Review OSPF Frequently Asked Questions

Contents

Introduction

Background Information

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

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

Q. How does OSPF calculate its metric or cost?

Q. What algorithm is used by OSPF if equal cost routes exist?

Q. Are OSPF routing protocol exchanges authenticated?

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

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

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

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

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

Q. Does theoffset-listsubcommand work for OSPF?

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

Q. Can I use thedistribute-list in/outcommand with OSPF to filter routes?

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

Q. Do I need to manually set up adjacenciesfor routers on the Switched Multimegabit Data Service (SMDS) cloud with the OSPFneighborsubcommand?

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

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

Q. Which address-wild-mask pair must I use to assignan unnumbered interface to an area?

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

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

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

Q. What do the statesDR,BDR, andDROTHERmean inshow ip ospf interfacecommand output?

Q. When I issue theshow ip ospf neighborcommand, why do I only seeFULL/DRandFULL/BDR, with all other neighbors that show2-WAY/DROTHER?

Q. Why do I not see OSPF neighbors asFULL/DRorFULL/BDRon my serial link?

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

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

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

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

Q. Why do OSPFshowcommands respond so slowly?

Q. What does theclear ip ospf redistributioncommand do?

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

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

Q. How do I prevent individual interfaces that developadjacenciesin an OSPF network?

Q. When I have two type 5 link-state advertisements (LSAs) for the same external network in the OSPF

database, which path must be installed in the IP routing table?

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

Q. Must I use the same process number when I configureOSPF on multiple routers within the same network?

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?

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

Q. How can you detect the topological changes rapidly?

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

Q. What does the error message %OSPF-4-FLOOD_WAR: Process process-id re-originates LSA IDip addresstype-2 adv-rtrip addressin areaarea idmeans?

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

Q. Is there a way to manipulate and prefer the Type 3 LSAs to originate from two different areas sent to the non-backbone area?

Q. Is there a drop/flap of an OSPF neighborship when an OSPF area type is changed from nssa no-summary to nssa?

Q. In the%OSPF-5-ADJCHG: Process ID, Nbr [ip-address] on Port-channel31 from FULL to EXSTART, SeqNumberMismatcherror message, what doesSeqNumberMismatchsignify?

Q. What is the maximum number of OSPF processes (VRF aware) on 7600/6500 platforms?

Q. How does ISPF impact or improve the OSPF network?

Q. Is there a way to compare Cisco NX-OS/Cisco IOS OSPF commands?

Q. Is there any feature of OSPF protocol for quick convergence and a slow re-convergence of routes?

Q. What doesBADSEQNUMin the%OSPF-5-NBRSTATE: ospf-101 [5330] Process 101, Nbr 10.253.5.108 on Vlan7 02 from FULL to EXSTART, BADSEQNUMOSPF log message mean?

Related Information

Introduction

The document describes the most frequently asked questions (FAQ) associated with the Open Shortest Path First (OSPF) protocol.

Background Information

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 to distribute 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 ofRFC 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.

For example, consider two routers (R1 and R2) connected via FastEthernet interface. R1 has the loopback configured with the **ip ospf network point-to-point** command and advertises the loopback in OSPF.

```
<#root>

interface Loopback0
 ip address 10.1.1.1 255.255.255.0


ip ospf network point-to-point
```

When checked in router R2 with the**show ip route ospf**command, the route 10.1.1.1 is seen as:

```
<#root>

!..output truncated


 10.0.0.0/24 is subnetted, 1 subnets

O       10.1.1.0 [110/11] via 10.1.1.1, 00:00:02, FastEthernet0/0
```

However, when the ip ospf network point-to-pointcommand is removed from R1 to 0 interface, the route 10.1.1.1 on R2 is seen as:

```
<#root>


10.0.0.0/32 is subnetted, 1 subnets


O       10.1.1.1 [110/11] via 10.1.1.1, 00:00:01, FastEthernet0/0
```

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 with the ospf auto-cost reference-bandwidthcommand under router ospf. By default, reference bandwidth is 100 Mbps. The ospf link-cost is a 16-bit number. Therefore, the maximum value supported is 65,535.

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 costis used on the interface, it overrides this formulated cost. For more information, refer toOSPF Cost.

Q. What algorithm is used by OSPF if equal cost routes exist?

A.If equal cost routes exist, OSPF uses CEF load balancing. For more information, refer toTroubleshooting Load Balancing Over Parallel Links with Cisco Express Forwarding.

Q. Are OSPF routing protocol exchanges authenticated?

A.Yes, OSPF can authenticate all packets exchanged between neighbors. Authentication can 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 authenticationcommand to the OSPF router configuration to enable authentication. (In the command,*x*is the area number.)

Cisco IOS Software Release 12.x also supports and enables 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 authenticationinterface 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). So, it sends 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 the non-backbone 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 with 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, which eliminates the need to calculate nonzero TOS paths.

Q. Does theoffset-listsubcommand 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 with the command default-information originateand 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 thedistribute-list in/outcommand with OSPF to filter routes?

A.The distribute-listcommands 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 when they enters the OSPF database. The distribute-list incommand only filters routes so that they do not enter the routing table; it does not prevent propagation of link-state packets. Therefore, this command does*not*help conserve router memory, and it does*not*prohibit a router that propagates from filtered routes to other routers.

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

The command**distribute-list out**works only on the routes that are 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 inter-area routes.

Refer to configuration example of distribute-list in OSPF,

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

A. From Section 11 of the RFC 2328, the order of preference for OSPF routes is:

- intra-area routes, O
- inter-area 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 that runs 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 OSPFneighborsubcommand?

A.In Cisco IOS Software releases earlier than Cisco IOS Software Release 10.0, the neighborcommand 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, which eliminates the need for the neighborcommand. If you do not use a fully meshed SMDS cloud, you must use the ip ospf network point-to-multipointcommand.

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 must 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 must 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 must I use to assign an unnumbered interface to an area?

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

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 installation 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 cannot support OSPF. This error message occurs most frequently with the Cisco 1600 series routers. If you are use a 1600 router, you need a Plus image to run OSPF.

Q. What do the statesDR,BDR, andDROTHERmean inshow ip ospf interfacecommand output?

A.DRmeans designated router.BDRmeans backup designated router.DROTHERindicates 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 theshow ip ospf neighborcommand, why do I only seeFULL/DRandFULL/BDR, with all other neighbors that show2-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 shows2-WAYfor all other routers.

Q. Why do I not see OSPF neighbors asFULL/DRorFULL/BDRon 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 must use the dialer mapcommand. When the dialer mapcommand is used, use the broadcastkeyword to indicate that broadcasts must 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 must use the async default routingcommand on the asynchronous interface. This command enables the router to pass routing updates to other routers over the asynchronous interface. Also, when the dialer mapcommand is used, use the broadcastkeyword to indicate that broadcasts must 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 inRFC 2178, was added in Cisco IOS Software Release 12.0(8).

Q. Can I control the P-bit when I import 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 must 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 do OSPFshowcommands respond so slowly?

A.You can experience a slow response when you issue the OSPF showcommands, but not with other commands. The most common reason for this delay is that you have the ip ospf name-lookupconfiguration command configured on the router. This command causes the router to look up the device Domain Name System (DNS) names for all OSPF showcommands, and makes it easier to identify devices, but results in a slowed response time for the commands. If you experience slow response on commands other than just OSPF showcommands, you need to look at other possible causes, such as the CPU utilization.

Q. What does theclear ip ospf redistributioncommand 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 can 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 is helpful if you use the ip unnumberedcommand, 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. For more information, refer toLink-State Advertisements.

Q. How do I prevent individual interfaces that develop adjacencies in an OSPF network?

A. So that routers do not become 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. To configure passive-interfaceon each of the 200 interfaces can be difficult. The solution in such situations is to configure all the interfaces as passive by default a single passive-interface defaultcommand. Then, configure individual interfaces where adjacencies are desired with the no passive-interfacecommand.

There are some known problems with the no passive-interface defaultcommand. Workarounds are listed inCisco bug ID CSCdr09263.

Note: Only registered Cisco users can access internal Cisco bug information.

Q. When I have two type 5 link-state advertisements (LSAs) for the same external network in the OSPF database, which path must 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-routerscommand to check the cost to the ASBR.

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** tool for more information on feature support.

Note: Only registered Cisco users can access internal Cisco bug information.

Q. Must I use the same process number when I configure 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 neighbor 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 Client 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-idOSPF mode command must 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 performs 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 toUnderstand how Load Balancing Works.

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 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.

Note:When OSPF Timers are tuned it can result in network as well device resource overhead. Cisco recommends that you use Bidirectional Forwarding Detection (BFD) and not tune the routing protocol timers. BFD also gives sub-second convergence. Refer to OSPF Support for BFD over IPv4 for more information.

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 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 toConfigure a GRE Tunnel over IPSec with OSPF.

Q. Is there a way to manipulate and prefer the Type 3 LSAs to originate from two different areas sent to the non-backbone area?

A.Type 3 LSA is originated by the Area Border Router (ABR) as a summary route.
The summary route cannot be manipulated in an ABR router.

Q. Is there a drop/flap of an OSPF neighborship when an OSPF area type is changed from nssa no-summary to nssa?

A.When the NSSA ABR is configured to move from nssa no-summary to nssa, the OSPF neighborship does not flap.

Q. In the`%OSPF-5-ADJCHG: Process ID, Nbr [ip-address] on Port-channel31 from FULL to EXSTART, SeqNumberMismatch`error message, what does`SeqNumberMismatch`signify?

A.The OSPF neighbor was changed state from FULL to EXSTART because of the receipt of a Database Description (DBD) packet from the neighbor with an unexpected sequence number.

`SeqNumberMismatch`means that a DBD packet during OSPF neighborship negotiation has been received that either:

- has an unexpected DBD sequence number

- unexpectedly has the Init bit set
- has an Options field that differs from the last Options field received in a Database Description packet.

Q. What is the maximum number of OSPF processes (VRF aware) on 7600/6500 platforms?

A.Cisco IOS has a limit of 32 routing processes. Two of these are saved for static and directly connected routes. The Cisco 7600 router supports 28 OSPF processes per VRF.

Q. How does ISPF impact or improve the OSPF network?

A.Incremental SPF is more efficient than the full SPF algorithm, which allows OSPF to converge faster on a new routing topology in reaction to a network event. The incremental SPF is designed in such a way that it only updates the affected nodes and does not rebuild the whole tree. This results in a faster convergence and saves CPU cycles because the unaffected nodes do not need to be processed. With a best practice ISPF would make more of a difference for a large OSPF domain.

Incremental SPF provides greater improvements in convergence time for networks with a high number of nodes and links. Incremental SPF also provides a significant advantage when the changes in the network topology are further away from the root of the SPT; for example, the larger the network the more significant the impact. A segment of 400-1000 nodes must see improvements. However, it can be hard to verify in a deployed production network without some kind of facility or tool to measure the end-to-end delay. For more information, refer toOSPF Incremental SPF.

Q. Is there a way to compare Cisco NX-OS/Cisco IOS OSPF commands?

A.Yes, refer toCisco NX-OS/Cisco IOS OSPF Comparison.

Q. Is there any feature of OSPF protocol for quick convergence and a slow re-convergence of routes?

A.The OSPF Shortest Path First Throttling feature makes it possible to configure SPF scheduling in millisecond intervals and to potentially delay SPF calculations during network instability. SPF is scheduled to calculate the Shortest Path Tree (SPT) when there is a change in topology.

Syntax of the command under OSPF:

timers throttle spf[spf-start] [spf-hold] [spf-max-wait]

Where:

- **spf-start**—Initial delay to schedule an SPF calculation after a change, in milliseconds. Range is from 1 to 600000.
- **spf-hold**—Minimum hold time between two consecutive SPF calculations, in milliseconds. Range is from 1 to 600000.
- **spf-max-wait**—Maximum wait time between two consecutive SPF calculations, in milliseconds. Range is 1 to 600000.

Q. What does`BADSEQNUM`in the`%OSPF-5-NBRSTATE: ospf-101 [5330] Process 101, Nbr 10.253.5.108 on Vlan7 02 from FULL to EXSTART, BADSEQNUM`OSPF log message mean?

A.This message is related to the DBD exchange process, which uses a sequence number for the synchronization of the database. For some reason a bad sequence number was reported in the DBD packet. This can occur because of transient conditions, which includes packet loss or packet corruption.

Related Information

- **OSPF Support Page**
- **Cisco Technical Support & Downloads**