


## OSPF Neighbor Adjacency

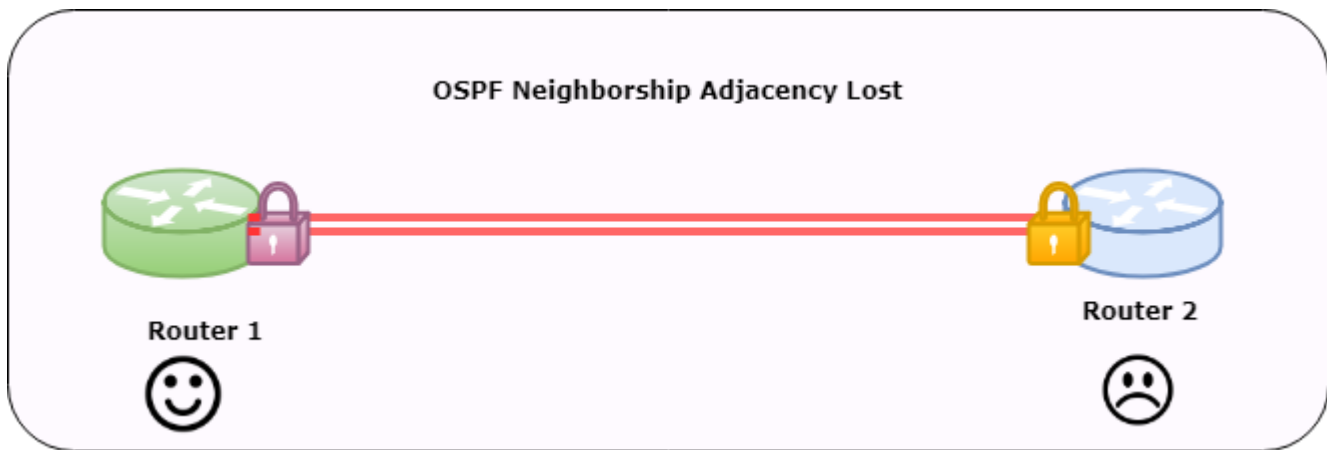
Once upon a time, we walked together holding hands, we enjoyed the breeze together, we called ourselves neighbors but could we say the same of our OSPF Neighbor Adjacency?



Troubleshooting OSPF Neighbor Adjacency is an important skill to have as a network engineer. OSPF is the most popular routing protocol loved by many.



Just like our behaviour in the real world, the fact that routers are neighbors is not sufficient enough to guarantee exchange of link-state updates; they must perform adjacencies which is based on a set of laid down requirements in order to exchange link-state updates. Adjacency is an advanced form of neighborhood formed by routers that are willing to participate in the exchange of routing information after negotiating parameters of such an exchange. Routers therefore can only 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 says hi, I have some more updates and so doing, exchanges a full copy of its link-state database with the router. After passing through several neighbor states, the routers become fully adjacent.

### **OSPF Adjacency Requirements,**

From the perspective of OSPF, there are a couple of things that must match for a OSPF neighborhood to establish; these include:

1. The devices must be in the same area
2. The devices must have the same authentication configuration
3. The devices must be on the same subnet
4. The devices hello and dead intervals must match
5. The devices must have matching stub flags

### **OSPF Adjacency Formation States**

During the adjacency formation process, two OSPF routers transition through several states, which include:

**Down.** Down is the starting state for all OSPF routers. A start event, such as configuring the protocol, transitions the router to the Init state. The local router may list a neighbor in this state when no hello packets have been received within the specified router dead interval for that interface.

**Init.** The Init state is reached when an OSPF router receives a hello packet but the local router ID is not listed in the received Neighbor field. This means that bidirectional communication has not been established between the peers.

**Attempt.** The Attempt state is valid only for Non-Broadcast Multi-Access (NBMA) networks. It means that a hello packet has not been received from the neighbor and the local router is going to send a unicast hello packet to that neighbor within the specified hello interval period.

**2-Way.** The 2-Way state indicates that the local router has received a hello packet with its own router ID in the Neighbor field. Thus, bidirectional communication has been established and the peers are now OSPF neighbors. (And this is where you stand and clap for the ingenuity of the neighborhood concept. Note that till the point of Neighborhood formation, databases haven't been exchanged. It is only before they become adjacent that they exchange databases which is explained below.)

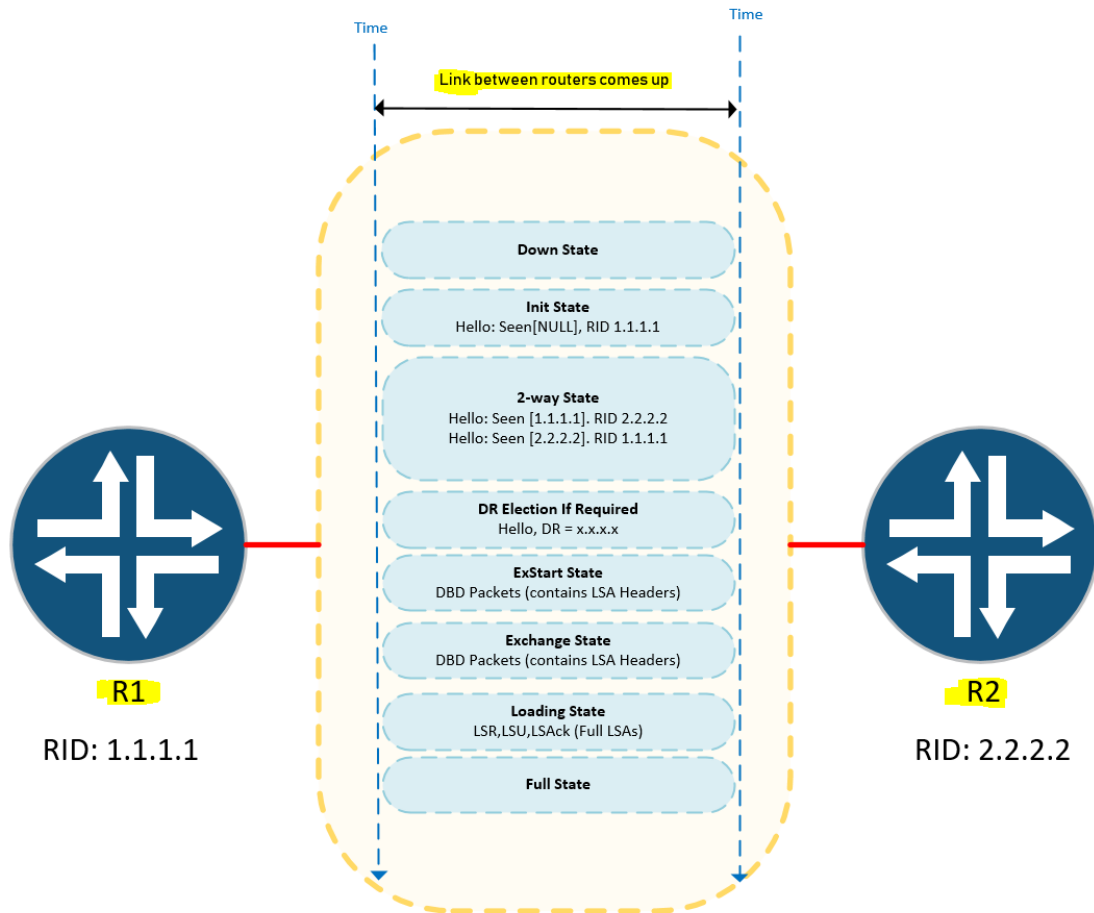
**ExStart.** In the ExStart state, the local router and its neighbor establish which router is in charge of the database synchronization process. The higher router ID of the two neighbors controls which router becomes the master.

**Exchange.** In the Exchange state, the local router and its neighbor exchange DD packets that describe their local databases.

Loading. Should the local router require complete LSA information from its neighbor, it transitions to the Loading state and begins to send link-state request packets.

**Full.** The Full state represents a fully functional OSPF adjacency, with the local router having received a complete link-state database from its peer. Both neighboring routers in this state add the adjacency to their local database and advertise the relationship in a link-state update packet.

Figure 1.0 – OSPF Neighbor Adjacency States and OSPF Neighbor Forming Process



Run the **show ip ospf neighbor** command to initiate the troubleshooting process for the lost OSPF neighbor adjacency,

```
Router1# show ip ospf neighbor
```

OSPF routers go through the seven states while building neighborship with other routers.

Down state  
Attempt/Init state  
Two ways state  
Exstart state  
Exchange state  
Loading state  
Full state

You can use the **show ip ospf neighbor** command in order 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

### **OSPF Neighborship Requirement;**

In order to become an OSPF neighbor, the following values must be match on both routers.

- Area ID
- Authentication
- Hello and Dead Intervals

Stub Flag  
MTU Size

**OSPF Stuck in INIT,**

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. Example 3-15 shows sample output of the show ip ospf neighbor command.

#### **OSPF Neighbor Stuck in INIT—Cause: Access List on One Side Is Blocking OSPF Hellos**

OSPF uses a multicast address of 224.0.0.5 for sending and receiving Hello packets. If an access list is defined on the interface and OSPF is enabled on that interface, this multicast address must be explicitly permitted in the access list; otherwise, it can produce problems such as stuck in INIT. The stuck in INIT problem occurs only if one side is blocking OSPF Hellos. If both sides are blocking OSPF Hellos, the output of show ip ospf neighbor returns an empty list.

#### **OSPF Neighbor Stuck in INIT—Cause: Multicast Capabilities Are Broken on One Side**

This is a specific situation that is valid only in the case of a Catalyst 6500 switch with the multilayer switch feature card (MSFC). The problem is that one side is sending OSPF Hellos that the other side does not receive. This situation is produced when the command set protocolfilter enabled is entered on the 6500 switch. By default, the protocol filter is disabled. Enabling this command begins altering the multicast frame to and from MSFC and port adapter within the FlexWan module of the 6500 switch. To fix this problem, disable the protocol filter on 6500 switch set protocolfilter disable.

#### **OSPF Neighbor Stuck in INIT—Cause: Cause: Authentication Is Enabled Only on One Side**

When authentication is used, it must be enabled on both sides; otherwise, one side will show the neighbor stuck in the INIT state. The router that has authentication enabled will reject all the nonauthenticated packets, and the adjacency will show stuck in INIT. The other side will not detect any problem because the authentication is turned on, so it will simply ignore the authentication in a packet and treat it as a normal packet.

#### **OSPF Neighbor Stuck in INIT—Cause: The Frame-Relay Map/dialer-map Statement on One Side Is Missing the broadcast Keyword**

OSPF uses a multicast address of 224.0.0.5 to send and receive OSPF Hellos. If one side is incapable of sending or receiving Hellos, the OSPF neighbor will be stuck in the INIT state. The important thing to note here is that only one side suffers from this multicast problem. R1 sees the neighbor in INIT state but can see the neighbor Hellos without any problem. When R1 sends the Hello to R2, it never reaches R2 because Layer 2 is incapable of sending any broadcast or multicast packets. This is because of the lack of the broadcast keyword in frame-relay map statement on R1. A similar problem can occur in the case of ISDN or dialer interface when the dialer map statement is configured without the broadcast keyword.

#### **OSPF Neighbor Stuck in INIT—Cause: Hellos Are Getting Lost on One Side at Layer 2**

This situation happens when there is a problem on the Layer 2 media; for example, the Frame Relay switch is blocking the multicast traffic for some reason. When R1 sends the Hello, R2 never receives it. Because R2 never saw Hellos from R1, the neighbor list of R2 will be empty. However, R1 sees the Hellos from R2, which does not list R1 as a valid neighbor; so, R1 declares this neighbor in the INIT state.

## **OSPF Stuck in LOADING State**

This is a rare problem in OSPF neighbor relationships. When a neighbor is stuck in the LOADING state, the local router has sent a link-state request packet to the neighbor requesting an outdated or missing LSA and is waiting for an update from its neighbor. If a neighbor doesn't reply or a neighbor's reply never reaches the local router, the router will be stuck in the LOADING state.

The most common possible causes of this problem are as follows:

- 1 Mismatched MTU
- 2 Corrupted link-state request packet

### **OSPF Neighbor Stuck in LOADING—Cause: Mismatched MTU Size**

This is a unique problem that happens when an MTU mismatch occurs. If the MTUs are not the same across the link, this problem occurs. Specifically, if a neighbor's MTU is greater than the local router's, the neighbor sends a large MTU packet as a link-state update. This packet never reaches the local router; as a result, the neighbor gets stuck in the LOADING state.

### **OSPF Neighbor Stuck in LOADING—Cause: Link-State Request Packet Is Corrupted**

When a link-state request packet is corrupted, the neighbor discards the packet and the local router never receives the response from the neighbor. This causes the OSPF neighbor to be stuck in the LOADING state.

Link-state request packets usually become corrupted because of the following reasons:

- 1 A device between the neighbors, such as a switch, is corrupting the packet.
- 2 The sending router's packet is invalid. In this case, either the sending router's interface is bad or the error is caused by a software bug.
- 3 The receiving router is calculating the wrong checksum. In this case, either the receiving router's interface is bad or the error is caused by a software bug. This is the least likely cause of this error message.

Most of the time, this problem is fixed by replacing hardware. This could be a simple bad port on the switch or a bad interface card on the sending/receiving router.

## **OSPF Neighbor Stuck in EXSTART/EXCHANGE**

This is an important state during the OSPF adjacency process. In this state, the router elects a master and a slave and the initial sequence number. The whole database also is exchanged during this state. If a neighbor is stuck in EXSTART/EXCHANGE for a long time, it is an indication of a problem.

The most common possible causes of this problem are as follows:

1. Mismatched interface MTU
2. Duplicate router IDs on neighbors
3. Inability to ping across with more than certain MTU size
4. Broken unicast connectivity because of the following:
5. Wrong VC/DLCI mapping in Frame Relay/ATM switch

6. Access list blocking the unicast
7. NAT translating the unicast
8. Network type of point-to-point between PRI and BRI/dialer

### **OSPF Neighbor Stuck in EXSTART/EXCHANGE—Cause: Mismatched Interface MTU**

OSPF sends the interface MTU in a database description packet. If there is a MTU mis-match, OSPF will not form an adjacency. The interface MTU option was added in RFC 2178. Previously, there was no mechanism to detect the interface MTU mismatch.

### **OSPF Neighbor Stuck in EXSTART/EXCHANGE—Cause: Duplicate Router IDs on Neighbors**

When OSPF sends a DBD packet to elect a master and a slave, the router with the highest router ID becomes the master. This happens in the EXSTART process. If there is any problem with election, the router will be stuck in the EXSTART/EXCHANGE state.

### **OSPF Neighbor Stuck in EXSTART/EXCHANGE—Cause: Can't Ping Across with More Than Certain MTU Size**

When OSPF begins forming an adjacency with its neighbor, it goes through several states. In EXSTART state, OSPF determines which will be the master and which will be the slave. After the routers decided this, they start exchanging the LSA header in the form of DBD packets. If the database is huge, OSPF uses the interface MTU and tries to send as much data as possible up to the limit of the interface MTU. If there is a problem with Layer 2 accepting large packets that are within the interface MTU range, the OSPF adjacency will be stuck in the EXCHANGE state.

### **OSPF Neighbor Stuck in EXSTART/EXCHANGE—Cause: Unicast Connectivity Is Broken**

When OSPF routers begin exchanging database information with each other, they send a unicast packet to each other in EXSTART/EXCHANGE state. This happens only if the network type is not a point-to-point link. In cases of a point-to-point link, OSPF sends all multicast packets. If unicast connectivity is broken, OSPF neighbor remains in EXSTART state.

This ping failure could occur for several reasons, including the following:

- 1 The wrong DLCI or VPI/VCI mapping exists in a Frame Relay or ATM switch, respectively.
- 2 An access list is blocking the unicast.
- 3 NAT is translating the unicast.

### **Wrong DLCI or VPI/VCI Mapping**

In cases of Frame Relay or ATM, this is a very common problem. The packet will be lost in the Frame Relay or ATM cloud. To further verify that this is the case, turn on debug ip packet detail with the access list on both routers.



### **Access List Blocking the Unicast**

If an access list is configured on a router, make sure that it's not blocking the unicast packet.

### **NAT Is Translating the Unicast**

This is another common problem that occurs when NAT is configured on the router. If NAT is misconfigured, it will start translating the unicast packet coming toward it, which will break the unicast connectivity.

### **OSPF Neighbor Stuck in EXSTART/EXCHANGE—Cause: Network Type Is Point-to-Point Between PRI and BRI/Dialer**

The network type on a PRI interface is point-to-point. This causes OSPF to send multicast packets even after the 2-WAY state. If only one BRI comes up as an OSPF neighbor, it will work fine. However, when multiple BRIs try to form an adjacency with the PRI, the PRI will complain because its network type is point-to-point. Because all OSPF packets are sent as multicast on a point-to-point link, the PRI receives DBD packets from multiple BRI neighbors, and this causes all the neighbors to get into the EXSTART/EXCHANGE state.

### **OSPF Troubleshooting Summary**

Neighborhood stuck in down stage : When an OSPF router first receives a hello packet, it verifies that the data in some fields matches its own locally configured information. Should any of the checked data be different, the hello packet is discarded and not processed. The data fields verified are the Area ID, Authentication, Network Mask (on broadcast networks), Hello Interval, Router Dead Interval, and Options fields. If this information doesn't match, then neighborhood is stuck in the down stage.

OSPF adjacency gets stuck in the ExStart state : This occurs due to a final check the routers perform. In the DD packet, each router advertises the IP MTU of the interface it is using. Should the local and remote routers not agree on the MTU of the network link, the database synchronization process stops and both neighbors remain in the ExStart state.

Visit [ExpertNetworkConsultant.com](http://ExpertNetworkConsultant.com) for expert articles on Networking.