



OSPF Timers – Hello and Dead Intervals



OSPF HelloInterval and RouterDeadInterval

OSPF timers are included in the OSPF Hello messages that are exchanged between OSPF-enabled routers. OSPF Hello messages are very important in that they enable OSPF devices to dynamically discover OSPF neighbors. OSPF Hello messages contain the information that must be agreed upon between the OSPF-enabled devices on the link in order for an OSPF neighbor relationship – and eventually an OSPF adjacency - to be formed. Without an adjacency, OSPF LSAs cannot be exchanged between devices and OSPF routing will fail.

One of the conditions included in the Hello messages that must be agreed upon is the OSPF timers. Specifically, the OSPF timers being agreed upon are the HelloInterval (Hello interval, Hello timer) and the RouterDeadInterval (router dead interval, dead interval, dead timer).

HelloInterval - The length of time, in seconds, between the Hello packets that the router sends on the interface. Advertised in Hello packets sent out this interface.

RouterDeadInterval - The number of seconds before the router's neighbors will declare it down, after they stop hearing the router's Hello Packets. Advertised in Hello packets sent out this interface.

***Hello Timer** - An interval timer that causes the interface to send a Hello packet. This timer fires every HelloInterval seconds.

***Wait Timer** - A single shot timer that causes the interface to exit the Waiting state, and as a consequence select a Designated Router on the network. The length of the timer is RouterDeadInterval seconds.



OSPF Neighbor Discovery

Before an OSPF router can send any OSPF messages, it must choose a unique 32-bit dotted-decimal identifier (note that does **not** necessarily need to be an IP address) called the OSPF router identifier (RID). Cisco routers use the following sequence to choose the OSPF RID:

1. Use the router ID manually configured under the OSPF process.
2. Use the highest numeric IP address on any currently “up and up” loopback interface*.
3. Use the highest numeric IP address on any currently “up and up” non-loopback interface*.

*These interfaces do not need to be OSPF-enabled nor advertised by OSPF.

We’ll take a much closer look at how Cisco devices choose an OSPF router ID and how that can affect OSPF routing in a future lesson, but this quick review suffices for our purpose.

Once an OSPF RID has been chosen, the router will start to send OSPF Hello Messages out all OSPF-enabled interfaces (excluding Loopback interfaces).



OSPF Neighbor Formation

When OSPF devices exchange OSPF Hello messages on a link they seek to form an OSPF neighbor relationship. In order to do so, they must agree on the following information:

- **IP subnet (including subnet mask)**
- **OSPF area**
- **Router IDs must be unique**
- **Stub flags (area type)**
- **OSPF timers**
- **Authentication**
- ***Compatible OSPF network types**

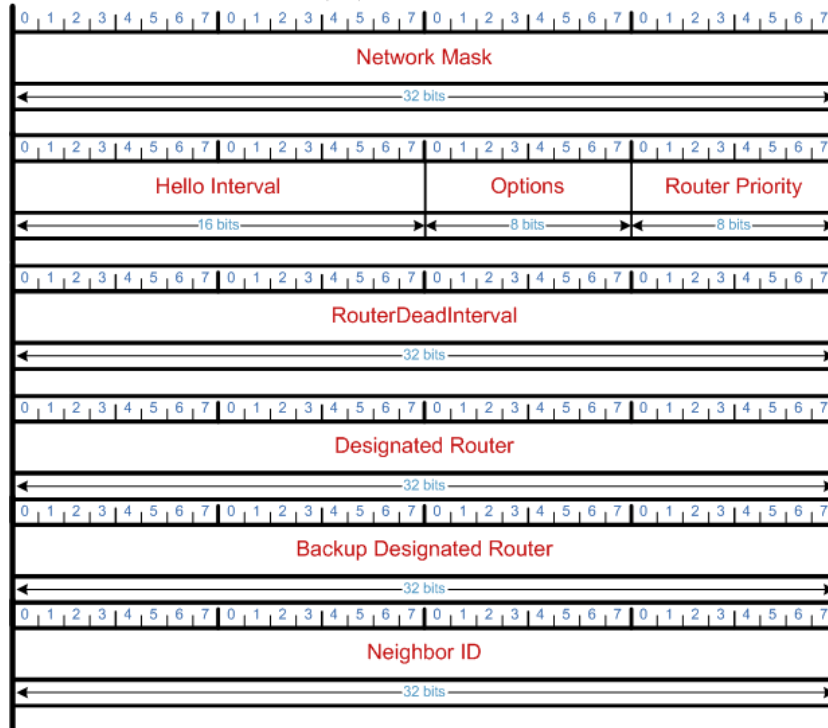
If one or more of these conditions are not agreed upon, the OSPF neighbor relationship fails.

*** Network type is not actually a requirement, but adjacencies will only form between compatible network types. Compatibility is based on the use of a DR/BDR.**



OSPF Hello Message

OSPF Hello





6	17.158000	17.158000	10.1.12.1	224.0.0.5	OSPF	Hello Packet
9	26.539000	26.539000	10.1.12.1	224.0.0.5	OSPF	Hello Packet
13	33.750000	33.750000	10.1.12.2	224.0.0.5	OSPF	Hello Packet

```
<
* Frame 6 (80 bytes on wire (80 bytes captured) on interface 0)
  * Cisco HDLC
  * Internet Protocol, Src: 10.1.12.1 (10.1.12.1), Dst: 224.0.0.5 (224.0.0.5)
  * Open Shortest Path First
    * OSPF Header
      * OSPF Version: 2
      * Message Type: Hello Packet (1)
      * Packet Length: 44
      * Source OSPF Router: 1.1.1.1 (1.1.1.1)
      * Area ID: 0.0.0.1
      * Packet checksum: 0xea9b [correct]
      * Auth Type: Null
      * Auth Data (none)
```

```
    * OSPF Hello Packet
      * Network Mask: 255.255.255.0
      * Hello interval: 10 seconds
      * options: 0x12 (L, E)
        * 0... .... = DN: DN-bit is NOT set
        * .0.. .... = O: O-bit is NOT set
        * ..0. .... = DC: Demand circuits are NOT supported
        * ...1 .... = L: The packet contains LLS data block
        * .... 0... = NP: Nssa is NOT supported
        * .... .0.. = MC: NOT multicast capable
        * .... ..1. = E: ExternalRoutingCapability
      * Router Priority: 1
      * Router Dead Interval: 40 seconds
      * Designated Router: 0.0.0.0
      * Backup Designated Router: 0.0.0.0
    * OSPF LLS Data Block
      * Checksum: 0xffff6
      * LLS Data Length: 12 bytes
    * Extended options TLV
      * Type: 1
      * Length: 4
      * options: 0x00000001 (LR)
        * ..... ..0. = RS: Restart signal (RS-bit) is NOT set
        * ..... ..1. = LR: LSDB Resynchronization (LR-bit) is SET
```



Default OSPF HelloInterval and RouterDeadInterval

OSPF Network Type	Default HelloInterval	Default RouterDeadInterval
Broadcast	10 seconds	40 seconds
Non-broadcast	30 seconds	120 seconds
Point-to-Point	10 seconds	40 seconds
Point-to-Multipoint	30 seconds	120 seconds
Point-to-Multipoint Non-broadcast	30 seconds	120 seconds
Loopback	N/A	N/A



OSFP Point-to-Point Network





Configuring OSPF HelloInterval

The default HelloInterval will be assigned based on OSPF network type and will be either 10 seconds (Ethernet, Token Ring, PPP, HDSL) or 30 seconds (Frame Relay and ATM main interfaces).

```
r1#show ip ospf interface s0/0
Serial0/0 is up, line protocol is up
  Internet Address 10.1.12.1/24, Area 1
  Process ID 1, Router ID 1.1.1.1, Network Type POINT_TO_POINT, Cost: 64
  Transmit Delay is 1 sec, State POINT_TO_POINT
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
<--Output truncated-->
```

```
r1#show ip ospf interface s0/0
Serial0/0 is up, line protocol is up
  Internet Address 10.1.12.1/24, Area 1
  Process ID 1, Router ID 1.1.1.1, Network Type NON_BROADCAST, Cost: 64
  Transmit Delay is 1 sec, State WAITING, Priority 1
  No designated router on this network
  No backup designated router on this network
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
<--Output truncated-->
```



Configuring OSPF HelloInterval

To change the OSPF HelloInterval (Hello Timer), you use the following interface level command:

ip ospf hello-interval

To specify the interval between hello packets that the Cisco IOS software sends on the interface, use the **ip ospf hello-interval** command in interface configuration mode. To return to the default time, use the **no** form of this command.

ip ospf hello-interval *seconds*

no ip ospf hello-interval

```
r1(config)#do show ip ospf interface s0/0 | i Timer
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
r1(config)#int s0/0
r1(config-if)#ip ospf hello-interval ?
  <1-65535>  Seconds

r1(config-if)#ip ospf hello-interval 6
r1(config)#do show ip ospf interface s0/0 | i Timer
  Timer intervals configured, Hello 6, Dead 24, Wait 24, Retransmit 5
```

Note that Cisco IOS automatically adjusts the RouterDeadInterval to 4x the value that you specify for the HelloInterval.



Configuring OSPF RouterDeadInterval

To change the OSPF RouterDeadInterval(Dead Timer), you use the following interface level command:

ip ospf dead-interval

To set the interval during which at least one hello packet must be received from a neighbor before the router declares that neighbor down, use the **ip ospf dead-interval** command in interface configuration mode. To restore the default value, use the **no** form of this command.

ip ospf dead-interval {*seconds* | **minimal hello-multiplier multiplier**}

no ip ospf dead-interval

Defaults

seconds: Four times the interval set by the **ip ospf hello-interval** command.

```
r1(config)#do show ip ospf interface s0/0 | i Timer
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
r1(config)#int s0/0
r1(config-if)#ip ospf dead-interval ?
<1-65535> Seconds
minimal Set to 1 second
```

```
r1(config-if)#ip ospf dead-interval 20
r1(config)#do show ip ospf interface s0/0 | i Timer
Timer intervals configured, Hello 10, Dead 20, Wait 20, Retransmit 5
```



OSPF Adjacencies

Remember that the HelloInterval and the RouterDeadInterval **must** be in agreement between OSPF neighbors in order for an OSPF neighbor relationship to form and OSPF LSAs to be exchanged. If either or both of these values are not matched, then the OSPF neighbor relationship – and by extension, the OSPF adjacency - fails:

```
r1#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
2.2.2.2	0	FULL/ -	00:00:38	10.1.12.2	Serial0/0

```
r1#debug ip ospf hello
```

```
r1(config)#int s0/0
```

```
r1(config-if)#ip ospf hello-interval 5
```

```
*Mar  3 20:00:23.839: OSPF: Send hello to 224.0.0.5 area 1 on Serial0/0 from 10.1.12.1  
*Mar  3 20:00:25.131: OSPF: Rcv hello from 2.2.2.2 area 1 from Serial0/0 10.1.12.2  
*Mar  3 20:00:25.135: OSPF: Mismatched hello parameters from 10.1.12.2  
*Mar  3 20:00:25.135: OSPF: Dead R 40 C 20, Hello R 10 C 5  
*Mar  3 20:00:45.975: %OSPF-5-ADJCHG: Process 1, Nbr 2.2.2.2 on Serial0/0 from FULL to DOWN,  
Neighbor Down: Dead timer expire
```

```
r1#show ip ospf neighbor
```

```
r1#
```



Why Change HelloInterval and/or RouterDeadInterval?

The HelloInterval and RouterDeadInterval are directly related to OSPF convergence time. An OSPF router will count down the RouterDeadInterval for each neighbor and will reset that counter each time that it receives a Hello message from that neighbor. If the router does not receive a Hello message from the neighbor before the RouterDeadInterval expires, that neighbor is set to Down state. By default, this means that if 4 consecutive Hello messages are missed then a neighbor is set to Down.

```
*Mar 3 20:40:58.154: OSPF: Send hello to 224.0.0.5 area 1 on Serial0/0 from 10.1.12.1
*Mar 3 20:41:07.642: OSPF: Send hello to 224.0.0.5 area 1 on Serial0/0 from 10.1.12.1
*Mar 3 20:41:17.286: OSPF: Send hello to 224.0.0.5 area 1 on Serial0/0 from 10.1.12.1
*Mar 3 20:41:26.906: OSPF: Send hello to 224.0.0.5 area 1 on Serial0/0 from 10.1.12.1
*Mar 3 20:41:33.862: %OSPF-5-ADJCHG: Process 1, Nbr 2.2.2.2 on Serial0/0 from FULL to DOWN,
Neighbor Down: Dead timer expired
```

With the default OSPF timers, this can mean a 40 to 120 second state where the neighbor is down, but the router does not recognize that it is down. You can adjust the HelloInterval (and/or the RouterDeadInterval) to reduce this convergence time.



Configuring HelloInterval

Adjusting the HelloInterval changes the interval that Hello messages are sent. When you configure the HelloInterval, the RouterDeadInterval is automatically set to four times the HelloInterval. If you want to speed up OSPF convergence, you can reduce the HelloInterval. The RouterDeadInterval will automatically be reduced as well. You will still need to miss 4 Hello messages from your neighbor, but the process of recognizing a down neighbor will be faster.

```
r1#show ip ospf interface s0/0 | i Timer
Timer intervals configured, Hello 1, Dead 4, Wait 4, Retransmit 5
```

```
r1#debug ip ospf hello
```

```
*Mar  3 20:59:36.630: OSPF: Send hello to 224.0.0.5 area 1 on Serial0/0 from 10.1.12.1
*Mar  3 20:59:37.538: OSPF: Send hello to 224.0.0.5 area 1 on Serial0/0 from 10.1.12.1
*Mar  3 20:59:38.462: OSPF: Send hello to 224.0.0.5 area 1 on Serial0/0 from 10.1.12.1
*Mar  3 20:59:39.366: OSPF: Send hello to 224.0.0.5 area 1 on Serial0/0 from 10.1.12.1
*Mar  3 20:59:40.934: %OSPF-5-ADJCHG: Process 1, Nbr 2.2.2.2 on Serial0/0 from FULL to DOWN,
Neighbor Down: Dead timer expired
```



Configuring RouterDeadInterval

Adjusting the RouterDeadInterval reduces the number of Hello message that can be missed before an OSPF neighbor relationship is declared as down. If you change the RouterDeadInterval, you will want to set it to a multiple of the HelloInterval.

```
r1#show ip ospf interface s0/0 | i Timer
  Timer intervals configured, Hello 10, Dead 20, Wait 20, Retransmit 5
```

```
r1#debug ip ospf hello
```

```
*Mar  3 20:47:06.610: OSPF: Send hello to 224.0.0.5 area 1 on Serial0/0 from 10.1.12.1
*Mar  3 20:47:16.090: OSPF: Send hello to 224.0.0.5 area 1 on Serial0/0 from 10.1.12.1
*Mar  3 20:47:18.566: %OSPF-5-ADJCHG: Process 1, Nbr 2.2.2.2 on Serial0/0 from FULL to DOWN,
Neighbor Down: Dead timer expired
```



OSPF Timer Considerations

If you change your OSPF HelloInterval (and/or RouterDeadInterval) that you need to change these timers on all OSPF-enabled devices on the link. It is also important to take into consideration the normal quality of your link. While you might be tempted to set your HelloInterval to 1 second and your RouterDeadInterval to 2 seconds (or to use sub-second OSPF Fast Hellos), make sure you take into consideration the quality (and possibly speed) of the link. You might improve convergence time at the expense of more OSPF neighbor down incidents or neighbor flapping.

Also, do not set both the HelloInterval and RouterDeadInterval to the same value. You could end up in a situation where the neighbor relationship never comes up or where it comes up only briefly (flaps). Set your RouterDeadInterval to at least two times your HelloInterval.

Also, don't set your RouterDeadInterval to a value that is not a multiple of your HelloInterval.



OSPF Fast Hellos

Remember that **ip ospf dead-interval minimal hello-multiplier** interface command? That command enables sub-second OSPF Hellos. This feature will be covered fully in another lesson.

```
r1(config)# int s0/0
r1(config-if)# ip ospf dead-interval minimal hello-multiplier ?
<3-20> Number of Hellos sent within 1 second
```



Summary

The HelloInterval and RouterDeadInterval are the two timers that you can adjust to speed up network convergence in an OSPF network. The HelloInterval determines the interval between sending OSPF Hello messages on an interface, while the RouterDeadInterval is the interval in which a router must receive an OSPF Hello message from a neighbor before it considers that neighbor to be down.

Cisco IOS assigns a default HelloInterval and RouterDeadInterval to OSPF enabled interfaces. Depending on the interface type, the HelloInterval will be either 10 seconds or 30 seconds. The RouterDeadInterval will be four times the HelloInterval (40 or 120 seconds). A Cisco OSPF-enabled device will maintain a count down timer for each neighbor based on the RouterDeadInterval. Each time receives a Hello message from a neighbor, it will reset this timer to the RouterDeadInterval. If it does not receive a Hello message before this timer expires, then the neighbor will be set to the OSPF DOWN state.

You can adjust the HelloInterval and RouterDeadInterval with the **ip ospf hello-interval** and **ip ospf dead-interval** interface configuration commands. Doing so can reduce OSPF convergence, but you should be careful to take into consideration the quality (and possibly speed) of your links before doing so.