Demand DoNotAge : 6                Helper Neighbors: 0
Graceful Restart: No               Restart Status  : No Restart
Graceful Helper : Yes              Strict Check    : No
Fast Convergence: Yes              Fast LSA Orig   : No

Area List (3 total):
0.0.0.0         2.0.0.0         3.0.0.4
[rock1200]rock1200#
```

Run the **`show process ospf`** command to check for process restarts and uptime.

```
[rock1200]rock1200#show process ospf

NAME           PID    SPAWN    MEMORY  TIME            %CPU  STATE     UP/DOWN
ospf          2038        1     7472K  00:00:15.61    0.00%  run       01:59:38
```

1.5 Step 2: Verify Port Status

Run the **show port** command to verify that the ports are up. To see "Admin" and "Line" states, run the **`show port detail`** command.

Before you check the status of a port, make sure that you understand the differences between the Admin state and the Line state:

- Admin state—Refers whether the port has been brought up (by using the **`no shutdown`** command) or is down (by using the **`shutdown`** command). If the Admin state is *shut down*, the port is down.

  Recommended Action: Run the **`no shutdown`** command on the port to bring up the port.

- Line state—Refers to the physical state of the port.

Recommended Action: When the Line state is *down*, use the checklist in Table 2.

Table 2 Line State Troubleshooting Checklist

#	Line State Troubleshooting Checklist	Checked?
1	Is the cable correctly connecting the two ports or nodes? In some cases, you might have ports looped together externally to connect different interfaces within separate contexts and share routing information between the two contexts through OSPF. In this case, check if both ports and their corresponding interfaces are up.	
2	Is there a fault in the cable?	
3	Are you using the right type of cable? For example, with Ethernet, are you using a cross-over cable instead of a straight cable?	
4	When the cable is connected to two nodes, is there a fault in one of the nodes?	
5	Is the card with a fiber port receiving light? Is the LOS LED in the port on?	
6	If you are using fiber optics, are you using the appropriate fiber type (for example, multimode or single mode) ?	
7	Is the other end port shut down?	
8	Is there a link speed or duplex setting mismatch?	
9	Is the SmartEdge router gigabit Ethernet port connected to an FE port? The SmartEdge router gigabit Ethernet traffic cards do not support FE speeds.	
10	Are the fibers correctly connected?	
11	Does the circuit configuration match?	
12	Is the line card configured correctly?	

If the Admin state is *Down*, the Line state is always down. For the port to be *Up*, the Admin state and Line state must both be *Up*. To see "Admin" and "Line" states, run the **show port detail** command. The **show port** command always returns real-time results. To see results in real time, use the **detail** keyword. You can use the **detail** or **live** parameters when verifying port counters or circuit counters. For detailed information about each field , see the *Command List*.

Use the following table to determine whether a port is *Up* or *Down*.

Table 3 Port States

Admin State (Configuration)	Line State (Physical)	Result
Up	Down	Down
Up	Up	Up
Down	Up	Down
Down	Down	Down

In the following example, the status of the Ethernet ports is *Up*.

```
[rock1200]rock1200#show port
Slot/Port:Ch:SubCh  Type                State
7/1                 ethernet            Up
9/1                 ethernet            Up
9/2                 ethernet            Up
9/3                 ethernet            Up
9/4                 ethernet            Up

[rock1200]rock1200#
```

In the following example, the status of the Ethernet port is *Down*. Although the Ethernet port is in a *no shutdown* state and the Admin state is *Up*, the cable has been unplugged from the Ethernet port `8/1` and, as a result, the Line state (the physical state) is *Down*.

```
[rock1200]rock1200#show port 9/1 detail
ethernet 9/1 state is Up
Description                 :
Line state                  : Down
Admin state                 : Up
Link Dampening              : disabled
Undampened line state       : Down
Dampening Count             : 0
Encapsulation               : dot1q
MTU size                    : 1500 Bytes
NAS Port Type               :
NAS-Port-Id                 :
MAC address                 : 00:30:88:11:d1:8d
Media type                  : 1000Base-LX
Auto-negotiation            : on                  state: fail
   Flc negotiated set       : tx  state: tx
   force                    : disabled            state: inactive
Flow control                : rx                  state: n/a
Speed                       : 1 Gbps
Duplex mode                 : full
Link Distance               : 15000 meters
Loopback                    : off
SFP Transceiver Status
 Wavelength                 : 1310.00 nm
 Diag Monitor               : Yes
 Tx Fault                   : No Fault
 Rx Fault                   : No Fault
 Tx Pwr measured[dbm]       : -11.90
 Rx Pwr measured[dbm]       : -6.72
 Temperature                : 49 C
 Vcc Measured               : 3.27 V
Active Alarms               : Link down
```

Each line card collects Layer 1, 2, and 3 statistics. Counters are updated every 60 seconds, unless you specify the **live** parameter. To check port counters, generate traffic on the port, and then run the **show port counters** command several times to determine if traffic is increasing on the port. For detailed information about each field displayed, see the *Command List*.

```
[rock1200]rock1200#show port counters
Port            Type
7/1             ethernet
packets sent       : 19293             bytes sent         : 1999412
packets recvd      : 42333             bytes recvd        : 3148785

9/1             ethernet
packets sent       : 448760            bytes sent         : 41552940
packets recvd      : 479000            bytes recvd        : 44527492
send packet rate   : 0.58              send bit rate      : 422.60
recv packet rate   : 0.60              recv bit rate      : 461.52
rate refresh interval : 60 seconds

9/2             ethernet
packets sent       : 889361            bytes sent         : 82364298
packets recvd      : 953289            bytes recvd        : 88298488
send packet rate   : 1.15              send bit rate      : 856.66
recv packet rate   : 1.13              recv bit rate      : 851.59
rate refresh interval : 60 seconds

9/3             ethernet
packets sent       : 889415            bytes sent         : 82332018
packets recvd      : 954644            bytes recvd        : 88457916
send packet rate   : 1.13              send bit rate      : 851.60
recv packet rate   : 1.15              recv bit rate      : 856.66
rate refresh interval : 60 seconds

9/4             ethernet
packets sent       : 0                 bytes sent         : 0
packets recvd      : 0                 bytes recvd        : 0
send packet rate   : 0.00              send bit rate      : 0.00
recv packet rate   : 0.00              recv bit rate      : 0.00
rate refresh interval : 60 seconds

[rock1200]rock1200#
```

1.6 Step 3: Verify Interfaces

1.6.1 Verify All Interfaces

Run the **`show ip interface brief`** command (in the local context) to check if the interfaces are enabled and *Up*. This command displays information about all interfaces, associated addresses, states, and bindings, including the interface bound to the Ethernet management port on the controller card.

An interface can be in any of the following states:

- Unbound—The reasons why the interface might be unbound are that the interface is not currently bound to any port or circuit. The binding is not valid. The VLANs need to match.

  Note: In some cases, an interface can have an Unbound state and still be valid; for example, multibind interfaces where no active PPPoE or CLIPS sessions are active.

- Bound—The interface is bound to at least one port or circuit; however, none of the bound circuits are up. Therefore, the interface is not up. The binding is valid. The state *Bound* is expected behavior for multibind interfaces where there are no active subscribers.

- Up—At least one of the bound circuits is in the up state; therefore, the interface is also up and traffic can be sent over the interface. The binding is valid.

If the interfaces are not *Up*, check the configuration and make sure you have enabled the interface. Both endpoints must have the same interface type; for example, point-to-multipoint or NBMA.

When OSPF passive mode is enabled, OSPF continues to advertise the interface IP subnet, but it does not send OSPF packets and drops all received OSPF packets. Enable OSPF passive mode by using the **`passive`** command for either of the following:

- An individual OSPF interface

- All interfaces within a defined OSPF area

When the interface is defined as passive, no adjacency will be formed and, as result, you will not see this interface in the **`show ospf neighbor`** command output. For detailed information about each field displayed, see the *Command List*.

```
[local]rock1200#show ip interface brief

Tue Dec 14 07:04:50 2010
Name              Address                  MTU   State    Bindings
loopback          10.10.10.2/32            1500  Up       (Loopback)
mgmt              10.18.17.103/24          1500  Up       ethernet 7/1
to-jazz           1.1.1.2/24               1500  Up       ethernet 9/1
[local]rock1200#
```

1.6.2 Verify OSPF Interfaces

Run the **show ospf interface** command to verify that the OSPF interfaces are *Up*.

Use the **detail** keyword to display detailed information about the interface.

if-name	Optional. Interface name. Displays information only for the specified interface.
ip-address	Optional. Name of a particular interface.
detail	Optional. Displays detailed OSPF interface information.

Run the **show ospf interface** output to display summary information about all configured OSPF interfaces in context `rock1200`:

```
[rock1200]rock1200#show ospf interface

  --- OSPF Interfaces for Instance 1/Router ID 11.0.0.0 ---

Addr            Len  NetworkType     Cost    Priority State    Area
0.0.0.4         0    virtual         2       1        P2P      0.0.0.0
11.11.11.1      24   broadcast       1       1        BDR      0.0.0.0
12.12.12.1      24   broadcast       1       1        BDR      2.0.0.0
13.13.13.1      24   broadcast       1       1        DR       3.0.0.4
[rock1200]rock1200#
```

Run the **show ospf interface** command with the argument ***if-name*** to display information about a specific OSPF interface, `to-jazz`:

```
[rock1200]rock1200#show ospf interface to-jazz

  --- OSPF Interface 11.11.11.1 Area 0.0.0.0 Instance 1 ---

Network Type   : broadcast          Mask           : 255.255.255.0
Cost           : 1                  Logical Intf   : to-jazz
MTU            : 1500               Physical Intf  : ethernet 9/1
State          : BDR                Priority       : 1
Hello Interval : 10                 Dead Interval  : 40
Transmit Delay : 1                  Retransmit Int : 5
DR Router ID   : 22.0.0.0           DR IP Address  : 11.11.11.2
BDR Router ID  : 11.0.0.0           BDR IP Address : 11.11.11.1
Ack Queued     : 1                  Flood Queued   : 0
Ack Delay      : 2                  Authentication : None
LSA Count      : 0                  LSA Checksum   : 0
Demand Circuit : No                 Flood Reduction: No
Neighbor Count : 1

Neighbor List (1 Adjacent):
22.0.0.0
[rock1200]rock1200#
```

Run the **`show ospf interface detail`** command to view detailed information about your OSPF interfaces:

```
[rock1200]rock1200#show ospf interface detail

  --- OSPF Interface 0.0.0.4 Area 0.0.0.0 Instance 1 ---

Network Type    : virtual           Mask            : 0.0.0.0
Endpoint Router: 0.0.0.4            Transit Area    : 3.0.0.4
Cost            : 2                 Logical Intf    : to-three
MTU             : 1500              Physical Intf   : ethernet 9/2
State           : P2P               Priority        : N/A
Hello Interval  : 10                Dead Interval   : 40
Transmit Delay  : 1                 Retransmit Int  : 5
Ack Queued      : 0                 Flood Queued    : 0
Ack Delay       : 2                 Authentication  : None
LSA Count       : 0                 LSA Checksum    : 0
Demand Circuit  : Yes               Flood Reduction: Yes
Suppress Allow  : Yes               Suppress Active: Yes
Neighbor Count  : 1

Neighbor List (1 Adjacent):
0.0.0.4

  --- OSPF Interface 11.11.11.1 Area 0.0.0.0 Instance 1 ---

Network Type    : broadcast         Mask            : 255.255.255.0
Cost            : 1                 Logical Intf    : to-jazz
MTU             : 1500              Physical Intf   : ethernet 9/1
State           : DR                Priority        : 1
Hello Interval  : 10                Dead Interval   : 40
Transmit Delay  : 1                 Retransmit Int  : 5
DR Router ID    : 11.0.0.0          DR IP Address   : 11.11.11.1
BDR Router ID   : 0.0.0.0           BDR IP Address  : 0.0.0.0
Ack Queued      : 0                 Flood Queued    : 0
Ack Delay       : 2                 Authentication  : None
LSA Count       : 0                 LSA Checksum    : 0
Demand Circuit  : No                Flood Reduction: No
Neighbor Count  : 0

  --- OSPF Interface 12.12.12.1 Area 2.0.0.0 Instance 1 ---

Network Type    : broadcast         Mask            : 255.255.255.0
Cost            : 1                 Logical Intf    : to-two
MTU             : 1500              Physical Intf   : ethernet 9/2
State           : DR                Priority        : 1
Hello Interval  : 10                Dead Interval   : 40
Transmit Delay  : 1                 Retransmit Int  : 5
DR Router ID    : 11.0.0.0          DR IP Address   : 12.12.12.1
BDR Router ID   : 0.0.0.2           BDR IP Address  : 12.12.12.2
Ack Queued      : 0                 Flood Queued    : 0
Ack Delay       : 2                 Authentication  : None
LSA Count       : 0                 LSA Checksum    : 0
Demand Circuit  : No                Flood Reduction: No
Neighbor Count  : 1

Neighbor List (1 Adjacent):
0.0.0.2

  --- OSPF Interface 13.13.13.1 Area 3.0.0.4 Instance 1 ---

Network Type    : broadcast         Mask            : 255.255.255.0
Cost            : 1                 Logical Intf    : to-three
MTU             : 1500              Physical Intf   : ethernet 9/2
State           : DR                Priority        : 1
Hello Interval  : 10                Dead Interval   : 40
Transmit Delay  : 1                 Retransmit Int  : 5
DR Router ID    : 11.0.0.0          DR IP Address   : 13.13.13.1
BDR Router ID   : 0.0.0.3           BDR IP Address  : 13.13.13.2
Ack Queued      : 0                 Flood Queued    : 0
Ack Delay       : 2                 Authentication  : None
LSA Count       : 0                 LSA Checksum    : 0
Demand Circuit  : No                Flood Reduction: No
Neighbor Count  : 1
```

```
Neighbor List (1 Adjacent):
0.0.0.3
:
```

1.7 Step 4: Verify Connectivity

Run the **ping** and **traceroute** commands to verify that links between routers are operational. Use the **ping** command to check directly connected interfaces and virtual links. Use the **traceroute** command to check for virtual links.

On router `rock1200`, the following example successfully pings and traces the routes of router `four` (context four) with the address `34.34.34.2`:

```
[rock1200]rock1200#ping 34.34.34.2
PING 34.34.34.2 (34.34.34.2): source 13.13.13.1, 36 data bytes,
timeout is 1 second
!!!!!

----34.34.34.2 PING Statistics----
5 packets transmitted, 5 packets received, 0.0% packet loss
round-trip min/avg/max/stddev = 1.784/2.000/2.258/0.172 ms
[rock1200]rock1200#


[rock1200]rock1200#traceroute 34.34.34.2
se_traceroute to 34.34.34.2 (34.34.34.2), 30 hops max, 40 byte packets
 1  13.13.13.2 (13.13.13.2)   3.178 ms  1.889 ms  1.988 ms
 2  34.34.34.2 (34.34.34.2)   3.914 ms  3.700 ms  2.975 ms
[rock1200]rock1200#
```

1.8 Step 5: Check the OSPF Configuration and Instances

Run the **`show configuration ospf`** command to verify the OSPF configuration in specific context. Because you are running OSPF in multiple contexts, run the **`show configuration ospf all-contexts`** command to verify the OSPF configuration in all contexts (instead of just running it from one context).

When examining a router configuration, use the following OSPF configuration checklist as a guide to isolate the fault.

Table 4 OSPF Configuration Checklist

#	Task	Checked
1	Do all interfaces have the correct addresses and masks?	
2	Are the interfaces enabled on the local router?	
3	Is OSPF configured on all neighboring interfaces, and do the OSPF parameters on the neighboring interfaces match?	
4	Are all the OSPF interfaces configured in the correct areas?	
5	Does the authentication type match on local and remote routers?	
6	Are both sides configured with the correct authentication key?	
7	Are both sides configured with matching area IDs? Neighboring interfaces must be in the same area to establish neighborhship.	
8	Did you check for any stub, transit, or NSSA mismatch?	
9	Do both endpoints must have the same interface type?	

Run the **show ospf global** command to view multiple routing OSPF instances in different routing contexts.

The following output displays the OSPF configuration in the context `rock1200`:

```
[rock1200]rock1200#show configuration ospf
Building configuration...

Current configuration:
!

!
!
context rock1200
!
 router ospf 1
  fast-convergence
  router-id 11.0.0.0
  area 0.0.0.0
   virtual-link 3.0.0.4 0.0.0.4
   interface to-jazz
  area 2.0.0.0
   area-type nssa
   interface to-two
  area 3.0.0.4
   interface to-three
!
! ** End Context **
!
end
[rock1200]rock1200#
```

The following output displays the OSPF configurations in all contexts:

```
[local]rock1200#show configuration ospf all-contexts
Building configuration...

Current configuration:
!
!  Configuration last changed by user 'test' at Tue Dec 14 06:41:04 2010
!

context local
!
 router ospf 1
  fast-convergence
  area 0.0.0.0
   interface loopback
   interface to-jazz
!
! ** End Context **
!

!
!
context rock1200
!
 router ospf 1
  fast-convergence
  router-id 11.0.0.0
  area 0.0.0.0
   virtual-link 3.0.0.4 0.0.0.4
   interface to-jazz
  area 2.0.0.0
   area-type nssa
   interface to-two
  area 3.0.0.4
   interface to-three
!
! ** End Context **
!
```

```
!
!
context two
!
 router ospf 2
  fast-convergence
  router-id 0.0.0.2
  area 2.0.0.0
   area-type nssa
   interface to-rock1200
  redistribute static
!
! ** End Context **
!

!
!
context three
!
 router ospf 3
  fast-convergence
  router-id 0.0.0.3
  area 3.0.0.4
   interface to-four
   interface to-rock1200
  redistribute static
!
! ** End Context **
!

!
!
context four
!
 router ospf 4
  fast-convergence
  router-id 0.0.0.4
  area 0.0.0.0
   interface to-five
   virtual-link 3.0.0.4 11.0.0.0
  area 3.0.0.4
   interface to-three
!
! ** End Context **
!

!
!
context five
!
 router ospf 5
  fast-convergence
  router-id 0.0.0.5
  area 0.0.0.0
   interface to-four
   interface to-jazz
!
! ** End Context **
!

!
!
context rock1200
!
 router ospf 1
  fast-convergence
  log-neighbor-up-down
!
! ** End Context **
!
end
[local]rock1200#
```

The following output displays the number of OSPF instances configured on the SmartEdge router in the local context:

```
[local]rock1200#show ospf global
  ---   OSPF Global Information   ---

Instance Count       : 6                Equal-Cost Paths     : 8
Sham Link Count      : 0                Schedule Delay usecs: 1000
Neighbors Exchanging: 0                 Exchanging Nbr Max   : 300
Restarted            : No               Restart reason       : Unknown
High Res Timers      : Yes              Receive Cfg EOF      : Yes
Shared Mem. Cleanup : No
[local]rock1200#
```

1.9 Step 6: Check OSPF Adjacency

When OSPF adjacency is formed, the router state changes as follows before it becomes fully adjacent with its neighbor:

- Down—No information has been received from the neighbor.
- Attempt—Attempt is valid for neighbors on NBMA network; the neighbor is being contacted, but no information has been received.
- Init—A Hello packet has been received from the neighbor, but the router is not listed in that Hello packet.
- 2-Way—The 2-way state indicates that the router has identified its own router ID in the Neighbor field of the neighbor's Hello packet. Receiving a database descriptor packet from a neighbor in the `init` state also causes a transition to the 2-way state. Receiving a hello with an OSPF neighbor in 2-way state is not a concern in broadcast multiaccess networks at the DR.. Other routers in these networks cannot progress beyond the 2-way state; in these networks, they achieve *Full* state only with Designated Router (DR) and Backup Designated Router (BDR). The adjacency must pass through 2-way state before proceeding to a *Full* state. Failure to pass through the 2-way state indicates there might be an error in the database exchange.
- ExStart—The routers and their Designated Router (DR) and Backup Designated Router (BDR) establish a master-slave relationship and choose the initial sequence number for adjacency formation.
- Exchange—OSPF routers exchange database descriptor (DD) packets. After the packet header, DD packets contain Link-state advertisement (LSA) headers only. They provide a summary of the sender's database contents to the receiving router.
- Loading—Routers send link-state request packets. During the database exchange, if a router receives a more recent or missing link-state advertisement (LSA), it requests that LSA by sending a link-state request packet.
- Full—Indicates that routers are fully adjacent with each other. All the router and network LSAs are exchanged, and the router databases are fully synchronized. `Full` is the normal state for an OSPF router. A router is stuck in another state indicates problems in forming adjacencies but not necessarily for broadcast and NBMA networks.

The following flowchart describes the general process for troubleshooting OSPF neighbor state issues.

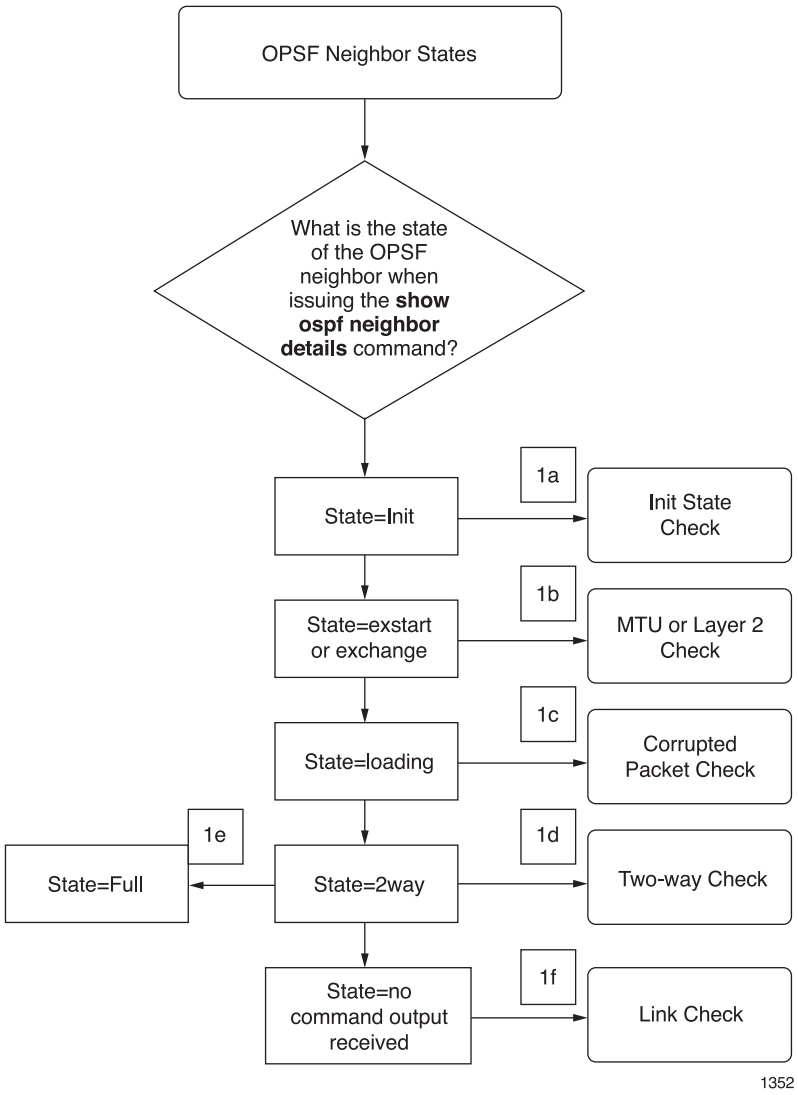

Figure 2 Troubleshooting OSPF Neighbor States

Before OSPF routers can exchange routing information, they must establish a neighborship.

Run the **show ospf neighbor** [*neighbor-id* | **interface** [*ip-addr* | *if-name*]] [**detail**] command to verify that every router has established neighborship.

Syntax	Description
neighbor-id	Optional. ID of the neighbor for which information is displayed.
interface	Optional. Displays information for the specified neighbor interface.
ip-addr	Optional. IP address of the interface.
if-name	Optional. Interface name.
detail	Optional. Displays detailed information.

1.9.1 Verify OSPF Neigbhorship

Run the **show ospf neighbor** command to verify that you have established adjacency with your neighbors. The configuration on both endpoints (the interface type; for example, point-to-multipoint or P2P) must match; otherwise, you cannot form an adjacency. This is a common issue.

Values other than `two-way` and `full` can indicate the following problems:

- Physical issues—If the MTU is too large, the routers attached to the link cannot negotiate the same MTU on the link, and an adjacency cannot form.

  For example, when the OSPF neighbor state is stuck in an `init` state, the router can receive packets from the neighbor but not transmit packets to the neighbor, resulting in a one-way link.

- Configuration Issues—If routers attached to the link cannot negotiate the link Area ID, the link Hello protocol timers, the link authentication type, or keys, adjacencies cannot form.

When examining adjacencies, use the following OSPF adjacency checklist to isolate the fault. For more information about how to troubleshoot OSPF adjacencies, see Section 2 on page 43.

Table 5 OSPF Neighborship Checklist

#	Task	Checked
1	Are Hello packets being sent from both neighbors?	
2	Are the dead timers set the same between neighbors?	
3	Are the interfaces configured on the same subnet (that is, do the address or mask pairs belong to the same subnet)?	
4	If authentication is being used, is the authentication type the same between neighbors? • Is authentication enabled on all routers within the area? • Are the passwords and the keys (in the case of MD5) the same?	
5	If the adjacency is across a virtual link, is the link configured within a stub area?	
6	Do neighbor MTUs match?	
7	Are router IDs unique within the entire internetwork? If not, no neighborship will form.	
8	Did you verify that the interface is not defined as passive in OSPF? If it is, no adjacency is formed. **Recommended Action**: Remove the **`passive`** command from the OSPF configuration.	
9	Is the network type the same for neighboring interfaces?	

The output of the **`show ospf neighbor`** command from `rock1200` indicates it has formed neighborships with routers `two`, `three`, `four`, and `jazz` as expected (has a `Full` state).

`Full` state indicates that routers are fully adjacent with each other. All the router and network LSAs are exchanged and the routers databases are fully synchronized. `Full` is the normal state for an OSPF router. `2-way` is a valid state for a neighbor on a broadcast or NBMA network when neither neighbor is Designated Router (DR) or Backup Designated Router (BDR). If a router is stuck in another state, there might be problems in forming adjacencies.

For more information about how to troubleshoot OSPF neighborhship, see Section 2 on page 43.

```
[rock1200]rock1200#show ospf neighbor

  --- OSPF Neighbors for Instance 1/Router ID 11.0.0.0 ---

NeighborID      NeighborAddress Pri State    DR-State IntfAddress     TimeLeft
0.0.0.4         34.34.34.2      1   Full     Other    0.0.0.4         0
22.0.0.0        11.11.11.2      1   Full     DR       11.11.11.1      35
0.0.0.2         12.12.12.2      1   Full     DR       12.12.12.1      39
0.0.0.3         13.13.13.2      1   Full     BDR      13.13.13.1      40
[rock1200]rock1200#
```

You can run the **`show ospf neighbor interface`** ***`if-name`*** command to display OSPF adjacency information for a specific interface.

The following output shows an adjacency that is *Up* on router `rock1200` on interface `to-jazz`.

```
[rock1200]rock1200#show ospf neighbor interface to-jazz

  --- OSPF Neighbors for Instance 1/Router ID 11.0.0.0 ---

NeighborID      NeighborAddress Pri State    DR-State IntfAddress     TimeLeft
22.0.0.0        11.11.11.2      1   Full     DR       11.11.11.1      34
```

1.9.2 Display Detailed OSPF Neighbor Information

Run the **`show ospf neighbor detail`** command to display detailed OSPF neighbor information. In the following output, all the expected neighbors are *Up* (in a Full state) and are working correctly. Full state indicates that routers are fully adjacent with each other.

```
[rock1200]rock1200#show ospf neighbor detail

  --- OSPF Neighbor 0.0.0.4 Area 0.0.0.0 Instance 1 ---

Address         : 34.34.34.2        Interface Addr : 0.0.0.4
State           : Full              DR State       : Other
Cost            : 2                 DR Priority    : 1
DR IP Address   : 0.0.0.0           BDR IP Address : 0.0.0.0
LSA Request     : 0                 LSA Retrans    : 0
DB Exchange     : 0                 Time Till Dead : 0
Hello Options   : E,DC              DD Options     : E,DC,O

  --- OSPF Neighbor 22.0.0.0 Area 0.0.0.0 Instance 1 ---

Address         : 11.11.11.2        Interface Addr : 11.11.11.1
State           : Full              DR State       : DR
Cost            : 1                 DR Priority    : 1
DR IP Address   : 11.11.11.2        BDR IP Address : 11.11.11.1
LSA Request     : 0                 LSA Retrans    : 0
DB Exchange     : 0                 Time Till Dead : 32
Hello Options   : E                 DD Options     : E,O

  --- OSPF Neighbor 0.0.0.2 Area 2.0.0.0 Instance 1 ---

Address         : 12.12.12.2        Interface Addr : 12.12.12.1
State           : Full              DR State       : DR
Cost            : 1                 DR Priority    : 1
DR IP Address   : 12.12.12.2        BDR IP Address : 12.12.12.1
LSA Request     : 0                 LSA Retrans    : 0
DB Exchange     : 0                 Time Till Dead : 40
Hello Options   : NP                DD Options     : NP,O

  --- OSPF Neighbor 0.0.0.3 Area 3.0.0.4 Instance 1 ---

Address         : 13.13.13.2        Interface Addr : 13.13.13.1
State           : Full              DR State       : BDR
Cost            : 1                 DR Priority    : 1
DR IP Address   : 13.13.13.1        BDR IP Address : 13.13.13.2
LSA Request     : 0                 LSA Retrans    : 0
DB Exchange     : 0                 Time Till Dead : 31
Hello Options   : E                 DD Options     : E,O

[rock1200]rock1200#
```

1.10 Step 7: Verify LSAs in the OSPF Database

Because OSPF is a link-state protocol, the link-state database should be the same for any router in the same area, except during brief periods of convergence. Each router has a link-state database with information about each router in the network and uses this information to build a network topology and calculate the best routes. All routers within the same area must have the same database content to help them identify the best routes. The LSA advertisement changes frequently, indicated by the fact that the sequence number is significantly higher than that of other LSAs.

Only the links that are appropriate for forwarding are included in the database. OSPF routers become fully adjacent with routers with which they have successfully completed the database synchronization process, during which OSPF routers exchange link-state information to populate their databases with the same information. An incomplete database results in an incomplete network view and, as a result, routing problems.

Use the OSPF database to draw a complete map of the network and observe the state of all the routers in the network. Examine the various LSAs; for example, if a link is unstable, the LSA advertising will change frequently, indicated by a sequence number that is significantly higher than that of other LSAs.

You can determine which parts of the network are changing the most by checking the LSA sequence numbers and the age of the LSA in the `LS Age` field. If the same LSA remains in the database, the LSA age increases. If LSAs are updated frequenly, the age remains low. The LSAge does not need to be the same in both router databases, but the other identifying elements of the LSA header should be the same.

The most common reasons for OSPF to not share the database information about a specific link include:

- The OSPF neighbor is not advertising routes.
- The OSPF neighbor (ABR) is not advertising the summary route.
- The OSPF neighbor is not advertising external routes.
- The OSPF neighbor is not advertising the default route.

To determine whether routers have a synchronized OSPF database, run the **`show ospf database`** command and compare the results of the shared areas:

1 Verify that the summaries of their LSA `checksum` fields are equal.

2 Determine that the two routers have the same number of LSAs (in the same area) in their link state databases.

If you have routing problems, make sure your interfaces are correctly configured.

Table 6 show ospf database Syntax Description

Field	Description
instance-id	Optional. OSPF instance ID. The range of values is 1 to 65,535.
area-id	Optional. Area ID. The range of values is 0 to 4,294,967,295.. Either a single integer or IP address format.
ip-addr	Optional. Area IP address. Either a single integer or IP address format.
databases-summary-network detail	Displays a count, grouped by type, of OSPF LSAs
database router	Displays information about OSPF router LSAs.
database network	Display information about OSPF network LSAs

When checking an area-wide issue, consider the following issues:

- Is the ABR correctly configured?
- Are all the routers configured for the same area type?
- When summarization is enabled, is it correctly configured?

  Two common problems are related to summarization in OSPF:

  – A router is not summarizing interarea routes.

  – A router is not summarizing external routes.

If you suspect that the link-state databases are corrupt or that the databases are not synchronized:

1 Run the **show ospf database database-summary** command to verify the number of LSAs in each router database. For a given area, the number of each LSA type should be the same in all routers.

2 Run the **show ospf database** command and verify that each LSA checksum is the same in a given area in every router database.

OSPF sends LSAs to all routers within the area. The LSA contains information about attached interfaces, link metrics and other variables.

In the following example, router `rock1200` and `two` are synchronized (all routers contain the same database) in area `2.0.0.0`. On context `rock1200`, the LSAs highlighted correctly match the LSAs in context `two` (router two).

To determine if you have the latest version of the LSAs, check the sequence number associated with the LSA in the **show ospf database** output. On each node, verify that each node OSPF database has synchronized this information to ensure that each router has the same view of your network.

```
[local]rock1200#context rock1200
[rock1200]rock1200#show ospf database

  --- OSPF Link State Database for Instance 1/Router ID 11.0.0.0 ---

             Router Link State Advertisements (Area 0.0.0.0)
LinkID           AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
0.0.0.4          0.0.0.4          800001c1    f5c5   E,DC     60     820
0.0.0.5          0.0.0.5          800001bf    5480   E,DC     60     92
11.0.0.0         11.0.0.0         800001c4    1eb2   E,DC     48     358
22.0.0.0         22.0.0.0         800001c2    ace4   E,DC     48     1703

             Network Link State Advertisements (Area 0.0.0.0)
LinkID           AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
11.11.11.2       22.0.0.0         800001b9    d64f   E,DC     32     217
15.15.15.1       22.0.0.0         800001ba    1b05   E,DC     32     1177

             Summary Network Link State Advertisements (Area 0.0.0.0)
LinkID           AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
12.12.12.0       11.0.0.0         800001bb    64b    E,DC     28     358
13.13.13.0       0.0.0.4          800001bb    371d   E,DC     28     250
13.13.13.0       11.0.0.0         800001bd    dd6e   E,DC     28     328
34.34.34.0       0.0.0.4          800001bb    36df   E,DC     28     305
34.34.34.0       11.0.0.0         800001bb    f419   E,DC     28     23

             Summary AS Border Router Link State Advertisements (Area 0.0.0.0)
LinkID           AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
0.0.0.3          0.0.0.4          800001ba    d8a0   E,DC     28     510
0.0.0.3          11.0.0.0         800001b8    91e2   E,DC     28     1338
11.0.0.0         0.0.0.4          800001b9    73fd   E,DC     28     1345

             Router Link State Advertisements (Area 2.0.0.0)
LinkID           AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
0.0.0.2          0.0.0.2          800001bd    fd1f   NP,DC    36     1225
11.0.0.0         11.0.0.0         800001bf    1cec   NP,DC    36     1408

             Network Link State Advertisements (Area 2.0.0.0)
LinkID           AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
12.12.12.2       0.0.0.2          800001ba    8ab9   NP,DC    32     1355

             Summary Network Link State Advertisements (Area 2.0.0.0)
LinkID           AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
11.11.11.0       11.0.0.0         800001b9    d37c   NP,DC    28     583
13.13.13.0       11.0.0.0         800001bb    87c0   NP,DC    28     1003
15.15.15.0       11.0.0.0         800001bb    49f7   NP,DC    28     483
34.34.34.0       11.0.0.0         800001bb    9a6d   NP,DC    28     358
45.45.45.0       11.0.0.0         800001ba    19cd   NP,DC    28     1138

             NSSA Link State Advertisements (Area 2.0.0.0)
LinkID           AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
0.0.0.0          11.0.0.0         800001bb    31b9   DC       36     898
88.88.88.88      0.0.0.2          800001ba    f570   NP,DC    36     845

             Router Link State Advertisements (Area 3.0.0.4)
LinkID           AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
0.0.0.3          0.0.0.3          800001c3    aa8c   E,DC     48     1356
0.0.0.4          0.0.0.4          800001be    6235   E,DC     36     1664
11.0.0.0         11.0.0.0         800001c1    b053   E,DC     36     203

             Network Link State Advertisements (Area 3.0.0.4)
LinkID           AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
13.13.13.1       11.0.0.0         800001bb    5be1   E,DC     32     888
34.34.34.1       0.0.0.3          800001ba    9677   E,DC     32     176

             Summary Network Link State Advertisements (Area 3.0.0.4)
LinkID           AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
11.11.11.0       0.0.0.4          800001b9    8dcd   E,DC     28     2019
11.11.11.0       11.0.0.0         800001b8    3027   E,DC     28     1478
12.12.12.0       0.0.0.4          800001bc    63f1   E,DC     28     1659
12.12.12.0       11.0.0.0         800001ba    84a    E,DC     28     298
15.15.15.0       0.0.0.4          800001bb    ee5f   E,DC     28     1234
15.15.15.0       11.0.0.0         800001ba    a5a2   E,DC     28     1828
45.45.45.0       0.0.0.4          800001bb    a84c   E,DC     28     689
45.45.45.0       11.0.0.0         800001bc    6f7b   E,DC     28     303

             External Link State Advertisements
```

```
LinkID          AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
88.88.88.88     11.0.0.0         800001b8    194e   E,DC     36     1523
99.99.99.99     0.0.0.3          800001b9    9ec9   E,DC     36     381


  --- OSPF Link State Database for Instance 1/Router ID 11.0.0.0 ---

            Router Link State Advertisements (Area 0.0.0.0)
LinkID          AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
0.0.0.4         0.0.0.4          800001c1    f5c5   E,DC     60     859
0.0.0.5         0.0.0.5          800001bf    5480   E,DC     60     131
11.0.0.0        11.0.0.0         800001c4    1eb2   E,DC     48     397
22.0.0.0        22.0.0.0         800001c2    ace4   E,DC     48     1742

            Network Link State Advertisements (Area 0.0.0.0)
LinkID          AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
11.11.11.2      22.0.0.0         800001b9    d64f   E,DC     32     256
15.15.15.1      22.0.0.0         800001ba    1b05   E,DC     32     1216

            Summary Network Link State Advertisements (Area 0.0.0.0)
LinkID          AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
12.12.12.0      11.0.0.0         800001bb    64b    E,DC     28     397
13.13.13.0      0.0.0.4          800001bb    371d   E,DC     28     289
13.13.13.0      11.0.0.0         800001bd    dd6e   E,DC     28     367
34.34.34.0      0.0.0.4          800001bb    36df   E,DC     28     344
34.34.34.0      11.0.0.0         800001bb    f419   E,DC     28     62

            Summary AS Border Router Link State Advertisements (Area 0.0.0.0)
LinkID          AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
0.0.0.3         0.0.0.4          800001ba    d8a0   E,DC     28     549
0.0.0.3         11.0.0.0         800001b8    91e2   E,DC     28     1377
11.0.0.0        0.0.0.4          800001b9    73fd   E,DC     28     1384

<< The LSAs highlighted in bold correctly match the LSAs in context two.

            Router Link State Advertisements (Area 2.0.0.0)
LinkID          AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
0.0.0.2         0.0.0.2          800001bd    fd1f   NP,DC    36     1264
11.0.0.0        11.0.0.0         800001bf    1cec   NP,DC    36     1447

            Network Link State Advertisements (Area 2.0.0.0)
LinkID          AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
12.12.12.2      0.0.0.2          800001ba    8ab9   NP,DC    32     1394

            Summary Network Link State Advertisements (Area 2.0.0.0)
LinkID          AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
11.11.11.0      11.0.0.0         800001b9    d37c   NP,DC    28     622
13.13.13.0      11.0.0.0         800001bb    87c0   NP,DC    28     1042
15.15.15.0      11.0.0.0         800001bb    49f7   NP,DC    28     522
34.34.34.0      11.0.0.0         800001bb    9a6d   NP,DC    28     397
45.45.45.0      11.0.0.0         800001ba    19cd   NP,DC    28     1177

            NSSA Link State Advertisements (Area 2.0.0.0)
LinkID          AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
0.0.0.0         11.0.0.0         800001bb    31b9   DC       36     937
88.88.88.88     0.0.0.2          800001ba    f570   NP,DC    36     884

            Router Link State Advertisements (Area 3.0.0.4)
LinkID          AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
0.0.0.3         0.0.0.3          800001c3    aa8c   E,DC     48     1395
0.0.0.4         0.0.0.4          800001be    6235   E,DC     36     1703
11.0.0.0        11.0.0.0         800001c1    b053   E,DC     36     242

            Network Link State Advertisements (Area 3.0.0.4)
LinkID          AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
13.13.13.1      11.0.0.0         800001bb    5be1   E,DC     32     927
34.34.34.1      0.0.0.3          800001ba    9677   E,DC     32     215

            Summary Network Link State Advertisements (Area 3.0.0.4)
LinkID          AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
11.11.11.0      0.0.0.4          800001b9    8dcd   E,DC     28     2058
11.11.11.0      11.0.0.0         800001b8    3027   E,DC     28     1517
12.12.12.0      0.0.0.4          800001bc    63f1   E,DC     28     1698
12.12.12.0      11.0.0.0         800001ba    84a    E,DC     28     337
15.15.15.0      0.0.0.4          800001bb    ee5f   E,DC     28     1273
15.15.15.0      11.0.0.0         800001ba    a5a2   E,DC     28     1867
```

```
45.45.45.0       0.0.0.4          800001bb    a84c   E,DC     28     728
45.45.45.0       11.0.0.0         800001bc    6f7b   E,DC     28     342

             External Link State Advertisements
LinkID           AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
88.88.88.88      11.0.0.0         800001b8    194e   E,DC     36     1562
99.99.99.99      0.0.0.3          800001b9    9ec9   E,DC     36     420
[rock1200]rock1200#]

[rock1200]rock1200#context two
[two]rock1200#show ospf database

  --- OSPF Link State Database for Instance 2/Router ID 0.0.0.2 ---

             Router Link State Advertisements (Area 2.0.0.0)
LinkID           AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
0.0.0.2          0.0.0.2          80000266    a9c9   NP,DC    36     1667
11.0.0.0         11.0.0.0         80000268    c797   NP,DC    36     1472

             Network Link State Advertisements (Area 2.0.0.0)
LinkID           AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
12.12.12.2       0.0.0.2          80000264    3465   NP,DC    32     262

             Summary Network Link State Advertisements (Area 2.0.0.0)
LinkID           AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
11.11.11.0       11.0.0.0         80000261    8126   NP,DC    28     1202
13.13.13.0       11.0.0.0         80000264    336b   NP,DC    28     822
15.15.15.0       11.0.0.0         80000263    f6a1   NP,DC    28     847
34.34.34.0       11.0.0.0         80000263    4817   NP,DC    28     1292
45.45.45.0       11.0.0.0         80000263    c478   NP,DC    28     102

             NSSA Link State Advertisements (Area 2.0.0.0)
LinkID           AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
0.0.0.0          11.0.0.0         80000264    dc64   DC       36     1847
88.88.88.88      0.0.0.2          80000262    a31a   NP,DC    36     1092
[two]rock1200#
```

1.11 Step 8: Check OSPF Routes

1.11.1 Verify OSPF Routes in the Control Plane

Run the **show ospf route** command to verify OSPF routes in the control plane. Use the **summary** keyword to display a summary of OSPF routes.

```
[rock1200]rock1200#show ospf route

       --- OSPF Routes for Instance 1/Router ID 11.0.0.0 ---

Destination        Type       Dest-Type/Proto   Cost  Nhops  Nhop
11.11.11.0/24      Intra      Net               1     1      to-jazz
12.12.12.0/24      Intra      Net               1     1      to-two
13.13.13.0/24      Intra      Net               1     1      to-three
15.15.15.0/24      Intra      Net               2     1      11.11.11.2
34.34.34.0/24      Intra      Net               2     1      13.13.13.2
45.45.45.0/24      Intra      Net               3     2      11.11.11.2
                                                             13.13.13.2
88.88.88.88/32     NSSA T2    Net               0     1      12.12.12.2
<--Redistributed routes that came from the NSSA area (area 2)

99.99.99.99/32     EXT T2     Net               0     1      13.13.13.2
[rock1200]rock1200#

[rock1200]rock1200#show ospf route summary

       --- OSPF Route Summary for Instance 1/Router ID 11.0.0.0 ---

Total routes          : 8            Redistributed routes  : 0
Intra-area routes     : 6            Inter-area routes     : 0
External type 1 routes: 0            External type 2 routes: 1
NSSA type 1 routes    : 0            NSSA type 2 routes    : 1
[rock1200]rock1200#
```

1.11.2 Verify OSPF Route Entries in the RIB

Run the **`show ip route ospf`** command to view all OSPF route entries in the RIB table (both active and standby paths).

```
[rock1200]rock1200#show ip route ospf

Codes: C - connected, S - static, S dv - dvsr, R - RIP, e B - EBGP, i B - IBGP
       A,H - derived hidden
       O   - OSPF, O3  - OSPFv3, IA - OSPF(v3) inter-area,
       N1  - OSPF(v3) NSSA external type 1, N2  - OSPF(v3) NSSA external type 2
       E1  - OSPF(v3) external type 1, E2  - OSPF(v3) external type 2
       i   - IS-IS, L1 - IS-IS level-1,  L2  - IS-IS level-2, N - NAT
       IPH - IP Host, SUB A - Subscriber address, SUB S - Subscriber static
       SUB P - AAA downloaded aggregate subscriber routes
       SUB N - Subscriber ND, SUB D - Subscriber DHCP-PD
       M F - Mobile Sub Foreign Agent, M H - Mobile Sub Home Agent
       M G - Mobile Sub GTP
       A - Derived Default, MeH - Media Nexthop
       >   - Active Route, * - LSP

Type    Network             Next Hop        Dist  Metric    UpTime  Interface
  O     11.11.11.0/24       11.11.11.1       110       1            to-jazz
  O     12.12.12.0/24       12.12.12.1       110       1            to-two
  O     13.13.13.0/24       13.13.13.1       110       1            to-three
> O     15.15.15.0/24       11.11.11.2       110       2      1w2d  to-jazz
> O     34.34.34.0/24       13.13.13.2       110       2      1w2d  to-three
> O     45.45.45.0/24       11.11.11.2       110       3      1w2d  to-jazz
>                           13.13.13.2                               to-three
> O N2  88.88.88.88/32      12.12.12.2       110       0      1w2d  to-two
> O E2  99.99.99.99/32      13.13.13.2       110       0      1w2d  to-three
[rock1200]rock1200#
```

1.11.3 View a Specific OSPF Route Entry

Run the **`show ospf route`** ***`address`*** command to view a specific OSPF route path—in this case, `11.11.11.0`, and the next hop taken in the OSPF route table through the `to-jazz` interface.

```
[rock1200]rock1200#show ospf route 11.11.11.0

OSPF longest prefix route lookup:

Instance ID     : 1                     Router ID       : 11.0.0.0
Destination     : 11.11.11.0/24         Type            : Intra
Dest-type       : Net                   Cost            : 1
SPF Version     : 35                    Distance        : 110
Area            : 0.0.0.0               Back Link Data  : 11.11.11.1
LSDB Type       : Net                   LSDB ID         : 11.11.11.2
LSDB Adv Router: 22.0.0.0               Next Hop Count  : 1
Next Hop 1 Intf: to-jazz                Next Hop 1 Addr: 11.11.11.1
Route Flags     : interface, intra-transit-net
[rock1200]rock1200#
```

1.12 Step 9: Verify IP Routes

This section describes the various commands to verify the route tables and OSPF route information in them. These commands are context-specific.

1.12.1 Verify the Active Routes in the RIB Table

Run the **show ip route** command to view the active (best) routes in the RIB. To view a specific address, specify the network prefix by using the ***ip-addr/prefix-length*** construct.

The following output shows an OSPF entry identified by the type value `O`.

The network `15.15.15.0/24` is reachable by the next hop `11.11.11.2` through the `to_jazz` interface.

The network `34.34.34.0/24` is reachable by the next hop `13.13.13.2` through the `to_three` interface.

The network `45.45.45.0/24` is reachable using two next hops, either `11.11.11.2` or `13.13.13.2`.

```
[rock1200]rock1200#show ip route

Codes: C - connected, S - static, S dv - dvsr, R - RIP, e B - EBGP, i B - IBGP
       O   - OSPF, O3  - OSPFv3, IA - OSPF(v3) inter-area,
       N1  - OSPF(v3) NSSA external type 1, N2  - OSPF(v3) NSSA external type 2
       E1  - OSPF(v3) external type 1, E2  - OSPF(v3) external type 2
       i   - IS-IS, L1 - IS-IS level-1,  L2  - IS-IS level-2, N - NAT
       IPH - IP Host, SUB A - Subscriber address, SUB S - Subscriber static
       SUB P - AAA downloaded aggregate subscriber routes
       SUB N - Subscriber ND, SUB D - Subscriber DHCP-PD
       M F - Mobile Sub Foreign Agent, M H - Mobile Sub Home Agent
       M G - Mobile Sub GTP
       A - Derived Default, MeH - Media Nexthop
       >   - Active Route, * - LSP

Type    Network             Next Hop        Dist  Metric    UpTime  Interface
> C     11.11.11.0/24                         0       0      1w2d  to-jazz
> C     12.12.12.0/24                         0       0      1w3d  to-two
> C     13.13.13.0/24                         0       0      1w3d  to-three
> O     15.15.15.0/24       11.11.11.2      110       2      1w2d  to-jazz
> O     34.34.34.0/24       13.13.13.2      110       2      1w2d  to-three
> O     45.45.45.0/24       11.11.11.2      110       3      1w2d  to-jazz

>                           13.13.13.2                             to-three
[rock1200]rock1200#
```

```
[rock1200]rock1200#show ip route 15.15.15.0
    Longest match Routing entry for 15.15.15.0/32 is 15.15.15.0/24 , version 16
    Route Uptime 1w5d
    Paths: total 1, best path count 1

    Route has been downloaded to following slots
    X-EP-NAME, 09/0

    Path information :

      Active path :
      Known via ospf 1, type-OSPF intra area, distance 110, metric 2,
      Tag 0, Next-hop 11.11.11.2, NH-ID 0x34500018, Adj ID: 0x8000021, Interface
 to-jazz
      Circuit 9/1:1023:63/1/2/9
[rock1200]rock1200#
```

To check redistribution problems for a static route:

1 Run the **show ip route static** command to view static routes.

2 Run the **show ospf route** command and verify that the static routes have been redistributed into OSPF.

1.12.2 Verify All Routes Stored in the RIB

Run the **`show ip route all`** command to view all the routes stored in the RIB from all routing protocols.

The following example displays the routes stored in the RIB on router `rock1200`.

```
[rock1200]rock1200#show ip route all
Codes: C - connected, S - static, S dv - dvsr, R - RIP, e B - EBGP, i B - IBGP
       A,H - derived hidden
       O   - OSPF, O3  - OSPFv3, IA - OSPF(v3) inter-area,
       N1  - OSPF(v3) NSSA external type 1, N2  - OSPF(v3) NSSA external type 2
       E1  - OSPF(v3) external type 1, E2  - OSPF(v3) external type 2
       i   - IS-IS, L1 - IS-IS level-1,  L2  - IS-IS level-2, N - NAT
       IPH - IP Host, SUB A - Subscriber address, SUB S - Subscriber static
       SUB P - AAA downloaded aggregate subscriber routes
       SUB N - Subscriber ND, SUB D - Subscriber DHCP-PD
       M F - Mobile Sub Foreign Agent, M H - Mobile Sub Home Agent
       M G - Mobile Sub GTP
       A - Derived Default, MeH - Media Nexthop
       >   - Active Route, * - LSP

Type    Network             Next Hop        Dist  Metric    UpTime  Interface
> C     11.11.11.0/24                         0       0      1w2d  to-jazz
  O     11.11.11.0/24       11.11.11.1      110       1            to-jazz
> C H   11.11.11.0/32                        0       0      1w2d  to-jazz
> C H   11.11.11.1/32                        0       0      1w2d  to-jazz
> A H   11.11.11.2/32       11.11.11.2      254       0      1w2d  to-jazz
> C H   11.11.11.255/32                      0       0      1w2d  to-jazz
> C     12.12.12.0/24                       0       0      1w3d  to-two
  O     12.12.12.0/24       12.12.12.1      110       1            to-two
> C H   12.12.12.0/32                       0       0      1w3d  to-two
> C H   12.12.12.1/32                       0       0      1w3d  to-two
> A H   12.12.12.2/32       12.12.12.2      254       0      1w3d  to-two
> C H   12.12.12.255/32                     0       0      1w3d  to-two
> C     13.13.13.0/24                       0       0      1w3d  to-three
  O     13.13.13.0/24       13.13.13.1      110       1            to-three
> C H   13.13.13.0/32                       0       0      1w3d  to-three
> C H   13.13.13.1/32                       0       0      1w3d  to-three
> A H   13.13.13.2/32       13.13.13.2      254       0      1w3d  to-three
> C H   13.13.13.255/32                     0       0      1w3d  to-three
> O     15.15.15.0/24       11.11.11.2      110       2      1w2d  to-jazz
> O     34.34.34.0/24       13.13.13.2      110       2      1w2d  to-three
> O     45.45.45.0/24       11.11.11.2      110       3      1w2d  to-jazz
>                           13.13.13.2                             to-three
[rock1200]rock1200#
```

1.13 Step 10: Check OSPF Statistics

Run the **show ospf statistics [*instance-id*] [detail]** command to verify traffic statistics accumulated for all OSPF processes running on the router.

Run the **show ospf statistics interface *if-name*]** command to verify OSPF statistics at the interface level. In the `Drop` column, check for dropped packets.

Run the **show port counters detail** command to check the interface counters. For more information about this command, see the *General Troubleshooting Guide*.

The following example displays output from the **show ospf statistics** command:

```
[rock1200]rock1200#show ospf statistics

  --- OSPF Statistics for Instance 1 ---

Max flood queue length   : 3           Interval            : 1w 3d 00:46:35
LSAs received            : 8891        LSAs sent           : 9335
LSAs changes received    : 36          LSA Retransmissions : 3
Packet Retransmissions   : 3           RIB initializations : 1
Routes downloaded        : 24          Routes deleted      : 3
Host Routes Queued       : 11          Normal Routes Queued: 16
Priority Routes Queued   : 0           Total Routes Queued : 27
Download Errors          : 0           RIB IPC messages    : 20
Download SPF Delays      : 0           SPF Download Delays : 0
DC Indicate originated   : 0           DC Indicate purged  : 0
DC DoNotAge purged       : 0

              Hello    DD       LSR      LSU      ACK
        Sent  259511   11       4        9253     8804
        Recv  258837   14       4        8843     9242
[rock1200]rock1200#


[rock1200]rock1200#show ospf statistics interface to-jazz

  --- OSPF Statistics for Interface 11.11.11.1/Instance 1 ---


Interface Events    : 4           Interval            : 3d 22:06:25

LSAs received       : 1595        LSAs sent           : 1230

Max ACK queue length: 11          Max flood queue size: 6


              Hello    DD       LSR      LSU      ACK

        Sent  33879    3        1        1212     1547

        Recv  33879    2        0        1565     1208

[rock1200]rock1200#
```

1.14 Step 11: Examine the OSPF SPF Log

Run the **`show ospf spf log`** command to display the SPF calculation timing calculation log. Verify that SPF ran and investigate network instabilities. The SPF log includes a description of what triggered SPF recalculations—for example, when a new LSA arrives. This command shows which LSAs change most frequently and what triggers the SPF calculations.

Note: If the SPF log is not running as expected, check the following SPF scheduling configuration parameters:

- fast-convergence
- spf-timers

```
[rock1200]rock1200#show ospf spf log

   --- OSPF SPF Route Calculation Timing Log  ---

     Maximum SPF-Phase Timings
When (elapsed)  Instance/Area          Phase     Duration
 1w 2d 23:53:28 1/3.0.0.4              Intra     1 ms
 1w 2d 23:18:44 1/3.0.0.4              Summary   < 1 ms
 1w 2d 23:19:38 1/N/A                  External  1 ms
 1w 2d 23:18:44 1/2.0.0.0              NSSA      < 1 ms

     Most Recent SPF-Phase Timings
When (elapsed)  Instance/Area          Phase     Duration
 1w 2d 23:18:45 1/2.0.0.0              NSSA      < 1 ms
 1w 2d 23:18:45 1/N/A                  External  < 1 ms
 1w 2d 23:18:45 1/3.0.0.4              Summary   < 1 ms
 1w 2d 23:18:45 1/0.0.0.0              Summary   < 1 ms
 1w 2d 23:18:45 1/0.0.0.0              Intra     < 1 ms
 1w 2d 23:18:45 1/3.0.0.4              Intra     < 1 ms
 1w 2d 23:18:45 1/2.0.0.0              Intra     < 1 ms
 1w 2d 23:18:49 1/2.0.0.0              NSSA      < 1 ms
 1w 2d 23:18:49 1/N/A                  External  < 1 ms
 1w 2d 23:18:49 1/3.0.0.4              Summary   < 1 ms
 1w 2d 23:18:49 1/0.0.0.0              Summary   < 1 ms
 1w 2d 23:18:49 1/0.0.0.0              Intra     < 1 ms
 1w 2d 23:18:49 1/3.0.0.4              Intra     < 1 ms
 1w 2d 23:18:49 1/2.0.0.0              Intra     < 1 ms
 1w 2d 23:18:54 1/2.0.0.0              NSSA      < 1 ms
 1w 2d 23:18:54 1/N/A                  External  < 1 ms
 1w 2d 23:18:54 1/3.0.0.4              Summary   < 1 ms
 1w 2d 23:18:54 1/0.0.0.0              Summary   < 1 ms
 1w 2d 23:18:54 1/0.0.0.0              Intra     < 1 ms
 1w 2d 23:18:54 1/3.0.0.4              Intra     < 1 ms
 1w 2d 23:18:54 1/2.0.0.0              Intra     < 1 ms
 1w 2d 23:18:55 1/2.0.0.0              NSSA      < 1 ms
 1w 2d 23:18:55 1/N/A                  External  < 1 ms
 1w 2d 23:18:55 1/3.0.0.4              Summary   < 1 ms
 1w 2d 23:18:55 1/0.0.0.0              Summary   < 1 ms
 1w 2d 23:18:55 1/0.0.0.0              Intra     < 1 ms
 1w 2d 23:18:55 1/3.0.0.4              Intra     < 1 ms
 1w 2d 23:18:55 1/2.0.0.0              Intra     < 1 ms
 1w 2d 23:18:59 1/2.0.0.0              NSSA      < 1 ms
 1w 2d 23:18:59 1/N/A                  External  < 1 ms

 . . .
```

1.15 Step 12: Check OSPF Logs

Use the **show log | grep ospf** command to filter the log for entries related to OSPF.

Note: You must configure the **log-neighbor-up-down** command to log an informational message when a neighbor transitions to or from full adjacency state.

1.16 Step 13: Monitor OSPF Events

Run the **monitor ospf** commands to monitor OSPF events in real time. You can use these commands to troubleshoot intermittent issues.

Table 7 Tasks to Monitor OSPF Events

Command	Description
monitor ospf interface	Display continuously updated information about OSPF interfaces.
monitor ospf neighbors	Display continuously updated information about OSPF neighbors.
monitor ospf spf last	Display continuously updated information about OSPF statistics.
monitor ospf spf statistics	Display continuously updated information about the most recent OSPF SPF calculation.

1.16.1 Monitor OSPF Interface

Run the **monitor ospf interface** command to display continuously updated information about OSPF interfaces.

```
[rock1200]rock1200#monitor ospf interface
  --- OSPF Neighbors for Instance 1/Router ID 11.0.0.0 ---

NeighborID      NeighborAddress Pri State    DR-State IntfAddress     TimeLeft
0.0.0.4         34.34.34.2      1   Full     Other    0.0.0.4         0
22.0.0.0        11.11.11.2      1   Full     DR       11.11.11.1      37
0.0.0.2         12.12.12.2      1   Full     DR       12.12.12.1      35
0.0.0.3         13.13.13.2      1   Full     BDR      13.13.13.1      36

% enter ctrl-C to exit monitor mode, monitor duration(sec): 600   (00:00:02)
```

1.16.2 Monitor OSPF Neighbors

Run the **monitor ospf neighbors** command to display continuously updated information about OSPF neighbors.

```
[rock1200]rock1200#monitor ospf neighbors
 --- OSPF Neighbors for Instance 1/Router ID 11.0.0.0 ---

NeighborID      NeighborAddress Pri State    DR-State IntfAddress     TimeLeft
0.0.0.4         34.34.34.2      1   Full     Other    0.0.0.4         0
22.0.0.0        11.11.11.2      1   Full     DR       11.11.11.1      36
0.0.0.2         12.12.12.2      1   Full     DR       12.12.12.1      34
0.0.0.3         13.13.13.2      1   Full     BDR      13.13.13.1      35

% enter ctrl-C to exit monitor mode, monitor duration(sec): 600   (00:00:02)
```

1.16.3 Monitor OSPF SPF Statistics

Run the **`monitor ospf spf last`** to display continuously updated information about the most recent OSPF SPF route calculation.

```
[rock1200]rock1200#monitor ospf spf last
  --- Most Recent OSPF SPF Route Calculation ---

When (elapsed)  Instance/Area          Phase     Duration
 2w 1d 00:28:38 1/2.0.0.0              NSSA      < 1 ms
 2w 1d 00:28:38 1/N/A                  External  < 1 ms
 2w 1d 00:28:38 1/3.0.0.4              Summary   < 1 ms
 2w 1d 00:28:38 1/0.0.0.0              Summary   < 1 ms
 2w 1d 00:28:38 1/0.0.0.0              Intra     < 1 ms
 2w 1d 00:28:38 1/3.0.0.4              Intra     < 1 ms
 2w 1d 00:28:38 1/2.0.0.0              Intra     < 1 ms

% enter ctrl-C to exit monitor mode, monitor duration(sec): 600   (00:00:02)
```

1.16.4 Monitor OSPF Statistics

Run the **`monitor ospf statistics`** to display continuously updated information about OSPF traffic statistics.

Make sure the `Sent` and `Received` columns are incrementing. In the `Drop` column, check for drop packets.

```
[rock1200]rock1200#monitor ospf statistics
 --- OSPF Statistics for Instance 1 ---

Max flood queue length   : 3           Interval            : 1w 3d 01:19:08
LSAs received            : 8912        LSAs sent           : 9357
LSAs changes received    : 36          LSA Retransmissions : 3
Packet Retransmissions   : 3           RIB initializations : 1
Routes downloaded        : 24          Routes deleted      : 3
Host Routes Queued       : 11          Normal Routes Queued: 16
Priority Routes Queued   : 0           Total Routes Queued : 27
Download Errors          : 0           RIB IPC messages    : 20
Download SPF Delays      : 0           SPF Download Delays : 0
DC Indicate originated   : 0           DC Indicate purged  : 0
DC DoNotAge purged       : 0

              Hello    DD       LSR      LSU      ACK
        Sent  260097   11       4        9274     8825
        Recv  259422   14       4        8864     9263

% enter ctrl-C to exit monitor mode, monitor duration(sec): 600   (00:00:04)
```

1.17 Step 14: Debug OSPF

Run the **debug opsf packet errors** command to check for MTU, authentication, interface, and area ID issues.

The following table lists error messages that can appear when you run this command, along with the recommended actions to resolve the issues.

Table 8 Debug OSPF Packet Error Messages

#	Error Message
1	"OSPF-1 NBR 1.2.3.4: MTU 567 less than our mtu (1500)" **Recommended Action:** 1. Check MTU configuration by running the **show ip interface** command on the interfaces that have the mismatched MTUs. 2. Change the router MTU to match the neighbor MTU.
2	"OSPF-1: Duplicate DD packet from slave NBR 1.2.3.4 flags: M,M,I options: 2 seq: 3450293" **Recommended Action:** The slave in a DD exchange has received a packet with a sequence earlier than the one it already has, which means there is a duplicate. Verify that the values are what you expect for the following options. The values need to match the interfaces between two communicating routers. • The initialize(I), more (M) and master(MS) bits—Options field, and DD sequence number contained in the last Database. Description packet received from the neighbor. Used to determine whether the next Database Description packet received from the neighbor is a duplicate. • ExternalRoutingCapability—Determines if this is a stub area. • E-bit —Describes the way AS-external-LSAs are flooded. • MC-bit —Describes whether IP multicast datagrams are forwarded. • N/P-bit —This bit describes the handling of Type-7 LSAs. • EA-bit —Describes the router's willingness to receive and forward. • External-Attributes-LSAs —Describes the router's handling of demand circuits.
3	"OSPF-1: Duplicate DD packet from master NBR 1.2.3.4 flags: M,M,I options: 2 seq: 345234098" **Recommended Action:** The master in a DD exchange has received a packet with a sequence earlier than the one it already has, which means there is a duplicate. Verify that the values are what you expect for the following options. The values need to match the interfaces between two communicating routers. • The initialize(I), more (M) and master(MS) bits—Options field, and DD sequence number contained in the last Database. Description packet received from the neighbor. Used to determine whether the next Database Description packet received from the neighbor is a duplicate. • ExternalRoutingCapability—Determines if this is a stub area. • E-bit —Describes the way AS-external-LSAs are flooded. • MC-bit —Describes whether IP multicast datagrams are forwarded. • N/P-bit —This bit describes the handling of Type-7 LSAs. • EA-bit —Describes the router's willingness to receive and forward. • External-Attributes-LSAs —Describes the router's handling of demand circuits.

Table 8 Debug OSPF Packet Error Messages

#	Error Message
4	`"OSPF-1 NBR 1.2.3.4: invalid LS recv Type 5 (AS-Ext) for this area/interface"`
	Recommended Action: Make sure the area is not a stub or an NSSA area. Type 5 LSAs are not allowed in these areas.
5	`"OSPF-1 LSU AS-Ext 9.10.11.12 [5.6.7.8] from 1.2.3.4: invalid type"`
	Recommended Action: Make sure the area is not a stub or an NSSA area. Type 5 LSAs are not allowed in these areas.
6	`"OSPF-1: No current keys in key-chain somePolicyName interface 11.0.0.1 (someIntfName)"`
	Recommended Action: Check the authentication configuration.
7	Dec 13 09:37:11: %OSPF-7-PCK_ERRORS: Recv invalid Hello packet 13.1.1.1->224.0.0.5 area 0.0.0.0: No interface
	Recommended Action: Verify that the interface for which you are receiving the Hello packet has OSPF enabled.
8	Nov 22 23:14:55.632: [0001]: %OSPF-7-PCK_ERRORS: OSPF-1: Recv int-to-UELKSRPP01 invalid Database Description packet 10.0.9.14->224.0.0.5 area 0.0.0.0: BFD DOWN Nov 22 23:14:11.331: [0001]: %OSPF-7-PCK_ERRORS: OSPF-1: Recv int-to-UELKSRPP01 invalid Hello packet 10.0.9.14->224.0.0.5 area 0.0.0.0: BFD DOWN
	Recommended Action: Check if BFD keepalives are being exchanged on OSPF neighbor interfaces. If not flap the ports or check if the BFD configuration is complete.

Caution!

Risk of performance loss. Enabling the generation of debug messages can severely affect system performance. To reduce the risk, exercise caution when enabling the generation of debug messages on a production system.

We highly recommend that you run this command during a maintenance window.

2 Troubleshooting Specific OSPF Issues

This section describes how to troubleshoot specific OSPF routing problems. For more information about troubleshooting OSPF adjacency, see Troubleshooting General OSPF Issues.

2.1 Troubleshooting OSPF Neighbor States

Use the following sample OSPF topology as a guide to troubleshoot specific OSPF issues that follow in subsequent sections.

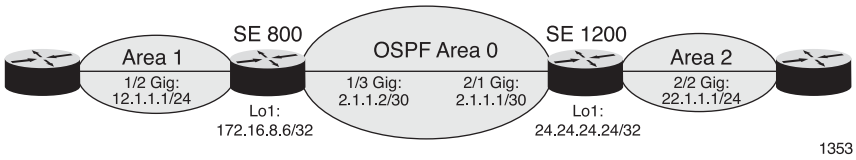

Figure 3 Sample OSPF Topology for Troubleshooting Specific Issues

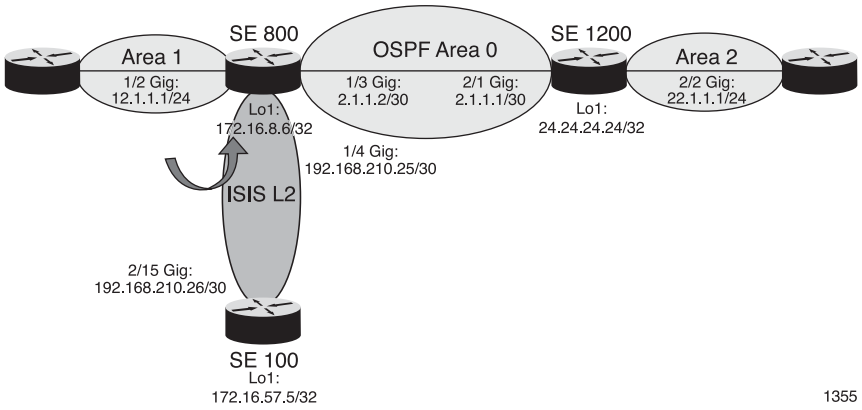

Figure 4 External Routes

2.1.1 Init State

The `Init` state specifies that the router has received a Hello packet from its neighbor, but the receiving router ID is not included in the Hello packet. When a router receives a Hello packet from a neighbor, it must list the sender's router ID in its Hello packet as an acknowledgment that it received a valid Hello packet.

When the neighbor is stuck in the `Init` state, a local router likely is not listed in a neighbor's Hello packets because the neighbor has not received Hello packets from the local router. Run the **`ping`** and **`traceroute`** commands to verify that the links between routers are operational. If a ping between routers is not successful, the link is not functioning properly and you need to troubleshoot it.

Consider the following when troubleshooting the `Init` state:

- If any access lists are defined on the neighbor interface, the destination IP of `224.0.0.5` must be permitted in the access list. OSPF Hello packets have a destination address of `224.0.0.5` (the **`all ospf routers multicast address`**).

  Recommended Action: To verify the access list, run the **`show access-group`** command.

- There may be a Layer 2 or configuration problem preventing multicast packets from reaching the neighboring router.

  Recommended Action: Type the **`ping`** command to the multicast address to see if responses are received from the neighboring routers if a ping to `224.0.0.5` from the `rock1200` context on the `rock1200` router does not return a response.

- Authentication is not enabled on both sides. The router on which authentication is not enabled still processes Hello packets from the neighbor and sees the neighbor in the `init` state.

  Recommended Action: To correct this problem, enable authentication on both sides. Make sure they are using the same authentication type.

- Hellos are getting lost on one side at Layer 2.

- For virtual links, authentication is enabled on only one side.
  Recommended Action: Enable authentication both sides.

In the following in example, check the state and reachablity of the OSPF neighbor.

```
[local]SE800_10.192.16.82#show ospf neighbor detail
  --- OSPF Neighbor 24.24.24.24 Area 0.0.0.0 Instance 100 ---
Address         : 2.1.1.1           Interface Addr : 2.1.1.2
State           : Init              DR State       : Other
Cost            : 1                 DR Priority    : 1
DR Router ID    : 0.0.0.0           BDR Router ID  : 0.0.0.0
LSA Request     : 0                 LSA Retrans    : 0
DB Exchange     : 0                 Time Till Dead : 39
Hello Options   : E,O               DD Options     :


[local]SE800_10.192.16.82#ping 2.1.1.1

PING 2.1.1.1 (2.1.1.1): source 2.1.1.2, 36 data bytes,
timeout is 1 second
!!!!!

----2.1.1.1 PING Statistics----
5 packets transmitted, 5 packets received, 0.0% packet loss
round-trip min/avg/max/stddev = 1.638/2.219/2.928/0.595 ms
[local]SE800_10.192.16.82#tra
[local]SE800_10.192.16.82#traceroute 2.1.1.1
se_traceroute to 2.1.1.1 (2.1.1.1), 30 hops max, 40 byte packets
 1  2.1.1.1 (2.1.1.1)   5.592 ms  3.223 ms  2.796 ms
[local]SE800_10.192.16.82#
```

In the following example, you run the **`ping size`** command to determine if the entire MTU can be used. Sometimes both routers have the same configured IP MTU but network equipment between the two routers will not pass the packets of maximum size.

You then run the **`show ospf interface detail`** command to display detailed information about the OSPF interface.

```
[local]SE800_10.192.16.82#ping 2.1.1.1 size 1500
PING 2.1.1.1 (2.1.1.1): source 2.1.1.2, 1500 data bytes,
timeout is 1 second
!!!!!   <--Indicates a successful ping.

----2.1.1.1 PING Statistics----
5 packets transmitted, 5 packets received, 0.0% packet loss
round-trip min/avg/max/stddev = 1.909/2.650/3.817/0.794 ms
[local]SE800_10.192.16.82#
```

```
[local]SE800_10.192.16.82#show ospf interface detaill

  --- OSPF Interface 2.1.1.2 Area 0.0.0.0 Instance 100 ---

Network Type    : broadcast          Mask            : 255.255.255.252
Cost            : 1                  Logical Intf    : 1/3
MTU             : 1500               Physical Intf   : ethernet 1/3
State           : BDR                Priority        : 1
Hello Interval  : 10                 Dead Interval   : 40
Transmit Delay  : 1                  Retransmit Int  : 5
DR Router ID    : 24.24.24.24        DR IP Address   : 2.1.1.1
BDR Router ID   : 2.1.1.2            BDR IP Address  : 2.1.1.2
Ack Queued      : 0                  Flood Queued    : 0
Ack Delay       : 2                  Authentication  : None  <--Make sure both sides use the
                                                             <--same authentication type.
LSA Count       : 0                  LSA Checksum    : 0
Demand Circuit  : No                 Flood Reduction: No
Neighbor Count  : 1

Neighbor List (1 Adjacent):
24.24.24.24
[local]SE800_10.192.16.82#
```

2.1.2 2-Way State

The 2-way state indicates that the router has seen its own router ID in the Neighbor field of the neighbor's Hello packet. Receiving a Database Description (DD) packet from a neighbor in the `init` state will also a cause a transition to 2-way state. The OSPF neighbor 2-way state is not a cause for concern.

This state indicates that bidirectional communication has been established between two routers—each router has seen the other's Hello packet, and the router receiving the Hello packet sees its own router ID in the received Hello packet's neighbor field. At this stage, a router can establish adjacency with this neighbor. On broadcast media and nonbroadcast multiaccess networks (NBMA), a router reaches full state only with the designated router (DR) and the backup designated router (BDR); it stays in the 2-way state with all other neighbors. On point-to-point and point-to-multipoint networks, a router reaches full state with all connected routers.

At the end of this stage, the DR and BDR for broadcast and nonbroadcast multiaccess networks are elected.

If the router interfaces are not in a waiting state, the router performs DR and BDR election. Once DR and BDR are elected, a router attempts to form a full adjacency with a neighbor if one of the two routers is the DR or BDR. OSPF routers become fully adjacent to routers with which they have successfully completed the database synchronization process. This is how OSPF routers within an area exchange link-state information to populate their databases with the same information. This database synchronization process occurs only if one of the two routers is a DR or BDR in the case of broadcast multiaccess networks. This database synchronization process is only executed between two routers if one of the two routers is the DR or BDR.

Note: If the OSPF neighbor is stuck in 2-way state, router-priority `0` might be configured on all routers.

The following example displays the state of OSPF interfaces.

```
[local]SE800_10.192.16.82#show ospf interface detail

  --- OSPF Interface 2.1.1.2 Area 0.0.0.0 Instance 100 ---

Network Type   : broadcast          Mask           : 255.255.255.252
Cost           : 1                  Logical Intf   : 1/3
MTU            : 1500               Physical Intf  : ethernet 1/3
State          : BDR                Priority       : 1
Hello Interval : 10                 Dead Interval  : use
Transmit Delay : 1                  Retransmit Int : 5
DR Router ID   : 24.24.24.24        DR IP Address  : 2.1.1.1
BDR Router ID  : 2.1.1.2            BDR IP Address : 2.1.1.2
Ack Queued     : 0                  Flood Queued   : 0
Ack Delay      : 2                  Authentication : None
LSA Count      : 0                  LSA Checksum   : 0
Demand Circuit : No                 Flood Reduction: No
Neighbor Count : 1

Neighbor List (1 Adjacent):
24.24.24.24
[local]SE800_10.192.16.82#
--- OSPF Neighbor 24.24.24.24 Area 0.0.0.0 Instance 100 ---

Address        : 2.1.1.1            Interface Addr : 2.1.1.2
State          : 2-Way              DR State       : Other
Cost           : 1                  DR Priority    : 1
DR Router ID   : 0.0.0.0            BDR Router ID  : 0.0.0.0
LSA Request    : 0                  LSA Retrans    : 0
DB Exchange    : 0                  Time Till Dead : 38
Hello Options  : E,O                DD Options     :
[local]SE800_10.192.16.82#sh ospf neig detail

  --- OSPF Neighbor 24.24.24.24 Area 0.0.0.0 Instance 100 ---

Address        : 2.1.1.1            Interface Addr : 2.1.1.2
State          : 2-Way              DR State       : Other
Cost           : 1                  DR Priority    : 1
DR Router ID   : 0.0.0.0            BDR Router ID  : 0.0.0.0
LSA Request    : 0                  LSA Retrans    : 0
DB Exchange    : 0                  Time Till Dead : 36
Hello Options  : E,O                DD Options     :
[local]SE800_10.192.16.82#
```

2.1.3 OSPF Neighbor Stuck in Exstart or Exchange State

A router gets stuck in ExStart or Exchange state more frequently when interoperating with another vendor's router.

Once the DR and BDR are elected, link-state information exchange can start between the routers and their DR and BDR. In the ExStart state, the routers and their DR and BDR establish a master-slave relationship and choose the initial sequence number to establish an adjacency. The router with the higher router ID becomes the master and starts the exchange, and is the only router that can increment the sequence number.

Note: Make sure that the router ID is unique within the entire internetwork; otherwise, OSPF neigborship will "not" form.

In the Exchange state, OSPF routers exchange DD packets. DD contain link-state advertisement (LSA) headers only and describe the contents of the entire link-state database. Each DD packet has a sequence number that can be incremented only by a master that is explicitly acknowledged by the slave. Routers also send link-state request packets and link-state update packets (which contain the entire LSA) in this state. The contents of the DD packets received are compared to the information contained in the routers link-state database to check if new or more current link-state information is available with the neighbor.

OSPF neighbors in Exstart or exchange state are trying to exchange DD packets. The adjacency should continue past this state. If it does not, there is a problem with the DD exchange, such as a maximum transmission unit (MTU) mismatch or receipt of an unexpected DD sequence number.

Note: Make sure that the router ID is unique within the entire internetwork; otherwise, no neighborship will form.

An OSPF neighbor can be stuck in ExStart or Exchange state for the following reasons:

- Duplicate router IDs on neighbors. This problem usually manifests before database exchange.
- Unable to ping without using specific maximum transmission unit (MTU) size.
- Broken unicast connectivity because of the following:
  - Wrong VC or DLCI mapping in the Frame Relay or ATM switch
  - An access list blocking the unicast
  - Incorrect NAT translation of the unicast packet
- Mismatched interface MTU

The problem occurs when the MTU settings for neighboring router interfaces do not match. If the router with the higher MTU sends a packet larger than the MTU set on the neighboring router, the neighboring router ignores the packet. When this occurs, the output of the **`show ospf neighbor detail`** command displays output similar to what is shown below.

```
[local]SE800_10.192.16.82#show ospf neighbor detail

  --- OSPF Neighbor 24.24.24.24 Area 0.0.0.0 Instance 100 ---

Address         : 2.1.1.1           Interface Addr : 2.1.1.2
State           : ExStart           DR State       : DR
Cost            : 1                 DR Priority    : 1
DR Router ID    : 2.1.1.1           BDR Router ID  : 2.1.1.2
LSA Request     : 0                 LSA Retrans    : 0
DB Exchange     : 0                 Time Till Dead : 36
Hello Options   : E,O               DD Options     : E,O

[local]SE800_10.192.16.82#show ospf neighbor detail

  --- OSPF Neighbor 24.24.24.24 Area 0.0.0.0 Instance 100 ---

Address         : 2.1.1.1           Interface Addr : 2.1.1.2
State           : Exchange          DR State       : DR
Cost            : 1                 DR Priority    : 1
DR Router ID    : 2.1.1.1           BDR Router ID  : 2.1.1.2
LSA Request     : 0                 LSA Retrans    : 0
DB Exchange     : 3                 Time Till Dead : 31
Hello Options   : E,O               DD Options     : E,O

[local]SE800_10.192.16.82#show ospf neighbor detail
--- OSPF Neighbor 24.24.24.24 Area 0.0.0.0 Instance 100 ---

Address         : 2.1.1.1           Interface Addr : 2.1.1.2
State           : Exchange          DR State       : DR
Cost            : 1                 DR Priority    : 1
DR Router ID    : 2.1.1.1           BDR Router ID  : 2.1.1.2
LSA Request     : 0                 LSA Retrans    : 0
DB Exchange     : 3                 Time Till Dead : 38
Hello Options   : E,O               DD Options     : E,O

local]SE800_10.192.16.82#show ospf neighbor detail

  --- OSPF Neighbor 24.24.24.24 Area 0.0.0.0 Instance 100 ---

Address         : 2.1.1.1           Interface Addr : 2.1.1.2
State           : Exchange          DR State       : DR
Cost            : 1                 DR Priority    : 1
DR Router ID    : 2.1.1.1           BDR Router ID  : 2.1.1.2
LSA Request     : 0                 LSA Retrans    : 0
DB Exchange     : 3                 Time Till Dead : 33
Hello Options   : E,O               DD Options     : E,O
[local]SE800_10.192.16.82#
```

After you have determined that OSPF is stuck in the Exchange or ExStart state, run the **debug ospf packet errors** command to determine the cause of the fault.

In the following example, the **debug ospf packet errors** output indicates mismatched MTUs.

Recommended Action:

1. Run the **show ip interface** command on the interfaces that have the mismatched MTUs.
2. Change either router MTU to match the neighbor MTU.

Caution!

Risk of performance loss. Enabling the generation of debug messages can severely affect system performance. To reduce the risk, exercise caution when enabling the generation of debug messages on a production system. We highly recommend that you issue this command during a maintenance window.

```
[local]SE800_10.192.16.82#debug ospf packet errors
[local]SE800_10.192.16.82#Dec 1 17:24:57: [0001]: %OSPF-7-PCK_ERRORS: OSPF-100 NBR 2.1.1.1:
MTU 1400 less than our mtu (1500)
Dec 1 17:25:02: [0001]: %OSPF-7-PCK_ERRORS: OSPF-100 NBR 2.1.1.1: MTU 1400 less than our mtu (1500)
Dec 1 17:25:26: [0001]: %OSPF-7-PCK_ERRORS: OSPF-100 NBR 2.1.1.1: MTU 1400 less than our mtu (1500)
Dec 1 17:25:31: [0001]: %OSPF-7-PCK_ERRORS: OSPF-100 NBR 2.1.1.1: MTU 1400 less than our mtu (1500)
Dec 1 17:25:36: [0001]: %OSPF-7-PCK_ERRORS: OSPF-100 NBR 2.1.1.1: MTU 1400 less than our mtu (1500)
Dec 1 17:25:41: [0001]: %OSPF-7-PCK_ERRORS: OSPF-100 NBR 2.1.1.1: MTU 1400 less than our mtu (1500)
Dec 1 17:25:46: [0001]: %OSPF-7-PCK_ERRORS: OSPF-100 NBR 2.1.1.1: MTU 1400 less than our mtu (1500)
Dec 1 17:25:51: [0001]: %OSPF-7-PCK_ERRORS: OSPF-100 NBR 2.1.1.1: MTU 1400 less than our mtu (1500)
Dec 1 17:25:56: [0001]: %OSPF-7-PCK_ERRORS: OSPF-100 NBR 2.1.1.1: MTU 1400 less than our mtu (1500)
Dec 1 17:26:01: [0001]: %OSPF-7-PCK_ERRORS: OSPF-100 NBR 2.1.1.1: MTU 1400 less than our mtu (1500)
Dec 1 17:26:06: [0001]: %OSPF-7-PCK_ERRORS: OSPF-100 NBR 2.1.1.1: MTU 1400 less than our mtu (1500)
Dec 1 17:26:11: [0001]: %OSPF-7-PCK_ERRORS: OSPF-100 NBR 2.1.1.1: MTU 1400 less than our mtu (1500)
Dec 1 17:26:36: [0001]: %OSPF-7-PCK_ERRORS: OSPF-100 NBR 2.1.1.1: MTU 1400 less than our mtu (1500

)
```

```
[local]SE800_10.192.16.82#show ip interface 1/3

Intf name:      1/3
Intf state:     Up                      MTU:            1500
IP address:     2.1.1.2                 Prefix len:     30
OSPF instance:  100                     OSPF net type:  broadcast
OSPF cost:      1                       OSPF state:     BDR
Resoln type:    Arp                     ARP timeout:    3600
ARP proxy:      Disabled                ARP secured:    Disabled
Number of Bound Circuits (incl. dynamic) = 1
Bindings: (Total Bound Circuits 1)
Encapsulation  Circuit
ethernet       1/3
[local]SE800_10.192.16.82#

[local]SE1200_10.192.17.224#show ip interface 2/1

Intf name:      2/1
Intf state:     Up                      MTU:            1400
IP address:     2.1.1.1                 Prefix len:     30
OSPF instance:  100                     OSPF net type:  broadcast
OSPF cost:      1                       OSPF state:     DR
Resoln type:    Arp                     ARP timeout:    3600
ARP proxy:      Disabled                ARP secured:    Disabled
Number of Bound Circuits (incl. dynamic) = 1
Bindings: (Total Bound Circuits 1)
Encapsulation  Circuit
ethernet       2/1
[local]SE1200_10.192.17.224#

[local]SE1200_10.192.17.224(config)#context local
[local]SE1200_10.192.17.224(config-ctx)#interface 2/1
[local]SE1200_10.192.17.224(config-if)#ip mtu 1500
[local]SE1200_10.192.17.224(config-if)#commit
Transaction committed.
[local]SE1200_10.192.17.224(config-if)#end
[local]SE1200_10.192.17.224#
```

2.1.4 Loading State

In the loading state, routers send link-state request packets. During the database exchange, if a router receives an outdated or missing LSA, it requests that LSA by sending a link-state request packet. Neighbors that do not transition beyond this state are most likely exchanging corrupted LSAs.

The most common causes of this problem include:

- Mismatched MTU

  Recommended Action:
  1. Run the **`show ip interface`** command on the interfaces that have the mismatched MTUs.
  2. Change either router MTU to match the neighbor MTU.

- Corrupted link-state request packet

  Recommended Action: Run the **`clear ospf neighbor all`** command to force adjacency reestablishment

2.1.5 Full State

In a *Full* state, routers are fully adjacent with each other. All the LSAs within the area are exchanged, and the router databases are fully synchronized. Full is the normal state for an OSPF router. Being stuck in another state indicates that the router cannot form adjacencies. The 2-way state is normal on broadcast and NBMA networks. In these networks, routers achieve the full state with their DR and BDR only. Other neighbors always see each other as 2-way.

```
[local]SE800_10.192.16.82#show ospf neighbor detail

  --- OSPF Neighbor 24.24.24.24 Area 0.0.0.0 Instance 100 ---

Address         : 2.1.1.1           Interface Addr : 2.1.1.2
State           : Full              DR State       : DR
Cost            : 1                 DR Priority    : 1
DR Router ID    : 2.1.1.1           BDR Router ID  : 2.1.1.2
LSA Request     : 0                 LSA Retrans    : 0
DB Exchange     : 0                 Time Till Dead : 32
Hello Options   : E,O               DD Options     : E,O
[local]SE800_10.192.16.82
```

2.1.6 show ospf neighbor Output Is Empty

When the **`show ospf neighbor detail`** command displays no output, run the **`show ip interface`** and **`show ip interface detail`** commands to verify the state of the interfaces.

Table 9 lists the common causes for this problem:

Table 9 OSPF Neighbor Checklist

#	Possible Reason	Checked?
1	OSPF is not enabled on the interface. Layer 1 or 2 is down.	
2	The interface is defined as passive under OSPF.	
3	A subnet number or mask has been mismatched over a broadcast links.	
4	The Hello or dead interval has been mismatched.	
5	The authentication type (plain text versus MD5) has been mismatched.	
6	An area ID has been mismatched.	
7	Stub, transit, or NSSA area options have been mismatched.	
8	An OSPF adjacency exists over an asynchronous interface.	
9	No network type or neighbor is defined over NBMA (Frame Relay).	

Run the **`show ip interface`** command to verify that the interface state is Up. Run the **`show ospf interface`** command to display detailed information about a specific OSPF interface.

```
[local]SE800_10.192.16.82#show ip interface 1/3

Intf name:       1/3
Intf state:      Up                        MTU:             1500
IP address:      2.1.1.2                   Prefix len:      30
OSPF instance:   100                       OSPF net type:   broadcast
OSPF cost:       1                         OSPF state:      BDR
Resoln type:     Arp                       ARP timeout:     3600
ARP proxy:       Disabled                  ARP secured:     Disabled
Number of Bound Circuits (incl. dynamic) = 1
Bindings: (Total Bound Circuits 1)
Encapsulation  Circuit
ethernet       1/3


[local]SE800_10.192.[local]SE800_10.192.16.82#show ospf interface 2.1.1.2 detail

  --- OSPF Interface 2.1.1.2 Area 0.0.0.0 Instance 100 ---

Network Type    : broadcast          Mask           : 255.255.255.252
Cost            : 1                  Logical Intf   : 1/3
MTU             : 1500               Physical Intf  : ethernet 1/3
State           : BDR                Priority       : 1
Hello Interval  : 10                 Dead Interval  : 40
Transmit Delay  : 1                  Retransmit Int : 5
DR Router ID    : 24.24.24.24        DR IP Address  : 2.1.1.1
BDR Router ID   : 2.1.1.2            BDR IP Address : 2.1.1.2
Ack Queued      : 0                  Flood Queued   : 0
Ack Delay       : 2                  Authentication : None
LSA Count       : 0                  LSA Checksum   : 0
Demand Circuit  : No                 Flood Reduction: No
Neighbor Count  : 1

Neighbor List (1 Adjacent):
24.24.24.24
[local]SE800_10.192.16.82#
16.82#
```

2.2 Troubleshooting OSPF Routing Tables

Most OSPF problems are caused by reachability problems. When no adjacency issues exist, check for route update problems. Route update problems are difficult to isolate because the routing table can be populated with routing information from multiple sources.

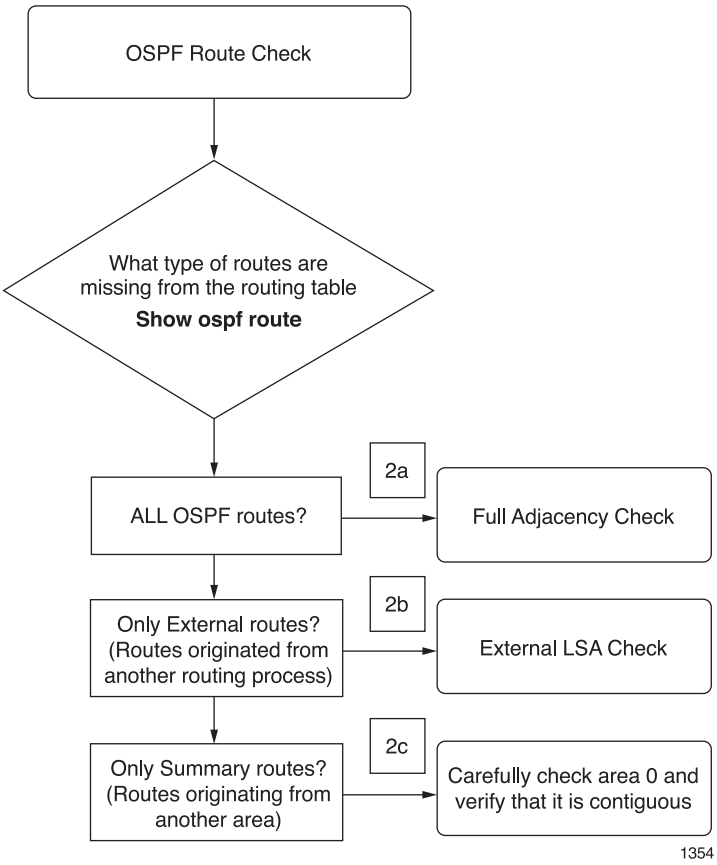

Figure 5 OSPF Routing Table Flowchart

The common causes of routing update problems for OSPF include missing OSPF routes and missing external routes.

If a route is not being learned in the rest of the area or domain, verify that the route in the OSPF process has been correctly configured.

2.2.1 Missing OSPF Routes

If all OSPF routes are missing, the OSPF neighbor likely is not in a *Full* state.

Do the following:

1. Check the physical link.
2. Ping the neighbor IP interface.
3. If you can ping the neighbor IP address, use the information in Troubleshoot OSPF Neighbor States to verify the state of neighbor adjacency.

Non-OSPF routes redistributed into OSPF are called external routes. In the following example, the SmartEdge SE100 and the SmartEdge SE800 routers are running IS-IS. The SE800 router redistributes its IS-IS routes into OSPF and, as a result, the SE1200 router can learn the route to reach the SE100 router through OSPF.

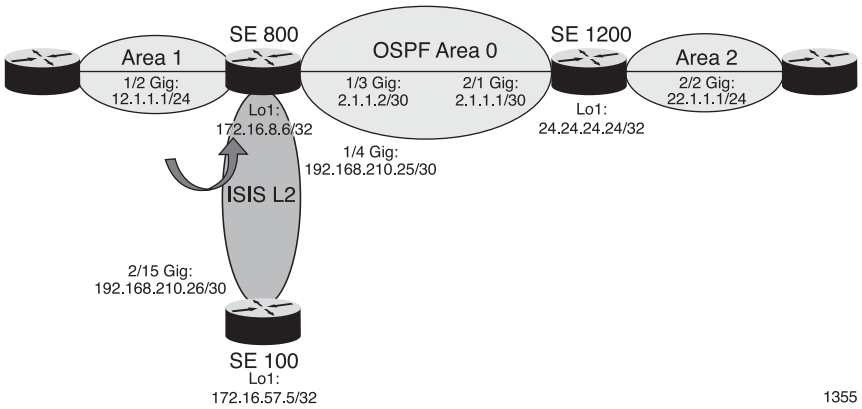

Figure 6 External Routes

When the external routes are missing, do the following:

1 Check if the External LSA Exists in the OSPF Database

2 Check if the IS-IS Routes from Originating Router

Make sure that originating router is **not** trying to redistribute into a stub area.

3 Check the Forwarding Address

2.2.1.1 Step 1: Check if the External LSA Exists in the OSPF Database

```
[local]SE800_10.192.16.82#show ospf database external detail

  --- OSPF Link State Database for Instance 100/Router ID 172.16.8.6 ---

        --- External LSA 172.16.57.5 ---

Link State Id     : 172.16.57.5        Advertising Router   : 172.16.8.6
Sequence Number   : 0x80000006         Checksum             : 0xba96
Options           : E,DC               Length               : 36
Mask              : 255.255.255.255    Metric               : 660
Forwarding Addr   : 0.0.0.0            Type                 : 2
Tag               : 0x0                Age                  : 1522


<-- 0.0.0.0 defaults to the originating router address (the advertising router).
```

2.2.1.2 Step 2: Check IS-IS routes from the Originating Router.

```
[local]se100_16_112#show isis database l2 detail

IS-IS level 2 link-state database for tag PBN1:
LSPID                 Sequence   Checksum   Holdtime   AT/OL   Len
SE800_10.192.1.00-00   0x323      0xa659     569        0/0     113
  Area Address: 49.0840
  NLPID:  IP
  Hostname: SE800_10.192.16.82
  Router ID: 172.16.8.6
  IP Address: 172.16.8.6
  M-Topology:
  Metric: 650         IS-Extended SE800_10.192.1.01
  Metric: 10          IP 172.16.8.6/32
  Metric: 650         IP 192.168.210.24/30
SE800_10.192.1.01-00   0x1d9      0x6fff     569        0/0     53
  Metric: 0           IS-Extended SE800_10.192.1.00
  Metric: 0           IS-Extended se100_16_112.00
se100_16_112.00-00*    0x188      0x8da      1178       0/0     107
  Area Address: 49.0840
  NLPID:  IP
  Hostname: se100_16_112
  Router ID: 172.16.57.5
  IP Address: 172.16.57.5
  M-Topology:
  Metric: 650         IS-Extended SE800_10.192.1.01
  Metric: 10          IP 172.16.57.5/32
  Metric: 650         IP 192.168.210.24/30

Total IS-IS LSP(s) for tag PBN1 in Level-2:   3
[local]se100_16_112#
```

2.2.1.3 Step 3: Check the Forwarding Address

The forwarding address for external route must be known as an internal OSPF route, and it must be an interarea or intra-area route.

1 If the AS-External LSA exists, run the **`show ospf database external detail`** command to check the forwarding address.

2 If the forwarding address field is `0.0.0.0`, run the **`show ospf border-routers asbr`** to determine whether the route to the advertising router exists.

3 If the forwarding address field is `a.b.c.d`, run the **`show ospf route a.b.c.d`** to determine whether a route to the specified forwarding address exists.

4 If a route does not exist, check the network state, and check the entire database.

2.2.2 Missing External Routes

Each OSPF network that is divided into different areas must follow these rules:

- A backbone area (area 0) that combines a set of independent areas into a single domain must exist.
- Each nonbackbone area must be directly connected to the backbone area; this connection can be a simple logical connection through a virtual link.
- If you have an issue, verify that the backbone area has **not** been partitioned. Given a failure condition, such as link or router down events, carefully check area 0 and verify that it is contiguous.

The most common causes of OSPF not installing external routes in the routing table include:

- The forwarding address is not known through the intra-area or inter-area route.
- The ABR is not generating Type 4 summary LSAs.

The following OSPF problems can happen during redistribution:

- ASBR does not advertise redistributed routes.
- OSPF does not install external routes in the routing table.

To troubleshoot missing summary routes:

1 Verify the OSPF Areas.

2 Identify and Check the ABR router.

3 Verify if the Inter-area Routes are Learned.

4 Verify if the Summary LSAs are Learned.

2.2.2.1 Step 1: Verify the Number of OSPF Areas

```
[local]SE800_10.192.16.82#show ospf area brief
  --- OSPF Areas for Instance 100/Router ID 172.16.8.6 ---
Area            Type             Intf-Count  LSA-Count   LSA cksum
0.0.0.0         Regular          2           5           0x00018e9d
0.0.0.1         Regular          1           5           0x0001f305
[local]SE800_10.192.16.82#
```

2.2.2.2 Step 2: Identify and Check the ABR Router

```
[[local]SE800_10.192.16.82#show ospf border-routers detail

  --- Border Routers for OSPF Instance 100/Router ID 172.16.8.6 ---

Destination    : 2.1.1.1              Type           : Intra
Dest-type      : ABR                  Cost           : 1
Area           : 0.0.0.0              Back Link Data : 2.1.1.1
LSDB Type      : Router               LSDB ID        : 2.1.1.1
LSDB Adv Router: 2.1.1.1              Next Hop Count : 1
SPF Version    : 8
Next Hop 1 Intf:                      Next Hop 1 Addr: 2.1.1.1
[local]SE800_10.192.16.82#

[local]SE800_10.192.16.82#show ospf interface

  --- OSPF Interfaces for Instance 100/Router ID 172.16.8.6 ---

Addr            Len  NetworkType     Cost    Priority State    Area
2.1.1.2         30   broadcast       1       1        BDR      0.0.0.0
172.16.8.6      32   loopback        1       1        Loopback 0.0.0.0
12.1.1.1        24   broadcast       1       1        DR       0.0.0.1
[local]SE800_10.192.16.82#


[local]SE800_10.192.16.82#show ospf route

       --- OSPF Routes for Instance 100/Router ID 172.16.8.6 ---

Destination        Type       Dest-Type/Proto   Cost  Nhops  Nhop
2.1.1.0/30         Intra      Net               1     1      1/3
12.1.1.0/24        Intra      Net               1     1      1/2
22.1.1.0/24        Inter      Sum-Net           2     1      2.1.1.1
24.24.24.24/32     Intra      Net               2     1      2.1.1.1
172.16.8.6/32      Intra      Net               1     1      lo1
```

2.2.2.3 Step 3: Verify that the Router Has Learned the Inter-area Routes

```
[local]SE800_10.192.16.82#show ospf route inter-area detail

       --- OSPF Routes for Instance 100/Router ID 172.16.8.6 ---

Destination     : 22.1.1.0/24           Type            : Inter
Dest-type       : Sum-Net               Cost            : 2
Version         : 7                     SPF Version     : 8
Area            : 0.0.0.0               Distance        : 110
Next Hop Count : 1
Next Hop 1 Intf:                        Next Hop 1 Addr: 2.1.1.1
[local]SE800_10.192.16.82#

[local]SE800_10.192.16.82#show ip  route ospf
Codes: C - connected, S - static, S dv - dvsr, R - RIP, e B - EBGP, i B - IBGP
       A,H - derived hidden
       O   - OSPF, O3  - OSPFv3, IA - OSPF(v3) inter-area,
       N1  - OSPF(v3) NSSA external type 1, N2  - OSPF(v3) NSSA external type 2
       E1  - OSPF(v3) external type 1, E2  - OSPF(v3) external type 2
       i   - IS-IS, L1 - IS-IS level-1,  L2  - IS-IS level-2, N - NAT
       IPH - IP Host, SUB A - Subscriber address, SUB S - Subscriber static
       MIP F - Mobile-IP Foreign Agent, MIP H - Mobile-IP Home Agent
       A - Derived Default, MH - Media Nexthop
       >   - Active Route, * - LSP

Type    Network             Next Hop        Dist  Metric    UpTime  Interface
  O     2.1.1.0/30          2.1.1.2          110       1            1/3
  O     12.1.1.0/24         12.1.1.1         110       1            1/2
> O IA  22.1.1.0/24         2.1.1.1          110       2  01:43:56  1/3
> O     24.24.24.24/32      2.1.1.1          110       2  02:14:14  1/3
  O     172.16.8.6/32       172.16.8.6       110       1            lo1

[local]SE800_10.192.16.82#show ospf route

       --- OSPF Routes for Instance 100/Router ID 172.16.8.6 ---

Destination         Type       Dest-Type/Proto   Cost  Nhops  Nhop
2.1.1.0/30          Intra      Net               1     1      1/3
12.1.1.0/24         Intra      Net               1     1      1/2
22.1.1.0/24         Inter      Sum-Net           2     1      2.1.1.1
24.24.24.24/32      Intra      Net               2     1      2.1.1.1
172.16.8.6/32       Intra      Net               1     1      lo1
[local]SE800_10.192.16.82#
```

2.2.2.4 Step 4: Verify that the Router Has Learned the Summary LSAs

```
[local]SE800_10.192.16.82#show ospf database summary-network all

  --- OSPF Link State Database for Instance 100/Router ID 172.16.8.6 ---

             Summary Network Link State Advertisements (Area 0.0.0.0)
LinkID          AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
12.1.1.0        172.16.8.6       80000004    ea75   E,DC     28     803
22.1.1.0        2.1.1.1          80000004    30eb   E,DC     28     1057

             Summary Network Link State Advertisements (Area 0.0.0.1)
LinkID          AdvertisingRtr   Sequence #  ChkSm  Option   Length LSAge
2.1.1.0         172.16.8.6       80000005    5913   E,DC     28     778
22.1.1.0        172.16.8.6       80000004    72e2   E,DC     28     888
24.24.24.24     172.16.8.6       80000005    52b9   E,DC     28     688
172.16.8.6      172.16.8.6       80000005    8220   E,DC     28     638
```

2.3 Troubleshooting OSPF Route Summarization Problems

OSPF can use two types of summarization:

- Interarea summarization that can be done on the ABR
- External summarization that can be done on the ASBR

Common problems related to summarization in OSPF include:

- A router does not summarize interarea routes.

  Cause: The **`area range`** command is not configured on the ABR.

- A router does not summarize external routes.

  Cause: The **`summary-address`** (OSPF) command is not configured on the ASBR.

2.4 Troubleshooting OSPF Not Advertising Routes

The most common reasons for OSPF to not share the database information about a specific link are:

- The neighbors's interface is configured to be passive, and so is not advertising routes.
- The OSPF neighbor (ABR) is not advertising the summary route.
- The OSPF neighbor is not advertising the default route.
- The OSPF neighbor is not advertising external routes.

2.5 Troubleshooting SPF Calculation and Route Flapping

The most common causes of SPF running constantly in the network include:

- An interface flap within the network.
- A neighbor flap within the network.
- A duplicate router ID.

2.6 OSPF Neighbor Not Advertising Default Routes

The most common causes for an OSPF router to not advertise the default route include:

- The **`originate-default`** command is missing.
- The default route is missing from the neighbor's routing table.
- A neighbor is trying to originate a default into a stub area.
- The **`default-route`** command at the router OSPF area level is not enabled.