

OSPF: Frequently Asked Questions

Document ID: 9237

Author: Syed Faraz Shamim

Questions

Introduction

Why are loopbacks advertised as /32 host routes in OSPF?

How do I change the reference bandwidth in OSPF?

How does OSPF calculate its metric or cost?

Are OSPF routing protocol exchanges authenticated?

What is the link-state retransmit interval, and what is the command to set it?

What is the purpose of the variable IP-OSPF-Transmit-Delay?

Is it true that only the static option of the virtual link in OSPF allows discontinuous networks, regardless of the mask propagation properties?

Are the multicast IP addresses mapped to MAC-level multicast addresses?

Does the Cisco OSPF implementation support IP TOS-based routing?

Does the offset-list subcommand work for OSPF?

Can an OSPF default be originated into the system based on external information on a router that does not itself have a default?

Can I use the distribute-list in/out command with OSPF to filter routes?

How can I give preference to OSPF interarea routes over intra-area routes?

Do I need to manually set up adjacencies for routers on the Switched Multimegabit Data Service (SMDS) cloud with the OSPF neighbor subcommand?

When routes are redistributed between OSPF processes, are all shortest path first algorithm (SPF) metrics preserved, or is the default metric value used?

How does Cisco accommodate OSPF routing on partial-mesh Frame Relay networks?

Which address-wild-mask pair should I use for assigning an unnumbered interface to an area?

Can I have one numbered side and leave the other side unnumbered in OSPF?

Why do I receive the "cannot allocate router id" error message when I configure Router OSPF One?

Why do I receive the "unknown routing protocol" error message when I configure Router OSPF One?

What do the states DR, BDR, and DROTHER mean in show ip ospf interface command output?

When I issue the show ip ospf neighbor command, why do I only see FULL/DR and FULL/BDR, with all other neighbors showing 2-WAY/DROTHER?

Why do I not see OSPF neighbors as FULL/DR or FULL/BDR on my serial link?

Do I need any special commands to run OSPF over BRI/PRI links?

Do I need any special commands to run OSPF over asynchronous links?

Which Cisco IOS Software release began support for per-interface authentication type in OSPF?

Can I control the P-bit when importing external routes into a not-so-stubby area (NSSA)?

Why are OSPF show commands responding so slowly?

What does the clear ip ospf redistribution command do?

Does OSPF form adjacencies with neighbors that are not on the same subnet?

How often does OSPF send out link-state advertisements (LSAs)?

How do I stop individual interfaces from developing adjacencies in an OSPF network?

When I have two type 5 link-state advertisements (LSAs) for the same external network in the OSPF database, which path should be installed in the IP routing table?

Why is it that my Cisco 1600 router does not recognize the OSPF protocol?

Why is it that my Cisco 800 router does not run OSPF?

Should I use the same process number while configuring OSPF on multiple routers within the same network?

I have a router that runs Cisco Express Forwarding (CEF) and OSPF, who does load-balancing when there are multiple links to a destination?

How does OSPF use two Multilink paths to transfer packets?

How can you detect the topological changes rapidly?

Does the 3825 Series Router support the OSPF Stub feature?

What does the error message %OSPF-4-FLOOD_WAR: Process process-id re-originate LSA ID *ip address* type-2 adv-rtr *ip address* in area *area id* mean?

Can we have OSPF run over a GRE tunnel?

Related Information

Introduction

The document addresses the most frequently asked questions (FAQ) associated with Open Shortest Path First (OSPF). The document covers OSPF version 2 only. OSPF version 3, introduced in Cisco IOS® Software Releases 12.0(24)S, 12.2(18)S, and 12.2(15)T, is used for distributing IP version 6 routing information; it is not explicitly covered in this document. In the scope of this document, "OSPF" refers to OSPF version 2 and "IP" refers to IP version 4.

Q. Why are loopbacks advertised as /32 host routes in OSPF?

A. Loopbacks are considered host routes in OSPF, and they are advertised as /32. For more information, refer to section 9.1 of RFC 2328. In Cisco IOS Software Releases 11.3T and 12.0, if the **ip ospf network point-to-point** command is configured under loopbacks, OSPF advertises the loopback subnet as the actual subnet configured on loopbacks. ISDN dialer interface advertises /32 subnet instead of its configured subnet mask. This is an expected behavior if **ip ospf network point-to-multipoint** is configured.

Q. How do I change the reference bandwidth in OSPF?

A. You can change the reference bandwidth in Cisco IOS Software Release 11.2 and later using the **ospf auto-cost reference-bandwidth** command under **router ospf**. By default, reference bandwidth is 100 Mbps.

Q. How does OSPF calculate its metric or cost?

A. OSPF uses a reference bandwidth of 100 Mbps for cost calculation. The formula to calculate the cost is reference bandwidth divided by interface bandwidth. For example, in the case of Ethernet, it is 100 Mbps / 10 Mbps = 10.

Note: If **ip ospf cost** *cost* is used on the interface, it overrides this formulated cost.

Q. Are OSPF routing protocol exchanges authenticated?

A. Yes, OSPF can authenticate all packets exchanged between neighbors. Authentication may be through simple passwords or through MD5 cryptographic checksums. To configure simple

password authentication for an area, use the command **ip ospf authentication-key** to assign a password of up to eight octets to each interface attached to the area. Then, issue the **area x authentication** command to the OSPF router configuration to enable authentication. (In the command, *x* is the area number.)

Cisco IOS Software Release 12.x also supports the enabling of authentication on a per-interface basis. If you want to enable authentication on some interfaces only, or if you want different authentication methods on different interfaces that belong to the same area, use the **ip ospf authentication** interface mode command.

Q. What is the link-state retransmit interval, and what is the command to set it?

A. OSPF must send acknowledgment of each newly received link-state advertisement (LSA). It does this by sending LSA packets. LSAs are retransmitted until they are acknowledged. The link-state retransmit interval defines the time between retransmissions. You can use the command **ip ospf retransmit-interval** to set the retransmit interval. The default value is 5 seconds.

Q. What is the purpose of the variable IP-OSPF-Transmit-Delay?

A. This variable adds a specified time to the age field of an update. If the delay is not added before transmission over a link, the time in which the link-state advertisement (LSA) propagates over the link is not considered. The default value is 1 second. This parameter has more significance on very low-speed links.

Q. Is it true that only the static option of the virtual link in OSPF allows discontinuous networks, regardless of the mask propagation properties?

A. No, virtual links in OSPF maintain connectivity to the backbone from nonbackbone areas, but they are unnecessary for discontinuous addressing. OSPF provides support for discontinuous networks because every area has a collection of networks, and OSPF attaches a mask to each advertisement.

Q. Are the multicast IP addresses mapped to MAC-level multicast addresses?

A. OSPF sends all advertisements using multicast addressing. Except for Token Ring, the multicast IP addresses are mapped to MAC-level multicast addresses. Cisco maps Token Ring to MAC-level broadcast addresses.

Q. Does the Cisco OSPF implementation support IP TOS-based routing?

A. Cisco OSPF only supports TOS 0. This means that routers route all packets on the TOS 0 path, eliminating the need to calculate nonzero TOS paths.

Q. Does the offset-list subcommand work for OSPF?

A. The **offset-list** command does not work for OSPF. It is used for distance vector protocols such as Interior Gateway Routing Protocol (IGRP), Routing Information Protocol (RIP), and

RIP version 2.

Q. Can an OSPF default be originated into the system based on external information on a router that does not itself have a default?

A. OSPF generates a default only if it is configured using the command **default-information originate** and if there is a default network in the box from a different process. The default route in OSPF is 0.0.0.0. If you want an OSPF-enabled router to generate a default route even if it does not have a default route itself, use the command **default-information originate always**.

Q. Can I use the **distribute-list in/out** command with OSPF to filter routes?

A. The **distribute-list** commands are supported in OSPF but work differently than distance-vector routing protocols such as Routing Information Protocol (RIP) and Enhanced Interior Gateway Routing Protocol (EIGRP). OSPF routes cannot be filtered from entering the OSPF database. The **distribute-list in** command only filters routes from entering the routing table; it does not prevent link-state packets from being propagated. Therefore, this command does *not* help conserve router memory, and it does *not* prohibit a router from propagating filtered routes to other routers.



Caution: Use of the **distribute-list in** command in OSPF may lead to routing loops in the network if not implemented carefully.

The command **distribute-list out** works only on the routes being redistributed by the Autonomous System Boundary Routers (ASBRs) into OSPF. It can be applied to external type 2 and external type 1 routes, but not to intra-area and interarea routes.

Q. How can I give preference to OSPF interarea routes over intra-area routes?

A. According to section 11 of RFC 2328, the order of preference for OSPF routes is:

- ◆ intra-area routes, O
- ◆ interarea routes, O IA
- ◆ external routes type 1, O E1
- ◆ external routes type 2, O E2

This rule of preference cannot be changed. However, it applies only within a single OSPF process. If a router is running more than one OSPF process, route comparison occurs. With route comparison, the metrics and administrative distances (if they have been changed) of the OSPF processes are compared. Route types are disregarded when routes supplied by two different OSPF processes are compared.

Q. Do I need to manually set up adjacencies for routers on the Switched Multimegabit Data Service (SMDS) cloud with the OSPF neighbor subcommand?

A. In Cisco IOS Software releases earlier than Cisco IOS Software Release 10.0, the **neighbor** command was required to establish adjacencies over nonbroadcast multiaccess (NBMA) networks (such as Frame Relay, X.25, and SMDS). With Cisco IOS Software

Release 10.0 and later, you can use the **ip ospf network broadcast** command to define the network as a broadcast network, eliminating the need for the **neighbor** command. If you are not using a fully meshed SMDS cloud, you must use the **ip ospf network point-to-multipoint** command.

Q. When routes are redistributed between OSPF processes, are all shortest path first algorithm (SPF) metrics preserved, or is the default metric value used?

A. The SPF metrics are preserved. The redistribution between them is like redistribution between any two IP routing processes.

Q. How does Cisco accommodate OSPF routing on partial-mesh Frame Relay networks?

A. You can configure OSPF to understand whether it should attempt to use multicast facilities on a multi-access interface. Also, if multicast is available, OSPF uses it for its normal multicasts.

Cisco IOS Software Release 10.0 includes a feature called subinterfaces. You can use subinterfaces with Frame Relay to tie together a set of virtual circuits (VCs) to form a virtual interface, which acts as a single IP subnet. All systems within the subnet should be fully meshed. With Cisco IOS Software Releases 10.3, 11.0 and later, the **ip ospf point-to-multipoint** command is also available.

Q. Which address-wild-mask pair should I use for assigning an unnumbered interface to an area?

A. When an unnumbered interface is configured, it references another interface on the router. When enabling OSPF on the unnumbered interface, use the address-wild-mask pair of interfaces to which the unnumbered interface is pointing.

Q. Can I have one numbered side and leave the other side unnumbered in OSPF?

A. No, OSPF does not work if you have one side numbered and the other side unnumbered. This creates a discrepancy in the OSPF database that prevents routes from being installed in the routing table.

Q. Why do I receive the "cannot allocate router id" error message when I configure Router OSPF One?

A. OSPF picks up the highest IP address as a router ID. If there are no interfaces in up/up mode with an IP address, it returns this error message. To correct the problem, configure a loopback interface.

Q. Why do I receive the "unknown routing protocol" error message when I configure Router OSPF One?

A. Your software may not support OSPF. This error message occurs most frequently with the

Cisco 1600 series routers. If you are using a 1600 router, you need a Plus image to run OSPF.

Q. What do the states DR, BDR, and DROTHER mean in show ip ospf interface command output?

A. DR means designated router. BDR means backup designated router. DROTHER indicates a router that is neither the DR or the BDR. The DR generates a Network Link-State Advertisement, which lists all the routers on that network.

Q. When I issue the show ip ospf neighbor command, why do I only see FULL/DR and FULL/BDR, with all other neighbors showing 2-WAY/DROTHER?

A. To reduce the amount of flooding on broadcast media, such as Ethernet, FDDI, and Token Ring, the router becomes full with only designated router (DR) and backup designated router (BDR), and it shows 2-WAY for all other routers.

Q. Why do I not see OSPF neighbors as FULL/DR or FULL/BDR on my serial link?

A. This is normal. On point-to-point and point-to-multipoint networks, there are no designated routers (DRs) or backup designated routers (BDRs).

Q. Do I need any special commands to run OSPF over BRI/PRI links?

A. In addition to the normal OSPF configuration commands, you should use the **dialer map** command. When using the **dialer map** command, use the **broadcast** keyword to indicate that broadcasts should be forwarded to the protocol address.

Q. Do I need any special commands to run OSPF over asynchronous links?

A. In addition to the normal OSPF configuration commands, you should use the **async default routing** command on the asynchronous interface. This command enables the router to pass routing updates to other routers over the asynchronous interface. Also, when using the **dialer map** command, use the **broadcast** keyword to indicate that broadcasts should be forwarded to the protocol address.

Q. Which Cisco IOS Software release began support for per-interface authentication type in OSPF?

A. Per-interface authentication type, as described in RFC 2178, was added in Cisco IOS Software Release 12.0(8).

Q. Can I control the P-bit when importing external routes into a not-so-stubby area (NSSA)?

A. When external routing information is imported into an NSSA in a type 7 link-state advertisement (LSA), the type 7 LSA has only area flooding scope. To further distribute the external information, type 7 LSAs are translated into type 5 LSAs at the NSSA border. The

P-bit in the type 7 LSA Options field indicates whether the type 7 LSA should be translated. Only those LSAs with the P-bit set are translated. When you redistribute information into the NSSA, the P-bit is automatically set. A possible workaround applies when the Autonomous System Boundary Router (ASBR) is also an Area Border Router (ABR). The NSSA ASBR can then summarize with the **not-advertise** keyword, which results in not advertising the translated type 7 LSAs.

Q. Why are OSPF show commands responding so slowly?

A. You may experience a slow response when issuing OSPF **show** commands, but not with other commands. The most common reason for this delay is that you have the **ip ospf name-lookup** configuration command configured on the router. This command causes the router to look up the device Domain Name System (DNS) names for all OSPF **show** commands, making it easier to identify devices, but resulting in a slowed response time for the commands. If you are experiencing slow response on commands other than just OSPF **show** commands, you may want to start looking at other possible causes, such as the CPU utilization.

Q. What does the clear ip ospf redistribution command do?

A. The **clear ip ospf redistribution** command flushes all the type 5 and type 7 link-state advertisements (LSAs) and scans the routing table for the redistributed routes. This causes a partial shortest path first algorithm (SPF) in all the routers on the network that receive the flushed/renewed LSAs. When the expected redistributed route is not in OSPF, this command may help to renew the LSA and get the route into OSPF.

Q. Does OSPF form adjacencies with neighbors that are not on the same subnet?

A. The only time that OSPF forms adjacencies between neighbors that are not on the same subnet is when the neighbors are connected through point-to-point links. This may be desired when using the **ip unnumbered** command, but in all other cases, the neighbors must be on the same subnet.

Q. How often does OSPF send out link-state advertisements (LSAs)?

A. OSPF sends out its self-originated LSAs when the LSA age reaches the link-state refresh time, which is 1800 seconds.

Q. How do I stop individual interfaces from developing adjacencies in an OSPF network?

A. To stop routers from becoming OSPF neighbors on a particular interface, issue the **passive-interface** command at the interface.

In Internet service provider (ISP) and large enterprise networks, many of the distribution routers have more than 200 interfaces. Configuring **passive-interface** on each of the 200 interfaces can be difficult. The solution in such situations is to configure all the interfaces as passive by default using a single **passive-interface default** command. Then, configure individual interfaces where adjacencies are desired using the **no passive-interface** command. For more information, refer to Default Passive Interface Feature.

There are some known problems with the **passive-interface default** command. Workarounds are listed in Cisco bug ID CSCdr09263 (registered customers only) .

Q. When I have two type 5 link-state advertisements (LSAs) for the same external network in the OSPF database, which path should be installed in the IP routing table?

A. When you have two type 5 LSAs for the same external network in the OSPF database, prefer the external LSA that has the shortest path to the Autonomous System Boundary Router (ASBR) and install that into the IP routing table. Use the **show ip ospf border-routers** command to check the cost to the ASBR.

Q. Why is it that my Cisco 1600 router does not recognize the OSPF protocol?

A. Cisco 1600 routers require the Plus feature set image of Cisco IOS Software to run OSPF. Refer to Table 3: Cisco 1600 Series Routers Feature Sets in the Release Notes for Cisco IOS Release 11.2(11) Software Feature Packs for Cisco 1600 Series Routers for more information.

Q. Why is it that my Cisco 800 router does not run OSPF?

A. Cisco 800 routers do not support OSPF. However, they do support Routing Information Protocol (RIP) and Enhanced Interior Gateway Routing Protocol (EIGRP). You can use the Software Advisor (registered customers only) tool for more information on feature support.

Q. Should I use the same process number while configuring OSPF on multiple routers within the same network?

A. OSPF, unlike Border Gateway Protocol (BGP) or Enhanced Interior Gateway Routing Protocol (EIGRP), does not check the process number (or autonomous system number) when adjacencies are formed between neighboring routers and routing information is exchanged. The only case in which the OSPF process number is taken into account is when OSPF is used as the routing protocol on a Provider Edge to Customer Edge (PE-CE) link in a Multiprotocol Label Switching (MPLS) VPN. PE routers mark OSPF routes with the domain attribute derived from the OSPF process number to indicate whether the route originated within the same OSPF domain or from outside it. If the OSPF process numbering is inconsistent on PE routers in the MPLS VPN, the **domain-id** OSPF mode command should be used to mark that the OSPF processes with different numbers belong to the same OSPF domain.

This means that, in many practical cases, you can use different autonomous system numbers for the same OSPF domain in your network. However, it is best to use consistent OSPF-process numbering as much as possible. This consistency simplifies network maintenance and complies with the network designer intention to keep routers in the same OSPF domain.

Q. I have a router that runs Cisco Express Forwarding (CEF) and OSPF, who does load-balancing when there are multiple links to a destination?

A. CEF works by performing the switching of the packet based on the routing table which is populated by the routing protocols such as OSPF. CEF does the load-balancing once the routing protocol table has been calculated. For more details on load balancing, refer to How

does load-balancing work?

Q. How does OSPF use two Multilink paths to transfer packets?

A. OSPF uses the metric aCost, which is related to the bandwidth. If there are equal cost paths (the same bandwidth on both multilinks), OSPF installs both routes in the routing table. The routing table tries to use both links equally, regardless of the interface utilization. If one of the links in the first multilink fails, OSPF does not send all the traffic down the second multilink. If the first multilink peaks 100%, OSPF does not send any traffic down the second multilink because OSPF tries to use both links equally, regardless of the interface utilization. The second is used fully only when the first multilink goes down.

Q. How can you detect the topological changes rapidly?

A. In order to have a rapid fault detection of topology changes, the hello timer value needs to be set to 1 second. The hold timer value, which is four times that of the hello timer, also needs to be configured. There is a possibility of more routing traffic if the hello and hold timer values are reduced from their default values.

Q. Does the 3825 Series Router support the OSPF Stub feature?

A. Yes, the 3800 Series Router that runs Advanced IPServices image supports the OSPF Stub feature.

Q. What does the error message %OSPF-4-FLOOD_WAR: Process process-id re-originates LSA ID *ip address* type-2 adv-rtr *ip address* in area *area id* means?

A. The error message is due to the some router that is flushing the network LSA because the network LSA received by the router whose LSA ID conflicts with the IP address of one of the router's interfaces and flushes the LSA out of the network. For OSPF to function correctly the IP addresses of transit networks must be unique. If it is not unique the conflicting routers reports this error message. In the error message the router with the OSPF router ID reported as adv-rtr reports this message.

Q. Can we have OSPF run over a GRE tunnel?

A. Yes, refer to Configuring a GRE Tunnel over IPsec with OSPF.

Related Information

- [OSPF Support Page](#)
- [Technical Support – Cisco Systems](#)

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2008 – 2009 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)