Troubleshoot OSPF Neighbor Problems

Contents

Introduction

Prerequisites

- Requirements
- Components Used
- Conventions

Background Information

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 describes common problems with how Open Shortest Path First (OSPF) neighbors become fully adjacent.

Prerequisites

Requirements

This document requires a basic knowledge of IP routing protocols and also the OSPF routing protocol. To learn more about IP routing protocols, refer to How to Configure Basic IP Routing. And 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 the software and hardware versions listed:

- OSPF can be configured in all routers, such as the Cisco 2900 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, ensure that you understand the potential impact of any command.

Conventions

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

Background Information

In a successful formation of OSPF adjacency, OSPF neighbors attain the `FULL` neighbor state. To examine the formation of adjacency, enter the debug ip ospf adj command. Before you enter **debug** commands, refer to Important Information on Debug Commands.

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 neighborship formed by routers that exchange routing information after parameters for such an exchange are negotiated. Routers reach a `FULL` state of adjacency when they have synchronized views on a link-state database.

The 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 become adjacent only with a subset of neighbor routers on the interface.

Once a router forms an adjacency with a neighbor, it begins with an exchange of 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 several neighbor states are passed, the routers become fully adjacent.

Neighbor States

Use the **show ip ospf neighbor** command to determine the state of the OSPF neighbor or neighbors. The output of this command reveals 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 of interest, then this router has not received 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? Enter the show interface command to check the interface status.
- Check for IP connectivity between the neighbor routers, as shown here:
    - Does the neighbor respond to a ping command? Ping the IP address assigned to the interface in question on the neighbor router. Enter the traceroute command to the same IP address and verify it takes no more than one hop to reach the destination.
    - Does the neighbor respond if you enter a ping 224.0.0.5 command? (224.0.0.5 is the IP address to which OSPF HELLOs are sent.)
    - Check for any inbound access lists or other devices (such as a switch) that can prohibit the IP packets to be forwarded from one neighbor to the other.
- Is OSPF enabled on both your interface and the interface of the neighbor/adjacent router? Enter the **show ip ospf interface** command to verify.
- Is OSPF configured as passive for the interface of the local or neighbor/adjacent router? Enter the **show ip ospf interface** command to verify HELLO packets are due to be sent out of the interface. An active OSPF interface displays a line similar to:

<#root>

```
Router#show ip ospf interface
GigabitEthernet0/0 is up, line protocol is up
Internet Address 10.1.1.1/30, Area 0, Attached via Network Statement
Process ID 1, Router ID 10.1.1.1, Network Type BROADCAST, Cost: 1
Topology-MTID Cost Disabled Shutdown Topology Name
0 1 no no Base
Transmit Delay is 1 sec, State BDR, Priority 1
Designated Router (ID) 10.1.1.2, Interface address 10.1.1.2
Backup Designated router (ID) 10.1.1.1, Interface address 10.1.1.1
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
oob-resync timeout 40
```

`Hello due in 00:00:05`

```
Supports Link-local Signaling (LLS)
Cisco NSF helper support enabled
IETF NSF helper support enabled
Index 1/1/1, flood queue length 0
Next 0x0(0)/0x0(0)/0x0(0)
Last flood scan length is 1, maximum is 1
Last flood scan time is 1 msec, maximum is 1 msec
Neighbor Count is 1, Adjacent neighbor count is 1
Adjacent with neighbor 10.1.1.2 (Designated Router)
Suppress hello for 0 neighbor(s)
```

- Verify that neighbor routers have different Router IDs. Router IDs are used to identify each router in

an OSPF network. Routers with the same Router ignores HELLOs sent by each other, and not become adjacent. The first line of the **show ip ospf** command output displays the current Router ID of each router.

- Verify these HELLO parameters match on the neighbor interfaces:
  - OSPF area number - Enter the **show ip ospf interface interface-name** command to verify.
  - OSPF area type, such as `stub` or `NSSA` – Enter the **show ip ospf** command to verify.
  - Subnet and subnet mask - Enter the **show interface** command to verify.
  - OSPF HELLO and Dead timer values - Enter the **show ip ospf interface interface-name** command to verify.
- If the problem is on a 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 the 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 cannot discover such a problem, but OSPF fails to establish adjacency. Use information provided by Cisco Discovery Protocol (CDP) to verify proper device interconnection. Enter the **show cdp neighbor interface-name** command to verify the name and PortID of a remote device match the 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 can contact Cisco TAC 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 deleted by OSPF process. For example, when OSPF does not receive HELLO packets from the neighbor for a period of time longer than the Dead timer interval, the `down` state is transient for such neighbors; they either advance to higher states or become deleted from the table of known neighbors. This is known as "forgotten".

Usually, neighbors 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 never received HELLO packets from the manually configured neighbor, or if no HELLO packets were received from the neighbor during the previous Dead timer interval, then the manually configured neighbor is listed as `down`.

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

- 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 a neighbor is in the `down` state, verify that the neighbor router is up, is active, 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 neighbor router with the show ip ospf neighbor command, and perform the same configuration verification actions listed earlier in this document in the No State Revealed section.

Neighbor in Init State

The `init` state indicates that a router receives 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 receive its own Router ID in the Neighbor field of the neighbor 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 received its own Router ID in the Neighbor field of the neighbor HELLO packet. Receipt of a Database Descriptor (DBD) packet from a neighbor in the `init` state also causes a transition to the `2-way` state. The OSPF neighbor `2-way` state is not a cause for concern in Broadcast and Non-Broadcast MultiAccess (NBMA) networks. 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 in the `exstart` or `exchange` state attempt to exchange DBD packets. The router and its neighbor form a primary and secondary relationship. The adjacency must 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. While adjacent, if a router receives an outdated or missed link-state advertisement (LSA), it sends a link-state request packet to request the LSA. Neighbors that do not transition beyond this state likely exchange corrupted LSAs. This problem is usually accompanied by a `%OSPF-4-BADLSA` console message. Because this problem is not common, contact Cisco TAC for assistance.

Typical Reasons for OSPF Neighbor Problems

This table lists reasons why OSPF neighbors have problems when they try to form an adjacency and lists some commands you can use to verify the problem.

Reason for Neighbor Adjacency Problem	Command to Diagnose 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 adjacent interfaces.	**show ip ospf interface**
MTU mismatch between neighbor interfaces.	**show interface <int-type><int-num>**
OSPF area-type is `stub` on one neighbor, but the adjacent neighbor in the same	**show running-config**

area is not configured for `stub`.	**show ip ospf interface**
OSPF neighbors have duplicate Router IDs.	**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 that prevents OSPF enablement on the interface.	**show ip ospf interface show running-config**
OSPF HELLOs are not processed due to a lack of resources, such as high CPU utilization or not enough memory.	**show memory summary show memory processor**
A Layer problem prevents receipt of OSPF HELLOs.	**show interface**

Note: To avoid the MTU check at OSPF adjacency establishment, you can configure the ip ospf mtu-ignore command in interface configuration mode. However, it is recommended to fix any MTU mismatch by reviewing the interface configuration instead of just bypassing the MTU check.

Related Information

- **Why Does not OSPF Form Adjacency on a PRI, BRI or Dialer Interface?**
- **What Does the show ip ospf neighbor Command Reveal?**
- **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**
- **Cisco Technical Support & Downloads**