

OSPF Neighbor Problems Explained

Document ID: 13699

Introduction

Prerequisites

Requirements

Components Used

Conventions

Adjacencies

Neighbor States

No State Revealed

Neighbor in down State

Neighbor in init State

Neighbor in 2-way State

Neighbor in exstart or exchange State

Neighbor in loading State

Typical Reasons for OSPF Neighbor Problems

Related Information

Introduction

This document explains common problems with getting Open Shortest Path First (OSPF) neighbors to become fully adjacent. In a successful formation of OSPF adjacency, OSPF neighbors will attain the FULL neighbor state.

In order to examine the formation of adjacency, issue the **debug ip ospf adj** command. Before you issue **debug** commands, refer to Important Information on Debug Commands.

Prerequisites

Requirements

This document requires a basic understanding of IP routing protocols and the OSPF routing protocol. In order to learn more about IP routing protocols, refer to Routing Basics. Refer to the Open Shortest Path First (OSPF) Support Page for more information about OSPF.

Components Used

The information in this document is based on these software and hardware versions:

- OSPF can be configured in all routers, such as the Cisco 2500 series, the Cisco 2600 series, and Layer 3 switches.
- OSPF is supported in Cisco IOS® Software Releases 10.0 and later.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Conventions

Refer to Cisco Technical Tips Conventions for more information on document conventions.

Adjacencies

The fact that routers are neighbors is not sufficient to guarantee an exchange of link–state updates; they must form adjacencies to exchange link–state updates. Adjacency is an advanced form of neighborhood formed by routers that are willing to exchange routing information after negotiating parameters of such an exchange. Routers reach a FULL state of adjacency when they have synchronized views on a link–state database.

Interface type plays a major role in how the adjacencies are formed. For example, neighbors on point–to–point links always try to become adjacent, while routers attached to broadcast media such as Ethernet can choose to become adjacent only with a subset of neighboring routers on the interface.

Once a router decides to form an adjacency with a neighbor, it starts by exchanging a full copy of its link–state database. The neighbor, in turn, exchanges a full copy of its link–state database with the router. After passing through several neighbor states, the routers become fully adjacent.

Neighbor States

You can use the **show ip ospf neighbor** command to determine the state of the OSPF neighbor or neighbors. The output of this command will most likely reveal one of these:

- nothing at all
- state = down
- state = init
- state = exstart
- state = exchange
- state = 2-way
- state = loading

There are other OSPF states, but those shown here are the most common ones seen in **show ip ospf neighbor** command output. Refer to OSPF Neighbor States for more information and an explanation of all the OSPF neighbor states.

No State Revealed

If the **show ip ospf neighbor** command reveals nothing at all or reveals nothing about the particular neighbor you are analyzing then this router has not seen any valid OSPF HELLOs from that neighbor. This means that OSPF either did not receive any HELLO packets from the neighbor or received HELLO packets that failed very basic sanity checks.

Check these:

- Is the interface up on the local router and neighboring router, with line protocol up? Issue the **show interface** command to check.
- Check for IP connectivity between the neighboring routers, as shown here:
 - ◆ Does the neighbor respond to a **ping** command? Ping the IP address assigned to the interface in question on the neighboring router. Issue a **traceroute** command to the same IP address and check that it takes no more than one hop to reach the destination.

- ◆ Does the neighbor respond if you issue a **ping 224.0.0.5** command? (224.0.0.5 is the address to which OSPF HELLOs are sent.)
- ◆ Check for any inbound access lists or other devices (such as a switch) that might prohibit the sending of IP packets from one neighbor to the other.
- Is OSPF enabled on both your interface and the interface of the neighboring router? Issue the **show ip ospf interface** command to check.
- Is OSPF configured as passive for the interface of the local or neighboring router? Issue the **show ip ospf interface** command to verify that HELLO packets are due to be sent out of the interface. An active OSPF interface displays a line similar to this:

```
Hello due in 00:00:07
```

- Verify that neighboring routers have different Router IDs. Router IDs are used to identify each router in an OSPF network. Routers with the same Router ID will ignore HELLOs sent by each other, which prevents them from forming adjacency. The first line of **show ip ospf** command output displays the current Router ID of each router.
- Verify that these HELLO parameters match on the neighboring interfaces:
 - ◆ OSPF area number Issue the **show ip ospf interface interface-name** command to check.
 - ◆ OSPF area type, such as `stub` or `NSSA` Issue the **show ip ospf** command to check.
 - ◆ Subnet and subnet mask Issue the **show interface** command to check.
 - ◆ OSPF HELLO and Dead timer values Issue the **show ip ospf interface interface-name** command to check.
- If the problem is on point-to-point link (such as PPP or High-Level Data Link Control [HDLC]) and there is more than one parallel link between this pair of routers, verify that lines are properly connected. Suppose you planned to connect interface Serial0/0 on one router with interface Serial0/0 on its neighbor and Serial1/0 with Serial1/0 on its neighbor, but you accidentally crossed them and connected Serial0/0 of each router with Serial1/0 on the other. The **ping** command might not discover such a problem, but OSPF will fail to establish adjacency. Use information provided by Cisco Discovery Protocol (CDP) to verify proper device interconnection. Issue the **show cdp neighbor interface-name** command to verify that the name and PortID of a remote device match network design.

Note: OSPF adjacencies only form over primary networks, not secondary networks.

If all of these checks are verified and the **show ip ospf neighbor** command still reveals nothing, then your problem is not very common and you might need to Contact Cisco for assistance.

Neighbor in down State

A neighbor that is discovered dynamically through reception of HELLO packets can fall back to a down state if it is being deleted, for example when OSPF does not receive HELLO packets from the neighbor for period of time longer than the Dead timer interval. Therefore, the down state is transient for such neighbors; they will either advance to higher states or be completely deleted from the table of known neighbors. This is known as being forgotten .

Usually, neighbors that are seen in the down state were manually configured with the **neighbor** command. Manually configured neighbors are always present in the OSPF neighbor table. If OSPF has never received HELLO packet from the manually configured neighbor, or if no HELLO packets were heard from the neighbor during the previous Dead timer interval, then the manually configured neighbor will be listed as down.

Note: The **neighbor** command can only be configured for directly attached neighbors on these types of networks:

- Non-Broadcast MultiAccess (NBMA) networks Interfaces configured with the **ip ospf network non-broadcast** command.
- Non-Broadcast Point-to-Multipoint networks Interfaces configured with the **ip ospf network point-to-multipoint non-broadcast** command.

If you see a neighbor in the down state, verify that the neighbor router is up, is running, and is properly configured for OSPF on this interface. Test connectivity between routers with the **ping** and **traceroute** commands. Check the OSPF neighbor table on the neighboring router with the **show ip ospf neighbor** command, and perform the same configuration verification actions listed in the No State Revealed section.

Neighbor in init State

The `init` state indicates that a router sees HELLO packets from the neighbor, but two-way communication has not been established. A Cisco router includes the Router IDs of all neighbors in the `init` (or higher) state in the Neighbor field of its HELLO packets. For two-way communication to be established with a neighbor, a router also must see its own Router ID in the Neighbor field of the neighbor's HELLO packets. For a more detailed example and explanation, refer to [Why Does the show ip ospf neighbor Command Reveal Neighbors in the Init State?](#)

Neighbor in 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 Descriptor (DBD) packet from a neighbor in the `init` state will also cause a transition to 2-way state. The OSPF neighbor 2-way state is not a cause for concern. For an explanation of the 2-way state, refer to [Why Does the show ip ospf neighbor Command Reveal Neighbors Stuck in 2-Way State?](#)

Neighbor in exstart or exchange State

OSPF neighbors that are in `exstart` or `exchange` state are trying to exchange DBD packets. The router and its neighbor form a master and slave relationship. The adjacency should continue past this state. If it does not, there is a problem with the DBD exchange, such as a maximum transmission unit (MTU) mismatch or the receipt of an unexpected DBD sequence number. For more information, refer to [Why Are OSPF Neighbors Stuck in Exstart/Exchange State?](#)

Neighbor in loading State

In the `loading` state, routers send link-state request packets. During the adjacency, if a router receives an outdated or missing link-state advertisement (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. This problem is usually accompanied by a `%OSPF-4-BADLSA` console message. Because this problem is not common, Contact Cisco for assistance.

Typical Reasons for OSPF Neighbor Problems

This table lists reasons why OSPF neighbors have problems forming an adjacency and lists some of the commands you can use to verify the problem. TAC Case Collection (registered customers only) can also help diagnose problems.

Reason for Neighbor Adjacency Problem	Commands for Diagnosing the Problem
---------------------------------------	-------------------------------------

OSPF is not configured on one of the routers.	show ip ospf
OSPF is not enabled on an interface where it is needed.	show ip ospf interface
OSPF HELLO or Dead timer interval values are mismatched.	show ip ospf interface
ip ospf network-type mismatch on the adjoining interfaces.	show ip ospf interface
MTU mismatch between neighboring interfaces.	show interface
OSPF area-type is <code>stub</code> on one neighbor, but the adjoining neighbor in the same area is not configured for <code>stub</code> .	<int-type><int-num> show running-config
OSPF neighbors have duplicate Router IDs.	show ip ospf interface show ip ospf show ip ospf interface
OSPF is configured on the secondary network of the neighbor, but not on the primary network. This is an illegal configuration which prevents OSPF from being enabled on the interface.	show ip ospf interface
OSPF HELLOs are not processed due to a lack of resources, such as high CPU utilization or not enough memory.	show running-config show memory summary
An underlying Layer problem is preventing OSPF HELLOs from being received.	show memory processor show interface

Note: In order to avoid the MTU check while you establish ospf adjacency, you possibly need to configure the **ip ospf mtu-ignore** command in interface configuration mode.

Related Information

- [Why Doesn't OSPF Form Adjacency on a PRI, BRI or Dialer Interface?](#)
 - [What Does the show ip ospf neighbor Command Reveal?](#)
 - [OSPF Design Guide](#)
 - [Initial Configurations for OSPF over Frame Relay Subinterfaces](#)
 - [Initial Configurations for OSPF over Non-Broadcast Links](#)
 - [Initial Configurations for OSPF over a Point-to-Point Link](#)
 - [Initial Configurations for OSPF Over Broadcast Media](#)
 - [OSPF: Frequently Asked Questions](#)
 - [Open Shortest Path First \(OSPF\) Support Page](#)
 - [Technical Support & Documentation – Cisco Systems](#)
-

