



**OSPF Fast Hellos**



# 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 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)
+ 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)
      .... .... = RS: Restart signal (RS-bit) is NOT set
      .... .... = 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





## Default 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
```



## Configuring OSPF Fast Hellos

OSPF fast hello packets refer to hello packets being sent at intervals of less than 1 second.

OSPF fast hellos are configured using the **ip ospf dead-interval minimal hello-multiplier** interface command. The dead interval is set to 1 second, and the hello-multiplier value is set to the number of hello packets you want sent during that 1 second. You can send between 3 (every 333 milliseconds) and 20 (every 50 milliseconds).\*

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

r1(config-if)#ip ospf dead-interval minimal hello-multiplier 5

r1#show ip ospf interface s0/0
Serial0/0 is up, line protocol is up
  Internet Address 10.1.12.1/24, Area 0
  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 200 msec, Dead 1, Wait 1, Retransmit 5
<--Output Truncated-->
```

\* May depend on platform/IOS version.



## Configuring OSPF Fast Hellos

When fast hello packets are configured on the interface, the hello interval advertised in the hello packets that are sent out this interface is set to 0. **The hello interval in the hello packets received over this interface is ignored.**

Filter: `ip.src == 10.1.12.1 && ospf` Expression... Clear Apply

No. .	Time	Time btwn packets	Source	Destination	Protocol	Info
15	25.000000	25.000000	10.1.12.1	224.0.0.5	OSPF	Hello Packet
16	25.316000	25.316000	10.1.12.1	224.0.0.5	OSPF	Hello Packet
17	25.604000	25.604000	10.1.12.1	224.0.0.5	OSPF	Hello Packet

< [ ]

- Frame 15 (84 bytes on wire, 84 bytes captured)
- 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 Hello Packet
      - Network Mask: 255.255.255.0
      - Hello Interval: 0 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: 1 seconds
      - Designated Router: 0.0.0.0
      - Backup Designated Router: 0.0.0.0
      - Active Neighbor: 2.2.2.2
    - OSPF LLS Data Block



## Configuring OSPF Fast Hellos

The dead interval must be consistent on a segment, whether it is set to 1 second (for fast hello packets) or set to any other value. The hello multiplier need not be the same for the entire segment as long as at least one hello packet is sent within the dead interval [remember: “The hello interval in the hello packets received over this interface is ignored.”].

```
r1#show ip ospf interface s0/0 | i Timer
Timer intervals configured, Hello 200 msec, Dead 1, Wait 1, Retransmit 5
```

```
r2#show ip ospf interface s0/0 | i Timer
Timer intervals configured, Hello 333 msec, Dead 1, Wait 1, Retransmit 5
```

```
r1# show ip ospf neighbor
Neighbor ID      Pri   State           Dead Time   Address      Interface
2.2.2.2          0    FULL/ -         660 msec   10.1.12.2   Serial0/0
```



## Benefits of OSPF Fast Hellos

The benefit of the OSPF Fast Hello Packets feature is that your OSPF network will experience faster convergence time than it would without fast hello packets. This feature allows you to detect lost neighbors within 1 second. It is especially useful in LAN segments, where neighbor loss might not be detected by the Open System Interconnection (OSI) physical layer and data-link layer.

```
r1(config)#int s0/0
r1(config-if)#ip ospf dead-interval minimal hello-multiplier ?
<3-20> Number of Hellos sent within 1 second
r1#show ip ospf interface s0/0 | i Timer
Timer intervals configured, Hello 333 msec, Dead 1, Wait 1, Retransmit 5

r2#show ip ospf interface s0/0 | i Timer
Timer intervals configured, Hello 333 msec, Dead 1, Wait 1, Retransmit 5

*Mar 1 01:07:51.555: OSPF: Rcv hello from 2.2.2.2 area 0 from Serial0/0 10.1.12.2
*Mar 1 01:07:51.555: OSPF: End of hello processing
*Mar 1 01:07:51.847: OSPF: Send hello to 224.0.0.5 area 0 on Serial0/0 from 10.1.12.1
*Mar 1 01:07:52.163: OSPF: Send hello to 224.0.0.5 area 0 on Serial0/0 from 10.1.12.1
*Mar 1 01:07:52.487: OSPF: Send hello to 224.0.0.5 area 0 on Serial0/0 from 10.1.12.1
*Mar 1 01:07:52.807: OSPF: Send hello to 224.0.0.5 area 0 on Serial0/0 from 10.1.12.1
*Mar 1 01:07:52.819: %OSPF-5-ADJCHG: Process 1, Nbr 2.2.2.2 on Serial0/0 from FULL to DOWN,
Neighbor Down: Dead timer expired
```



## OSPF Fast Hellos Considerations

If you find a need for OSPF fast hellos, you will need to ensure that your link will be able to handle the volume and – more importantly – meet the latency criteria for OSPF fast hellos. You’re not going to want to use this feature on slow links as you will only have a RouterDeadInterval of 1 second. OSPF fast hellos are great for LAN segments that require fast convergence (MANs and Data Centers).

I don’t see much advantage in cranking up the number of hello messages sent per second because the RouterDeadInterval is fixed at 1 second regardless of the number of hellos sent per second. I generally just configure 4 hello messages per second (250 milliseconds) on my OSPF fast hello links.



## 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.

OSPF fast hellos allow you to configure sub-second convergence. The RouterDeadInterval is fixed at 1 second while hello messages are sent in fractions of a second. While all OSPF neighbors must agree to user fast hellos, the HelloInterval does not have to match (technically it will match as it will be set to 0). OSPF fast hellos are great for LAN segments that require fast convergence (MANs and Data Centers).