6AOSCG0050-29A
May 2013

ADTRAN

Configuration Guide

OSPFv3 In AOS

This configuration guide describes the Open Shortest Path First version 3 (OSPFv3) for Internet Protocol version 6 (IPv6) configuration and use in ADTRAN Operating System (AOS) products. This guide includes an overview of OSPFv3 and its differences from OSPFv2, configuration of OSPFv3 globally and per-interface on the AOS product, and OSPFv3 troubleshooting information. In addition, an OSPFv3 configuration example is included.

This guide contains the following sections:

IPv6 OSPFv3 Overview

OSPFv3 is the newest version of OSPF and includes IPv6 support by using IPv6 messages to calculate IPv6 routes. Dual stack networks can use an OSPFv2 instance to support IPv4 connectivity and an OSPFv3 instance to support IPv6 connectivity. OSPFv3 functions by allowing routers to resolve paths through the network using OSPFv3 messages sent within IPv6 packets. *Figure 1* illustrates OSPFv3 function in a simple network. In this illustration, a host on the 2001:DB8:1::/64 network communicates with a host on the 2001:DB8:2::/64 network using an IPv6 connection where routes are resolved using OSPFv3.

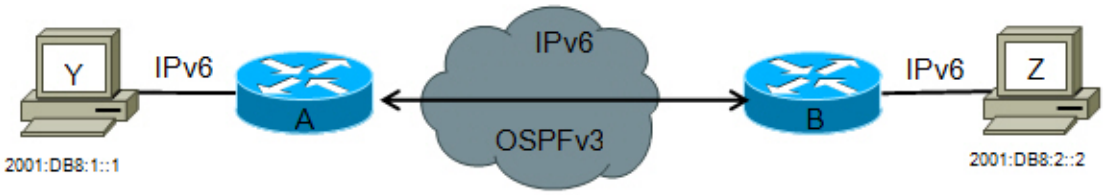

Figure 1. OSPFv3 Over an IPv6 Network

OSPFv3 functions in a similar manner to OSPFv2, with a few crucial exceptions. The relationship of OSPFv3 to OSPFv2 is outlined in the following section.

OSPFv3 versus OSPFv2

OSPFv3 and OSPFv2 are entirely independent, and they can both be enabled on a system or a link without interacting. Both protocols:

- Use protocol 89
- Use the same basic packet types (Hello, Database Description (DBD), Link State Request (LSR), Link State Update (LSU), Link State Acknowledgement (LSAck))
- Form neighbors and adjacencies in the same manner (maximum transmission unit (MTU) settings must match between interfaces on a link in order to form adjacencies)
- Have the same election and operation of designated router and backup designated router
- Have the same Link State Advertisement (LSA) flooding and aging mechanisms
- Use the same interface types
- Have the same default administrative distances for routes (110)
- Use equivalent multicast groups (FF02::5 for OSPFv3)
- Use equivalent DR multicast groups (FF02::6 for OSPFv3)
- Use the same algorithms

However, the major difference between OSPFv2 and OSPFv3 is that OSPFv2 depends heavily on IPv4 addresses for its operation, whereas OSPFv3 must rely on other sources because it uses IPv6. The differences between the two OSPF versions are outlined in *Table 1*.

Table 1. OSPFv2 and OSPFv3 Comparison

OSPFv2	OSPFv3
Protocol processing is per-subnet.	Protocol processing is per-link.
Interface is OSPF enabled through a network command within the router OSPF configuration.	Interface is OSPF enabled through direct configuration on the interface.
Options field in Hello and Database Decription packets is 8 bits.	Options field in Hello and Database Description packets is 24 bits.
Uses a network mask for adjacencies and does not include an interface ID in the Hello packet.	Uses an interface ID instead of a network mask for adjacencies, and includes the interface ID in Hello packets to identify the originating router's interface to the link.
No R-bit in the Options field.	R-bit is included in the Options field, and is used to allow an OSPF speaker to participate in OSPF topology distribution without being used to forward transit traffic.
No V6-bit in the Options field.	V6-bit is included in the Options field, and is used to specialize the R-bit. If the V6 bit is clear, an OSPF speaker can participate in OSPF topology distribution without being used to forward IPv6 datagrams. If the R-bit is set, and the V6 bit is clear, then IPv6 datagrams are not forwarded.
Does not support IPv6.	Supports IPv6. In addition, even when OSPFv3 may convey IPv4 or other route information, all messages are sent over IPv6 packets.
Supports a single OSPF protocol instance.	Supports multiple OSPF protocol instances on a single link.
Does not support OSPF address families.	Supports OSPF address families to separate OSPFv3 instances.
Have AuType and Authentication fields in the OSPF packet header and support authentication.	Authentication has been removed from the OSPF protocol, and the AuType and Authentication fields have been removed from the OSPF packet header.
Supports only area and autonomous systems (AS)-wide link state advertisements (LSAs) flooding and a minimal link-local flood scope.	Supports a generalized flooding scope for LSAs and has the scope explicitly coded in the LSA's link state type field.
Routing code for OSPF Intra-area: O	Routing code for OSPF Intra-area: O
Routing code for OSPF Inter-area: IA	Routing code for OSPF Inter-area: OI
Routing code for OSPF External Type 1: E1	Routing code for OSPF External Type 1: OE1
Routing code for OSPF External Type 2: E1	Routing code for OSPF External Type 2: OE2
Routing code for OSPF NSSA External Type 1: N1	Routing code for OSPF NSSA External Type 1: ON1
Routing code for OSPF NSSA External Type 2: N2	Routing code for OSPF NSSA External Type 2: ON2

OSPFv3 Considerations

There are many considerations to process when beginning to configure OSPFv3. The following sections outline changes to OSPF that relate to OSPFv3 only.

IPv6 Interface Addresses and Prefixes

When IPv6 is enabled on an interface, all IPv6 addresses whose prefix is on-link and assigned to that interface are included in OSPF processing using the prefix length specified for each address. If an IPv6 address is configured with a /128 prefix, OSPFv3 advertises the prefix using the /128 prefix. Routes that are created by adding on-link prefixes to an interface are NOT included in OSPF processing. For example, if you use the **ipv6 nd prefix** command from the interface's configuration mode, you will create an on-link prefix that adds a directly connected route to the IPv6 route table. However, even if the interface has OSPF enabled, that route is not introduced to the OSPF process.

If multiple routers are connected and configured for OSPF on a given link, and at least one router has an IPv6 address in a given prefix, any router that does not have an address in that prefix will do one of the following:

- Learn the prefix using OSPF.
- Insert an OSPF route for that prefix into the IPv6 route.
- Set the route's next hop to the router's own interface that is connected to the link (effectively treating the prefix as if it's on-link).
- Make no changes to the interface's prefix list and only make changes to the route table.

OSPFv3 for IPv6 assumes that each router has been assigned a link-local unicast address on each of its interfaces. On all OSPF interfaces (except virtual links), OSPF messages are sent using the interface's link-local unicast address as the source address, which allows OSPFv3 interfaces to communicate with other routers on the same link regardless of if or how prefixes are assigned to the interfaces. Networks learned through OSPFv3 are stored in the route table with a link-local next-hop address, which allows routers to learn the link-local addresses of all other routers attached to its link and to use these addresses as next-hop information during packet forwarding. Link-local addresses appear in OSPF link-LSAs, however, link-local addresses are not allowed in other OSPF LSA types. Link-local addresses MUST NOT be advertised in inter-area-prefix-LSAs, AS-external-LSAs, not-so-stubby-area (NSSA)-LSAs, or intra-area-prefix-LSAs.

OSPFv3 Authentication

In OSPFv3 for IPv6, authentication has been removed from the OSPF protocol. When running over IPv6, OSPF relies on the IP Encapsulating Security Payload (ESP) and IP Authentication Header (AH) protocols to ensure integrity and authentication or confidentiality of routing exchanges (refer to RFC 4552). Protection of OSPF packet exchanges against accidental data corruption is provided by the standard IPv6 upper layer checksum, which covers the entire OSPF packet and prepended IPv6 pseudo-header.

In addition, because routers act as hosts when performing OSPF, they are required to support IPsec transport modes. Therefore, support for ESP is included in OSPFv3, and AH is allowed. Authentication is required using either ESP or AH, while encryption is optional. However, only ciphers suitable for manual keys are allowed. AOS supports IPsec for OSPFv3, and thus relies on the IPv6 AH and ESP to provide integrity, authentication, and confidentiality to OSPFv3 packets.

In AOS, OSPFv3 authentication is completed using IPsec. IPsec protection of OSPF can be configured at the interface or OSPF area level. Area-level authentication eases configuration management when the same security is to be applied to multiple interfaces. When protection is specified at the interface level, the configured security parameter index (SPI) is used only at that interface. When protection is specified at the area level, the SPI is used at each interface in that area.

To balance management advantages of area-level protection with concerns over interoperability and consistency when using multiple OSPFv3 instances at an interface, AOS supports protection at the area level, along with a set of rules to choose one set of security associations (SAs) for an interface when multiple protections are configured. When multiple OSPFv3 protections are specified that affect the same interface, one set is selected to protect all OSPFv3 instances on that interface. The set is chosen using the following rules processed in this order:

- The protection configured at the interface level (including null), then everything else
- The protection configured at the IPv6 address family instance, then everything else
- The protection configured at the IPv4 address family instance, then everything else (IPv4 address family is not supported as of AOS firmware release R10.5.0)
- No protection

When OSPFv3 authentication is enabled, received OSPFv3 packets that are not protected with AH or ESP, or packets that fail authentication checks, are discarded. Similarly, when OSPFv3 confidentiality is enabled, OSPFv3 packets that are not protected with ESP, or that fail the confidentiality checks, are discarded.

In AOS, IPsec uses the IP header to distinguish between IPv4 and IPv6 OSPF packets. Each AOS product that supports IPsec has a limited number of IPsec SAs it can support. In previous firmware versions, that limit consisted of only IPv4 IPsec SAs. With the addition of IPv6 IPsec and IPsec protection of OSPFv3, that limit is now a shared space of SA types. The space can be consumed by all IPv4 SAs, all IPv6 SAs, all SAs for OSPF, or any other combination. You should keep this in mind when configuring IPsec SAs for your networking applications.

In the AOS CLI, IPsec encryption and authentication algorithm keys are used in OSPFv3 commands. These keys are specified as hex characters, and their sizes for each algorithm are outlined in *Table 2*.

Table 2. IPsec Encryption/Authentication Key Types and Sizes

Key Type	Size (bytes)
3DES	48
AES-CBC 128	32
AES-CBC 192	48
AES-CBC 256	64
DES	16
MD5	32
SHA1	40

OSPFv3 Firewall and VPN Interactions

Except when using IPsec protection, OSPFv3 traffic is subject to firewall policies, and associations are created to track the traffic. When IPsec protected, OSPFv3 traffic is not subject to firewall policies, and the firewall does not interact with OSPFv3.

When OSPFv3 traffic is not wrapped in either ESP or AH, the following will occur:

- Unicast OSPFv3 packets to the router are processed by the firewall. If the firewall is enabled, a policy must be configured that allows OSPFv3 so that it can function properly.
- Multicast OSPFv3 traffic is not processed by the firewall. Therefore, it is possible that partial OSPF interaction might occur if a policy is not set to also allow unicast OSPFv3 traffic.
- Any OSPFv3 traffic that is destined to the receiving interface, or to be routed through the device, if the traffic is unencapsulated and the packet matched a stateful or stateless allow policy-class entry, then a new association with protocol 89 and Layer 3 source and destination information is created, and the packet is allowed. If the traffic is unencapsulated and the packet matches an explicit or implicit discard policy-class entry, the packet is dropped. If the traffic is encapsulated in ESP or AH, then traffic to or from the unit is decrypted/encrypted, and firewall sessions are not created.

Addressing Semantics

All addressing semantics have been removed from OSPF packet headers in OSPFv3. All addressing information is contained in the various LSA types only. IPv6 addresses are not present in OSPF packets, except in LSA payloads carried by the link state update packets. Therefore, router-LSAs and network-LSAs no longer contain network addresses, but rather simply express network topology information. OSPF router IDs, area IDs, and LSA link state IDs remain at the IPv4 size of 32 bits and are not assigned IPv6 addresses.

In OSPFv3, neighboring routers are always identified by router ID, rather than by an IPv4 address as with OSPFv2. This change affects the reception of OSPF packets, the lookup of neighbors, and the reception of Hello packets. In addition, the router ID of **0.0.0.0** is reserved and SHOULD NOT be used.

Stub and NSSA Area Support

OSPFv3 stub and NSSA areas function similarly to stub and NSSA areas in OSPFv2. The difference is that with IPv6, of the mandatory LSA types for OSPF, stub areas carry only these: router-LSAs, network-LSAs, inter-area-prefix-LSAs, link-LSAs, and intra-area-prefix-LSAs. NSSA areas are restricted to these types, with the addition of NSSA-LSAs.

Flooding Scope

For OSPFv3, flooding scope for LSAs has been generalized and is explicitly encoded in the LSA's link-scope type field. There are now three separate flooding scopes for LSAs:

- Link-local scope, in which the LSA is only flooded on the local-link and no further. This is used for link LSAs.
- Area scope, in which the LSA is only flooded throughout a single OSPF area. This is used for router, network, inter-area-prefix, inter-area-router, intra-area-prefix, and NSSA LSAs.
- AS scope, in which the LSA is flooded throughout the routing domain except in stub and NSSA areas. This is used for AS-external LSAs. A router that originates AS-scoped LSAs is considered an AS boundary router (ASBR) and will set its E-bit in router-LSAs for regular areas.

Handling Unknown LSA Types

In OSPFv2, unknown LSAs were discarded. In OSPFv3, the protocol allows a mixture of router capabilities on a single link, so these LSAs are not discarded. Rather, the handling of unknown LSA types has been made more flexible in OSPFv3. Based on the link-scope type, unknown LSA types are either treated as having link-local flooding scopes, or they are stored and flooded as if they were understood. This behavior is explicitly coded in the LSA handling bit (U-bit) of the link-state header's link-state type field. When set, the LSA is flooded within its encoded scope as if it were understood. When this bit is not set, the LSA is flooded within the link-scope.

OSPFv3 and Address Families

In OSPFv3, the concept of address families is introduced. Address families use entirely separate OSPF processes for each family. Each process maps its messages to or from the wire using a different instance ID. The instance ID is a field in the OSPF header with a value from **0** to **255**. The value itself conveys no information, other than the instance number; however, the instance ID adheres to a range of values that convey address family information (refer to RFC 5838). *Table 3* outlines the address family information associated with each instance ID value.

Table 3. Instance ID and Address Family Information

Instance ID Range	Address Family	Default Value
0 to 31	IPv6 Unicast	0
32 to 63	IPv6 Multicast	32
64 to 95	IPv4 Unicast	64
96 to 127	IPv4 Multicast	96
128 to 255	Unassigned	

As of AOS firmware release R10.5.0, only IPv6 unicast address families are supported for OSPFv3. This configuration is backward compatible OSPFv3 devices that support only IPv6 but do not support address families, as long as those devices use the instance ID for the ***Base IPv6 Unicast AF*** *(value of* ***0****). The AOS router only accepts OSPF messages that do not set the address family bit when the instance ID is 0.*

Before an OSPFv3 process address family can be created, routing for the address family must be enabled at the global level (using the command **ipv6 unicast-routing**). Before an OSPFv3 process can be enabled at an interface, routing for the address family must be enabled at that interface (using the command **ipv6** or a variation of the **ipv6 address** command). Globally disabling routing for an address family does not remove existing OSPFv3 commands associated with that address family. Disabling an address family at an interface removes OSPFv3 commands associated with that address family at that interface.

OSPFv3 Message Encapsulation

OSPFv3 messages are sent in IPv6 packets with the following characteristics:

- All messages use IPv6 protocol 89 (indicated in the Next header field).
- Messages sent to the AllSPFRouters address (FF02::5) are received by all routers running OSPF. Hello packets are always sent to this destination, and certain OSPF protocol packets are sent to this address during a flooding procedure.
- Messages sent to the AllDRouters address (FF02::6) are received by both the designated router and designated backup router. Certain OSPF protocol packets are sent to this address during a flooding procedure.
- The traffic class field of OSPFv3 messages is set to DSCP value **CS6** (0 x C0).
- Large OSPFv3 messages can be sent as multiple separate messages when possible (preferred in AOS), or they can use the local stack's IPv6 fragmentation (up to **65535** bytes).

OSPFv3 Routing

To analyze an OSPFv3 metric of a given path in the network, follow the path from the link being represented to the desired point of reference. Each time the path enters a router interface, add the OSPFv3 cost of that interface to the path cost. Only metrics of the same type are directly comparable. OSPFv3 routes are preferred in the following (decreasing) order: OSPF intra-area, OSPF inter-area, OSPF external type 1, OSPF external type 2, OSPF NSSA external type 1, and OSPF NSSA external type 2. The default administrative distance for routes generated by OSPFv3 is **110**.

In AOS, OSPFv3 class-specific capabilities are not supported. This means that OSPFv3 in AOS does not make policy or routing decisions based on the value of a packet's IPv6 traffic class or IPv4 type of service (ToS) or precedence field.

VRFs and OSPFv3

OSPFv3 supports per-virtual routing and forwarding (VRF) operation, meaning the VRF is specified when creating the OSPFv3 process. All interfaces enabled with this OSPFv3 process must already be a member of that VRF. Routes resolved by a specific OSPFv3 process are placed in the route table belonging to that VRF. Changing the VRF settings at an interface removes all OSPF configuration from that interface. Deleting a VRF removes all OSPFv3 configurations associated with that VRF globally across the router. *Figure 2 on page 9* illustrates the relationships between interfaces, address families, OSPFv3 processes, and VRFs.

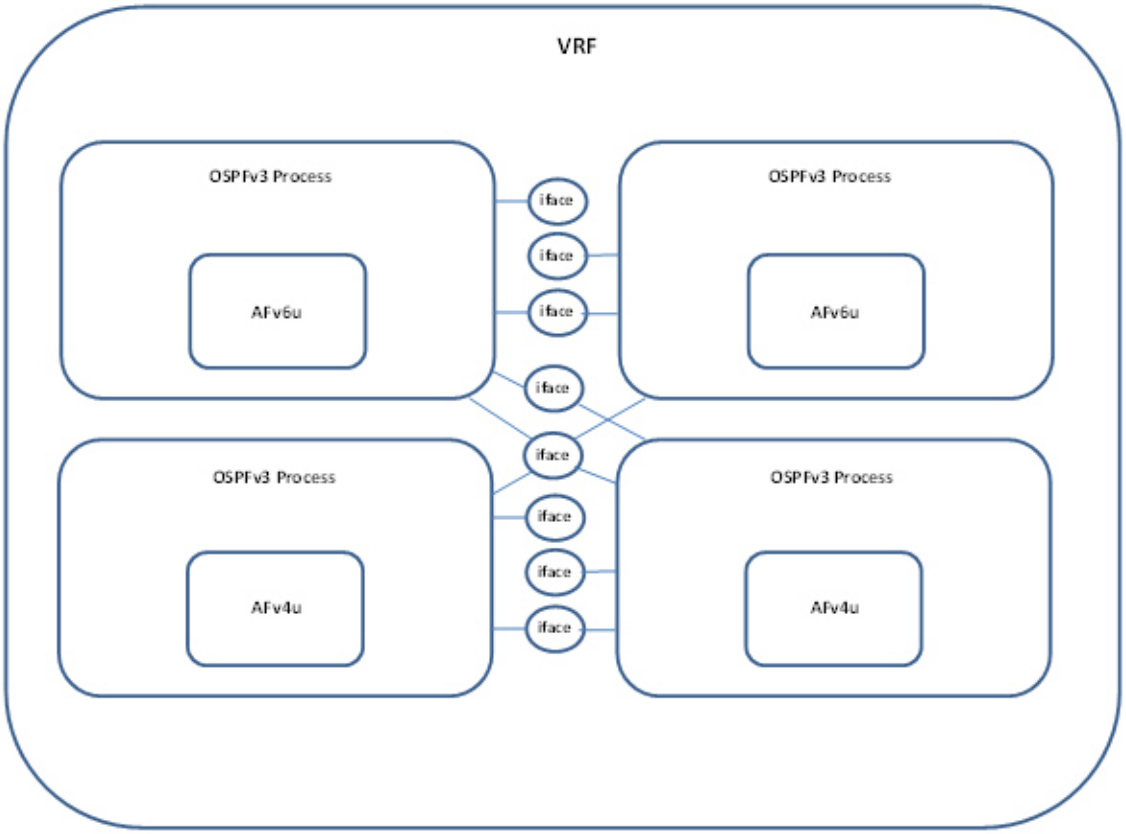

Figure 2. Relationship between VRF, OSPF Processes, Address Families, and Interfaces

OSPFv3 Redistribution

Redistribution is supported between OSPFv3 and other sources or targets that are in the same VRF instance. You can redistribute into OSPFv3 or redistribute from OSPFv3. Redistribute commands are issued at the address family level, and are used to redistribute routes from another source into OSPFv3.

Hardware and Software Requirements and Limitations

OSPFv3 is available on AOS devices using firmware R10.5.0 or later. For a complete listing of products that support OSPFv3, refer to the *Product Feature Matrix*, available online at https://supportforums.adtran.com.

In AOS firmware R10.8.0 or later, OSPFv3 supports OSPF processes that are not in the default VRF instance. In addition, multiple OSPFv3 process can be created, whether they are in the same or different VRF instances.

IPv6 is only supported in specific AOS devices using firmware 18.1.00 or later. For a complete listing of products that support IPv6, refer to the *Product Feature Matrix*, available online at https://supportforums.adtran.com.

AOS implements a dual IP stack method for supporting IPv4 and IPv6. Not all AOS features are compatible with IPv6. Refer to the *Product Feature Matrix*, available online at https://supportforums.adtran.com, for more information about the supported IPv6 features in AOS.

For more information about IPv6 and its configuration in AOS, refer to the configuration guide *Configuring IPv6 in AOS*, available online at http://supportforums.adtran.com.

Configuring OSPFv3 in AOS

OSPFv3 is configured in AOS using the CLI. The commands for OSPFv3 are separated into three main categories: global OSPFv3 settings, address family OSPFv3 settings, and interface OSPFv3 settings. The global OSPFv3 settings are used to specify the OSPFv3 process number, the VRF on which OSPFv3 is operating, and other general OSPFv3 global settings. The address family is used to configure OSPFv3 settings specific to the address family, such as the OSPFv3 area for the family, maximum paths, summary prefixes, and redistribution information. The interface OSPFv3 settings include enabling OSPFv3 on the interface, specifying packet intervals, route costs, and area options. The following sections describe the commands used in OSPFv3 configuration.

OSPFv3 Global Configuration

To configure the global parameters for OSPFv3 over IPv6 in AOS, use the commands described in this section.

1. To enable OSPFv3, create an OSPFv3 process, and enter the router's OSPFv3 Configuration mode, enter the **router ospfv3** <*process id*> **[vrf** <*name*>**]** command from the Global Configuration mode prompt. This command creates a new OSPFv3 process, or allows you to edit a previously configured OSPFv3 process. By default, OSPFv3 is disabled. The <*process id*> parameter creates a process ID for the OSPFv3 process. These IDs are globally unique across the AOS device. The ID value must be unique among all OSPFv3 processes on the device, and it does not correlate or interact with any OSPFv2 processes. Valid ID range is **1** to **65535**. By default, no process ID exists. The optional **vrf** <*name*> parameter specifies a nondefault VRF on which to enable OSPFv3. If the VRF is not specified, OSPFv3 is enabled and configured on the default (unnamed) VRF. Using the **no** form of this command removes the OSPFv3 process and all of its settings at both the global and interface level. To enable OSPFv3 and create an OSPFv3 process on the default VRF instance, enter the command as follows:

   ```
   (config)#router ospfv3 5
   (config-ospfv3)#
   ```

   

   The process ID represents an instance of an OSPFv3 process running on the local system, but the value defined for the process ID is not the same as the instance ID. The instance ID is used in the header of OSPFv3 packets transmitted on a link, and is defined and mapped to a process ID when enabling OSPFv3 at a link.

2. Specify the cost of the default summary route that is injected into the specified OSPFv3 area (when the area is a stub or NSSA) using the **area** <*area id* | *ipv4 address*> **default-cost** <*cost*> command. This command cannot be applied to the backbone area, since the backbone area cannot be a stub or NSSA. The <*area id* | *ipv4 address*> parameter specifies the ID of the area that receives the default summary route. The area ID can be specified as an integer or as an IPv4 address. Valid integer area ID range is **1** to **4294967295**. Enter IPv4 addresses in dotted decimal notation (for example, **10.10.10.1**). The <*cost*>

parameter is the advertised cost of the default summary route. Valid cost range is **0** to **16777214**. By default, the cost is set to **0**. Using the **no** form of this command returns the cost of the summary route to the default value. To specify the cost of the default summary route, enter the command from the OSPFv3 Configuration mode as follows:

```
(config-ospfv3)#area 10 default-cost 100
(config-ospfv3)#
```

3. Specify a symmetrical, bidirectional SA that uses ESP for encryption and authenticates all OSPFv3 messages that are sent and received on each interface in the specified area using the **area** <*area id | ipv4 address*> **encryption ipsec spi** <*spi*> **esp [null |** <*encryption type*> <*encryption key*>**] [**<*authentication type*> <*authentication key*>**]** command. This command allows you to specify OSPFv3 security at the area level, which eases configuration and management when the same security settings are desired at multiple interfaces. It also allows the same SPIs to be used at multiple interfaces, which is not possible when specifying OSPFv3 protection at the interface level. By default, there is no security for the OSPFv3 messages in an area. The <*area id | ipv4 address*> parameter specifies the ID of the area on which encryption and authentication is used. The area ID can be specified as an integer or as an IPv4 address. Valid integer area ID range is **1** to **4294967295**. Enter IPv4 addresses in dotted decimal notation (for example, **10.10.10.1**). The **spi** <*spi*> parameter specifies the SPI for the SA. The value specified must not be in use by another IPsec function on the system or an error message is generated. If the same SPI is already in use in the same OSPFv3 area, entering this command with the same value will overwrite the current configuration. Valid SPI range is **256** to **4294967295**. The **null** parameter specifies that OSPFv3 messages are not encrypted, but will be authenticated. The <*encryption type*> parameter specifies the type of algorithm used to encrypt OSPFv3 messages. Valid values for encryption are **3des**, **aes-cbc 128**, **aes-cbc 192**, **aes-cbc 256**, and **des**. After specifying the encryption type, specify the <*encryption key*>. The size of the encryption key is determined by the respective algorithm. The <*authentication type*> parameter specifies the algorithm for authenticating OSPFv3 message. Valid values for authentication are **md5** and **sha1**. After specifying the authentication type, specify the <*authentication key*>. The size of the encryption key is determined by the respective algorithm. Use the **no** form of this command to remove IPsec protection of OSPFv3 messages in this area. To configure encryption and authentication for OSPFv3 messages, enter the command from the OSPFv3 Configuration mode. In the following example, OSPFv3 messages in area **10** are configured with the SPI **100**, no encryption, and **md5** as an authentication method:

```
(config-ospfv3)#area 10 encryption ipsec spi 100 esp null md5 NeWtStpsswdKEY
(config-opsfv3)#
```

After this command has been entered, and the SAs are in place, if the command is reentered with different settings an attempt is made to create the new SA. If SA creation fails, an error message is displayed, the command is not accepted into the running configuration on the device, and the original command and SAs remain in place. When OSPFv3 protection is specified at multiple locations that affect the same interface, one protection set is selected to protect all OSPFv3 instances on that interface. The set is chosen using the following rules processed in this order:

- The protection configured at the interface level (including null), then everything else
- The protection configured at the IPv6 address family instance, then everything else
- The protection configured at the IPv4 address family instance, then everything else (IPv4 address family is not supported as of AOS firmware release R10.5.0)
- No protection

4. Specify the area type as a stub or total stub using the **area** <*area id | ipv4 address*> **stub [no-summary]** command from the OSPFv3 Configuration mode. When an area is specified as a stub, the area border router withholds inter-area-router-LSAs and AS-external-LSAs, and instead injects a default summary route. When an area is a total stub, the area border router also withholds inter-area-prefix-LSAs. The cost for the injected route is defined by the command **area** <*area id | ipv4 address*> **default-cost** <*cost*> (refer to *page 10*). By default, the area type is normal, and is not specified as a stub. The backbone area cannot be used as a stub. Using the **no** form of this command returns the area to the default area type. The <*area id | ipv4 address*> parameter specifies the ID of the area that you are defining. The area ID can be specified as an integer or an IPv4 address. Valid integer area ID range is **1** to **4294967295**. Enter IPv4 addresses in dotted decimal notation (for example, **10.10.10.1**). The optional **no-summary** command specifies the area is a total stub. To specify the area as a total stub, enter the command as follows:

```
(config-ospfv3)#area 10 stub no-summary
(config-ospfv3)#
```

> NOTE *All routers connected to an area must have the same configuration for the area type.*

5. Specify the reference value used to calculate the OSPFv3 cost of an interface using the **auto-cost reference-bandwidth** <*bandwidth*> command from the OSPFv3 Configuration mode. In OSPFv3, the cost of an interface is the reference value divided by the interface's bandwidth. Certain conditions require a change to the system's reference value, such as an increase in interface speeds. The default reference of 100 Mbps can become ineffective since 1 is the lowest possible cost value. This command allows the reference value to be customized for your specific network design. All OSPFv3 participants in the network should use the same reference value. The <*bandwidth*> parameter is the reference bandwidth used to calculate the OSPFv3 cost on the AOS device's OSPFv3 interfaces. Valid range is **1** to **4294967** Mbps. By default, the reference value is set to **100** Mbps. Using the **no** form of this command returns the reference to the default value. To change the default value, enter the command as follows:

```
(config-ospfv3)#auto-cost reference-bandwidth 400
(config-ospfv3)#
```

6. Specify the value to be used by the OSPFv3 process as the router ID using the **router-id** <*ipv4 address*> command. An OSPFv3 router ID is selected using a well defined algorithm that is also used for BGP router IDs. The default router ID value is chosen by first looking at the configured router ID (if there is one), then the highest value IPv4 address assigned to a loopback interface in the same VRF, and then the highest value IPv4 address assigned to a non-loopback interface. The <*ipv4 address*> parameter of this command is a 32-bit value (represented in an IPv4 address format), that is used by this specific OSPFv3 instance as the router ID. The value must be unique for the OSPF domain to which the OSPFv3 process belongs. If the AOS device detects that another router is using this router ID, and duplicate router ID detection is enabled, a message is displayed. Though the value uses an IPv4 address in general format, it does not actually use the IPv4 address. IPv4 addresses should be entered in dotted decimal notation (for example, **10.10.10.1**). Valid IPv4 address range for this command is **0.0.0.1** to **255.255.255.255**. Using the **no** form of this command returns the router ID for this OSPFv3 process to the default (which is determined as explained previously). To specify a router ID for this OSPFv3 process, enter the command as follows:

```
(config-ospfv3)#router-id 10.10.10.2
```



```
(config-ospfv3)#
```

7. Specify the OSPFv3 shortest path first (SPF) calculation and hold interval timers using the **timers spf** <*delay*> **[**<*hold time*>**]** command from the OSPFv3 Configuration mode. The <*delay*> parameter specifies the time (in seconds) between receipt of OSPFv3 topology changes and the beginning of SPF calculations. Valid range is **0** to **65535** seconds, with a default value of **5** seconds. The optional <*hold time*> parameter specifies the time (in seconds) between consecutive SPF calculations. Valid range is **0** to **65535** seconds, with a default value of **10** seconds. Using the **no** form of this command returns the SPF timers to the default value. To change the SPF timers, enter the command as follows:

   ```
   (config-ospfv3)#timers spf 3 7
   (config-ospfv3)#
   ```

8. Specify an address family for the OSPFv3 process, and enter the address family's configuration mode using the **address-family ipv6 unicast** command from the OSPFv3 Configuration mode. If the address family has been previously configured, this command allows you to edit its configuration. Once an address family is created, it is permanently associated with the parent OSPFv3 process until the process itself is removed. Therefore, there is not a **no** version of this command. By default, no address family exists. To create an address family for the OSPFv3 process, enter the command as follows:

   ```
   (config-ospfv3)#address-family ipv6 unicast
   (config-ospfv3-ipv6)#
   ```

OSPFv3 IPv6 Address Family Configuration

After creating an address family for the OSPFv3 process on the AOS device, you can configure the specific settings for the address family. IPv6 address family settings are configured using the commands described in this section.

1. Control the route summarization between OSPFv3 areas (inter-area prefixes of type 3 LSAs) using the **area** <*area id* | *ipv4 address*> **range** <*ipv6 address/prefix-length*> **[advertise | not-advertise]** command from the OSPFv3 IPv6 AF Configuration mode. Operation of this command is the same as the existing OSPFv2 command, with the exception that the range is specified as an IPv6 prefix and length, rather than an IPv4 address with subnet and mask. By default, route summarization is disabled. Using the **no** form of this command removes the specified summarization, and returns to advertisement of individual summary prefixes between areas. The <*area id* | *ipv4 address*> parameter specifies the ID of the area of which the router is a member, and that contains the prefixes being summarized. The summaries are advertised into other areas of which this router is also a member. The area ID can be specified as an integer or as an IPv4 address. Valid integer area ID range is **1** to **4294967295**. Enter IPv4 addresses in dotted decimal notation (for example, **10.10.10.1**). The <*ipv6 address/prefix-length*> parameter is the IPv6 prefix and length to be advertised. This value should be a summary of other, more specific prefixes that exist in the area. Enter IPv6 addresses in colon hexadecimal format (**X:X:X:X::X/<Z>**), for example, **2001:DB8::1/64**. The prefix length (<**Z**>) is an integer with a value between **0** and **128**. The optional **advertise** and **not-advertise** parameters specify whether the summary will or will not be advertised to other areas. The summary is advertised by default, unless otherwise specified.

   To configure the route summarization between OSPFv3 areas, enter the command as follows:

   ```
   (config-ospfv3-ipv6)#area 10 range 2001:DB8::1/64 not-advertise
   (config-ospfv3-ipv6)#
   ```


2. Specify that a default external route is injected into this OSPFv3 process using the **default-information originate [always] [metric** <*value*>**] [metric-type** <**1** | **2**>**]** from the OSPFv3 IPv6 Address Family Configuration mode. A router that injects an external default route becomes an OSPFv3 ASBR because it sources routes from outside the OSPF AS. The optional **always** parameter allows the ASBR to inject a default route into the OSPF AS. The optional **metric** <*value*> parameter assigns an OSPFv3 metric value to the route being injected into OSPFv3. If the metric is not specified, then the **default-metric** command is used for the default route metric (refer to Step 3 below). If the **default-metric** command is not configured, then the metric of **10** is used. If the metric is specified as **0** using this command, it means that the metric is set to an unconfigured value, and the **default-metric** command setting is used. Valid metric range is **0** to **16777214**, with a default value of **10**. The optional **metric-type** parameter specifies the external metric type for the route being injected into OSPFv3 by this command. Metric type **1** specifies that when external routes are assigned a metric they begin with the metric value specified by this command, and add the cost of the OSPF path as they are advertised throughout the AS. Metric type **2** is not affected by the OSPF path cost, and retains the original metric values. By default, metric type **2** is used. These parameters can be entered in any order. By default, no external route is injected into OSPFv3. Using the **no** form of this command removes the default route. To specify an external route is injected into OSPFv3, using the default metric of **10** and the default metric type of **2**, enter the command as follows:

   ```
   (config-ospfv3-ipv6)#default-information originate always
   (config-ospfv3-ipv6)#
   ```

3. Specify the metric value used for redistributed routes when the value is not otherwise specified by entering the **default-metric** <*value*> command from the OSPFv3 IPv6 Address Family Configuration mode. The metric for redistributed routes can be specified using the **redistribute** command (refer to Step 8 on *page 15*), or using a route map. When the value is not set in one of the locations, the default metric is used. This setting does not affect the metric of the default route injected using the **default-information originate** command, although this value is used if the metric is not specified in the **default-information originate** command. The <*value*> parameter specifies the OSPFv3 metric value assigned to the route being injected into OSPFv3. Valid range is **0** to **6777214**. By default, the metric is set to **20**. Using the **no** form of this command returns the default metric to the default value. To change the default metric value, enter the command as follows:

   ```
   (config-ospfv3-ipv6)#default-metric 30
   (config-ospfv3-ipv6)#
   ```

4. Specify the administrative distance for OSPFv3 routes using the **distance [**<*distance*> | **ospf [external** <*distance*> | **inter-area** <*distance*> | **intra-area** <*distance*>**]]** command in the OSPFv3 IPv6 Address Family Configuration mode. The distance can be set once for all OSPFv3 route types or individually for each route type. The optional **ospf external**, **inter-area**, and **intra-area** parameters allow you to specify the administrative distance for routes in the specified area (external, inter-area, or intra-area). The <*distance*> parameter specifies the administrative distance to be used by the route type (whether all routes or routes in a specified area). Valid distance range is **0** to **255**. By default, all routes have an administrative distance of **110**. Using the **no** form of this command returns the administrative distance to the default value.

   To change the administrative distance for all OSPFv3 routes, enter the command as follows:

   ```
   (config-ospfv3-ipv6)#distance 150
   (config-ospfv3-ipv6)#
   ```


5. Enable the ability to detect when a duplicate router ID is found in the OSPFv3 link state database using the **duplicate-routerid-detection** command in the OSPFv3 IPv6 Address Family Configuration mode. When enabled, if a received LSA contains the router ID of the OSPFv3 process to which the LSA belongs, a warning event is displayed. By default, this feature is enabled. Using the **no** form of this command disables duplicate router ID detection. To enable this feature, enter the command as follows:

   ```
   (config-ospfv3-ipv6)#duplicate-routerid-detection
   (config-ospfv3-ipv6)#
   ```

6. Specify the maximum number of equal cost routes to a given prefix that OSPFv3 can enter into the route table using the **maximum-paths** <*value*> command in the OSPFv3 IPv6 Address Family Configuration mode. The <*value*> parameter specifies the maximum number of equal cost routes. Valid range is **1** to **6**, with a default value of **4**. Using the **no** form of this command returns the maximum paths to the default value. To change the maximum number of equal cost routes, enter the command as follows:

   ```
   (config-ospfv3-ipv6)#maximum-paths 5
   (config-ospfv3-ipv6)#
   ```

7. Control route summarization and route advertisement that is redistributed into this OSPFv3 process (external prefixes of type 5 LSAs) using the **summary-prefix** <*ipv6 address/prefix-length*> **[not-advertise]** command in the OSPFv3 IPv6 Address Family Configuration mode. By default, the cost applied to the summary route is that of the lowest cost route in the set it summarizes. The command can be entered multiple times to summarize different prefixes. Individual routes being summarized, as well as the resulting summary prefix, bypass route maps specified in **redistribute** commands (below). In addition, if multiple summaries exist where one summary subsumes another, the prefix with the shortest length is used. The <*ipv6 address/prefix-length*> parameter specifies the IPv6 prefix and length that is advertised. The value should be a summary of other, more specific prefixes that are being redistributed. Enter IPv6 addresses in colon hexadecimal format (**X:X:X:X::X/<Z>**), for example, **2001:DB8::1/64**. The prefix length (<**Z**>) is an integer with a value between **0** and **128**. The optional **not-advertise** parameter specifies that the summary will not be advertised to other OSPFv3 areas. Use the **no** form of this command to remove the specified route summary, and return to advertising the individual prefixes that are being redistributed. To create a route summary, enter the command as follows:

   ```
   (config-ospfv3-ipv6)#summary-prefix 2001:DB8::1/64
   (config-ospfv3-ipv6)#
   ```

8. To redistribute routes from a specified source into the OSPFv3 process, enter the **redistribute [bgp | connected | static | ospf** <*process id*> **[include-connected]] [metric** <*value*>**] [metric-type [1 | 2]] [route-map** <*name*>**]** command in the OSPFv3 IPv6 Address Family Configuration mode. A router that performs redistribution becomes an OSPFv3 ASBR, because it is sourcing routes from outside the OSPFv3 AS. The specified source from which routes are redistributed must be in the same VRF and of the same AF as the OSPFv3 instance being configured. This command can be entered multiple times (at most, once for each source). Reentering the command with the same source replaces any existing command with that source. The **bgp**, **connected**, and **static** parameters define the source from which routes are redistributed. The **ospf** <*process id*> parameter specifies the OSPF process from which to redistribute routes. The optional **include-connected** parameter allows you to redistribute the corresponding connected routes. When this parameter is used, all connected routes on interfaces participating in the OSPF process are also redistributed. The optional **metric** <*value*> parameter assigns an OSPFv3 metric value to the route being redistributed into OSPFv3. If the metric is not specified, then the **default-metric** command is used for the default route metric (refer to Step 3 on

page 14). If the **default-metric** command is not configured, then the metric of **10** is used. If the metric is specified as **0** using this command, it means that the metric is set to an unconfigured value, and the **default-metric** command setting is used. Valid metric range is **0** to **16777214**, with a default value set by the **default-metric** command. The optional **metric-type** parameter specifies the external metric type for the route being redistributed into OSPFv3 by this command. Metric type **1** specifies that when external routes are assigned a metric they begin with the metric value specified by this command, and add the cost of the OSPF path as they are advertised throughout the AS. Metric type **2** is not affected by the OSPF path cost, and retains the original metric values. By default, metric type **2** is used. The optional **route-map** parameter specifies a route map that is applied to routes being redistributed by this command. A route map can impose granular control on routes being redistributed. These parameters can be entered in any order.

When using redistribution, keep in mind the following:

- OSPFv3 settings, such as metric and metric types, that are specified in a route map entry override those settings in the **redistribution** command for the routes that match the route map entry.
- When using the **summary-prefix** command (refer to Step 7 on *page 15*), individual routes that are summarized, as well as the resulting summary, bypass route maps in **redistribute** commands.
- By default, no route map is used.

By default, no routes are redistributed into the OSPFv3 process. Using the **no** form of this command removes the redistribution from the specified source. To redistribute routes from another source, enter the command as follows:

```
(config-ospfv3-ipv6)#redistribute bgp
(config-ospfv3-ipv6)#
```

To redistribute a prefix (**1234::/24**) in the OSPFv3 process **2**, along with all the routes learned from OSPFv3 process **1**, enter the command as follows:

```
(config)#interface ethernet 0/1
(config-eth 0/1)#ipv6 address 1234::1/24
(config-eth 0/1)#ospfv3 1 area 0 ipv6
(config-eth 0/1)#exit
(config)#router ospfv3 2
(config-ospfv3)#address family ipv6 unicast
(config-ospfv3-ipv6)#redistribute ospf 1 include-connected
```

If the **include-connected** parameter is not used, only the routes learned through the OSPF process **1** are redistributed and the **1234::/24** would not be redistributed.

OSPFv3 Interface Configuration

After configuring the global OSPFv3 settings, and any OSPFv3 IPv6 Address Family settings, you must configure the OSPFv3 interface. When configuring the interface, you will enable OSPFv3 on the interface, associate the interface with an OSPFv3 process, specify an OSPFv3 cost for the interface, and configure other OSPFv3 settings that are specific to the interface. Each OSPFv3 interface configuration requires a process ID, which allows you to match the OSPFv3 settings you are configuring with an OSPFv3 process that is defined at the global level (refer to *OSPFv3 Global Configuration on page 10*). OSPFv3 interface configuration commands are described in the following section.

1. Specify a symmetrical, bidirectional SA that uses ESP for encryption and authenticates all OSPFv3 messages that are sent and received on the interface using the **ospfv3 encryption [null | ipsec spi** <*spi*> **esp [null |** <*encryption type*> <*encryption key*>**] [**<*authentication type*> <*authentication key*>**]]** command. This command allows you to specify OSPFv3 security at the interface level. By default, there is no security for the OSPFv3 messages on an interface. Protection specified with this command overrides any area-level OSPFv3 protection that might apply to the interface. The **spi** <*spi*> parameter specifies the SPI for the SA. The value specified must not be in use by another IPsec function on the system, or an error message is generated. If the same SPI is already in use in the same OSPFv3 area, entering this command with the same value will overwrite the current configuration. Valid SPI range is **256** to **4294967295**. The **null** parameter specifies that OSPFv3 messages on this interface are not encrypted when used in the **ospfv3 encryption null** format (even when specified by OSPFv3 area configuration). When used in the **ospfv3 encryption ipsec spi** <*spi*> **esp null** format, **null** indicates that messages on the interface will not by encrypted, but will be authenticated. The <*encryption type*> parameter specifies the type of algorithm used to encrypt OSPFv3 messages. Valid values for encryption are **3des**, **aes-cbc 128**, **aes-cbc 192**, **aes-cbc 256**, and **des**. After specifying the encryption type, specify the <*encryption key*>. The size of the encryption key is determined by the respective algorithm. The <*authentication type*> parameter specifies the algorithm for authenticating OSPFv3 message. Valid values for authentication are **md5** and **sha1**. After specifying the authentication type, specify the <*authentication key*>. The size of the encryption key is determined by the respective algorithm. Use the **no** form of this command to remove IPsec protection of OSPFv3 messages on this interface. To configure encryption and authentication for OSPFv3 messages on the interface, enter the command from the interface's configuration mode. In the following example, OSPFv3 messages on the Ethernet interface are configured with the SPI **120**, no encryption, and **md5** as an authentication method:

```
(config-eth 0/1)#ospfv3 encryption ipsec spi 120 esp null md5 NeWtStpsswdKEY
(config-eth 0/1)#
```

To apply protection to OSPFv3 messages on an interface with settings that are different than those configured at the OSPFv3 area, protection must be defined at the interface and with a unique SPI.

2. Add an interface to an OSPFv3 process using the **ospfv3** <*process id*> **area** <*area id*> **ipv6 [instance** <*instance*>**]** command from the interface's configuration mode. This command places the interface in the specified area for the specified address family (IPv6), and optionally defines the instance ID that is used to represent this OSPFv3 process in messages on the interface's link. The <*process id*> parameter specifies the OSPFv3 routing process this interface is to join, for the specified address family. The process ID is locally significant to the device. This command can also be used to create the OSPFv3 process instance (rather than with the global configuration command), but the value must be unique among all OSPFv3 processes on the device. Valid process ID range is **1** to **65535**. By default, a process ID is not defined. The **area** <*area id*> parameter specifies the ID of the area to which this interface is assigned for the given OSPFv3 process. Valid range is **0** to **4294967295**. The **ipv6** parameter identifies the OSPFv3 address family. The optional **instance** <*instance*> parameter specifies the value to use in the instance ID field of messages sent or received by this OSPFv3 process on the interface's link. Valid range is **0** to **31**. By default, the interface is not configured to be part of an OSPFv3 process. Using the **no** form of this command removes the interface from the OSPFv3 process.

The following rules apply to this command:

- The interface must have the address family enabled on the interface. This means that either the **ipv6** or a form of the **ipv6 address** command must be configured on the interface. If the address family is not enabled on the interface, the command is rejected, and an error is displayed.
- The interface and the specified OSPFv3 process (if defined in the global configuration) must be in the same VRF or the command will fail.
- The address family must match that specified for the OSPFv3 process if the process has been defined in the global configuration or the command will fail.
- If the OSPFv3 process identified by the process ID does not exist in the global configuration, it is automatically created, along with the specified address family, and it is assigned to the VRF of which the interface is a member.
- If the specified OSPFv3 process is already at its maximum limit of processes or address families, the command fails.
- If the specified OSPFv3 process already exists in the global configuration, but its configuration does not include an address family, the specified address family is added to the OSPFv3 router configuration.
- A given OSPFv3 process can only have one address family.
- Multiple OSPFv3 instances per address family, per VRF, can be created and can be assigned to a given interface. However, each OSPFv3 process on an interface must have a different instance ID.
- If the interface's VRF changes, all OSPFv3 assignments are removed.
- To change an OSPFv3 process's VRF, the process must first be removed and then recreated. Removing the process removes all OSPFv3 assignments for that process from all interfaces.

To add an interface to the OSPFv3 process, enter the interface's configuration mode and enter the command as follows:

```
(config-eth 0/1)#ospfv3 5 area 10 ipv6 instance 10
(config-eth 0/1)#
```

3. Specify the OSPFv3 cost of the interface using the **ospfv3** <*process id*> **cost** <*cost*> command from the interface's configuration mode. This command specifies a value that represents the cost of sending a packet over the interface. By default, the cost of the interface is automatically computed. The automatic cost computation formula is the reference bandwidth divided by the interface bandwidth. The reference bandwidth is set by the **auto-cost reference-bandwidth** command (refer to Step 5 on *page 12*) and defaults to 100 Mbps. The <*process id*> parameter specifies the OSPFv3 routing process this interface is to join. Valid process ID range is **1** to **65535**. The <*cost*> parameter is the OSPFv3 cost of the interface. This value overrides any computed cost value. Valid range is **1** to **65535**. By default, the cost is automatically computed. Use the **no** form of this command to return the cost to the default value. To change the OSPFv3 cost of the interface, enter the command as follows:

   ```
   (config-eth 0/1)#ospfv3 5 cost 10
   (config-eth 0/1)#
   ```

4. Specify the maximum interval allowed between OSPFv3 Hello packets on this interface by entering the **ospfv3** <*process id*> **dead-interval** <*value*> command from the interface's configuration mode. If the maximum interval is exceeded, the device determines that the neighbor is down. This value must be the same across all interfaces on a link. The <*process id*> parameter specifies the OSPFv3 routing process this interface is to join. Valid process ID range is **1** to **65535**. If the process ID has not already been created, entering this command will not create the process. The <*value*> parameter is the maximum

number of seconds allowed between OSPFv3 Hello packets. It is recommended that this value be **4** times the Hello packet interval value. Valid range is **1** to **65535** seconds, with a default value of **40** seconds. Using the **no** form of this command returns to the dead interval to the default value. To change the dead interval on the interface, enter the command as follows:

```
(config-eth 0/1)#ospfv3 5 dead-interval 100
(config-eth 0/1)#
```

5. Specify the interval between OSPFv3 Hello packets sent on this interface by entering the **ospfv3** <*process id*> **hello-interval** <*value*> command from the interface's configuration mode. This value must be the same across all interfaces on the link. The <*process id*> parameter specifies the OSPFv3 routing process this interface is to join. Valid process ID range is **1** to **65535**. If the process ID has not already been created, entering this command will not create the process. The <*value*> parameter is the number of seconds allowed between OSPFv3 Hello packets. Valid range is **1** to **65535** seconds, with a default value of **10** seconds. Using the **no** form of this command returns to the Hello packet interval to the default value. To change the Hello packet interval on the interface, enter the command as follows:

```
(config-eth 0/1)#ospfv3 5 hello-interval 20
(config-eth 0/1)#
```

6. Specify the OSPFv3 network type for the interface by entering the **ospfv3** <*process id*> **network [broadcast | point-to-point]** command from the interface's configuration mode. The <*process id*> parameter specifies the OSPFv3 routing process this interface is to join. Valid process ID range is **1** to **65535**. If the process ID has not already been created, entering this command will not create the process. The **broadcast** parameter sets the OSPFv3 network type to broadcast, and the **point-to-point** parameter sets the OSPFv3 network type to point-to-point. By default, Ethernet interfaces are set to broadcast, and PPP, Frame Relay, and loopback interfaces are set to point-to-point. Using the **no** form of this command returns the interface's network type to the default value. To change the interface's network type, enter the command as follows:

```
(config-eth 0/1)#ospfv3 5 network point-to-point
(config-eth 0/1)#
```

7. Specify the interface OSPFv3 priority by entering the **ospfv3** <*process id*> **priority** <*value*> command from the interface's configuration mode. Priority is used in the election of the designated router and the designated backup router on multi-access networks. Interfaces connected to multi-access networks (such as Ethernet interfaces) perform an election for a designated and backup designated router. The router interface with the highest OSPFv3 priority on the link becomes the designated router for that link. The interface with the next highest priority becomes the designated backup router. In the event there is a tie, the router interface with the highest router ID takes precedence. A value of **0** indicates the router is ineligible to become either the designated or backup designated router. The <*process id*> parameter specifies the OSPFv3 routing process this interface is to join. Valid process ID range is **1** to **65535**. If the process ID has not already been created, entering this command will not create the process. The <*value*> parameter specifies the OSPFv3 priority of the interface. Valid range is **0** to **255**, with a default value of **1**. Using the **no** form of this command returns the interface's priority to the default value. To change the interface's priority, enter the command as follows:

```
(config-eth 0/1)#ospfv3 5 priority 6
(config-eth 0/1)#
```

8. Specify the interval between unacknowledged OSPFv3 LSAs sent on the interface by entering the **ospfv3** <*process id*> **retransmit-interval** <*value*> command from the interface's configuration mode. The <*process id*> parameter specifies the OSPFv3 routing process this interface is to join. Valid process

ID range is **1** to **65535**. If the process ID has not already been created, entering this command will not create the process. The <*value*> parameter specifies the number of seconds between OSPFv3 LSAs sent on the interface. Valid range is **1** to **65535** seconds, with a default value of **5** seconds. Using the **no** form of this command returns the LSA interval to the default value. To change the interface's LSA interval, enter the command as follows:

```
(config-eth 0/1)#ospfv3 5 retransmit-interval 10
(config-eth 0/1)#
```

9. Specify the estimated time that is required to propagate an LSA on the interface by entering the **ospfv3** <*process id*> **transmit-delay** <*value*> command from the interface's configuration mode. The <*process id*> parameter specifies the OSPFv3 routing process this interface is to join. Valid process ID range is **1** to **65535**. If the process ID has not already been created, entering this command will not create the process. The <*value*> parameter specifies the number of seconds required to send LSAs from the interface. Valid range is **1** to **65535** seconds, with a default value of **1** second. Using the **no** form of this command returns the transmit delay to the default value. To change the interface's transmit delay, enter the command as follows:

```
(config-eth 0/1)#ospfv3 5 transmit-delay 2
(config-eth 0/1)#
```

10. Disable an OSPFv3 process on the interface by entering the **ospfv3** <*process id*> **shutdown** command from the interface's configuration mode. When this command is used, the OSPFv3 commands remain in place, but logically it appears to the interface as though the OSPFv3 process was removed from the configuration. This command is beneficial when troubleshooting. The <*process id*> parameter specifies the OSPFv3 routing process to shut down. Valid process ID range is **1** to **65535**. If the process ID has not already been created, entering this command will not create the process. Using the **no** form of this command reinstates the OSPFv3 process on the interface. To shutdown an OSPFv3 process on an interface, enter the command as follows:

```
(config-eth 0/1)#ospfv3 5 shutdown
(config-eth 0/1)#
```


OSPFv3 CLI Configuration Example

The following configuration example shows the relevant configuration necessary to configure OSPFv3 with an IPv6 address family, area backbone router (ABR), and intra-area roles. This example should be used for illustrative purposes only. You will need to make configuration changes to ensure the configuration will function in your network. The network topology for this example is shown in *Figure 3*.

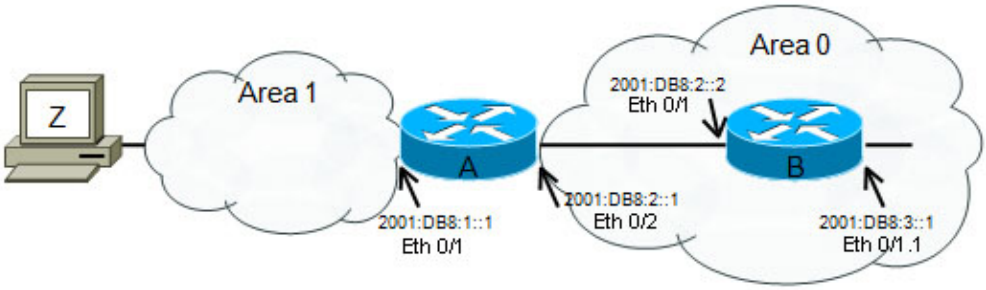

Figure 3. OSPFv3 with IPv6 Address Family, ABR, and Intra-Area Roles

The following is the sample configuration for both Router A and Router B:

Router A Configuration

```
router ospfv3 61
    router-id 202.202.202.61
    area 0 encryption ipsec spi 1000 esp des 1234567890123456 sha1
    12345678901234567890123456789012345678901234567890
    address-family ipv6 unicast
!
interface Eth 0/1
    ipv6 address 2001:DB8:1::1/64
    ospfv3 61 area 1 ipv6
!
interface Eth 0/2
    ipv6 address 2001:DB8:2::1/64
    ospfv3 61 area 0 ipv6
!
```

Router B Configuration

```
!
router ospfv3 61
    router-id 204.204.204.61
    area 0 encryption ipsec spi 1000 esp des 1234567890123456 sha1
    12345678901234567890123456789012345678901234567890
    address-family ipv6 unicast
!
!
interface Eth 0/1
    ipv6 address 2001:DB8:2::2/64
```



```
    ospfv3 61 area 0 ipv6
!
interface Eth 0/1.1
    ipv6 address 2001:DB8:3::1/64
    ospfv3 61 area 0 ipv6
!
```

OSPFv3 Command Summary

The following tables summarize the commands used to configure OSPFv3 on an AOS product.

Table 4. OSPFv3 Global Configuration Commands

Prompt	Command	Description
(config)#	**[no] router ospfv3** <*process id*> **[vrf** <*name*>**]**	Enables OSPFv3, creates an OSPFv3 process, and enters the OSPFv3 Configuration mode. OSPFv3 is disabled by default. Valid ID range is **1** to **65535**.
(config-ospfv3)#	**[no] area** <*area id* | *ipv4 address*> **default-cost** <*cost*>	Specifies the cost of the default summary route injected into the specified OSPFv3 area. Area IDs can be an integer, with a range of **1** to **4294967295**, or an IPv4 address. IPv4 addresses are entered in dotted decimal notation (**10.10.10.1**). Valid cost range is **0** to **16777214**. Default cost value is **0**.
(config-ospfv3)#	**[no] area** <*area id* | *ipv4 address*> **encryption ipsec spi** <*spi*> **esp [null |** <*encryption type*> <*encryption key*>**] [**<*authentication type*> <*authentication key*>**]**	Specifies a symmetrical, bidirectional SA that uses ESP for encryption and authenticates all OSPFv3 messages sent and received on each interface in the specified area. Area IDs can be an integer, with a range of **1** to **4294967295**, or an IPv4 address. IPv4 addresses are entered in dotted decimal notation (**10.10.10.1**). Valid **spi** range is **256** to **4294967295**. The **null** keyword specifies that OSPFv3 messages are authenticated, but not encrypted. Valid encryption values are **3des**, **aes-cbc 128**, **aes-cbc 192**, **aes-cbc 256**, and **des**. Valid authentication values are **md5** and **sha1**.
(config-ospfv3)#	**[no] area** <*area id* | *ipv4 address*> **stub [no-summary]**	Specifies the area type as a stub or total stub. By default, area types are normal and are not specified as a stub. Area IDs can be an integer, with a range of **1** to **4294967295**, or an IPv4 address. IPv4 addresses are entered in dotted decimal notation (**10.10.10.1**). The **no-summary** parameter specifies the area as a total stub.

Table 4. OSPFv3 Global Configuration Commands *(Continued)*

Prompt	Command	Description
(config-ospfv3)#	**[no] auto-cost reference-bandwidth** <*bandwidth*>	Specifies the reference value used to calculate the OSPFv3 cost of an interface. Default value is **100 Mbps**. Valid <*bandwidth*> range is **1** to **4294967** Mbps.
(config-ospfv3)#	**[no] router-id** <*ipv4 address*>	Specifies the value to be used by the OSPFv3 process as the router ID. The <*ipv4 address*> parameter is a 32-bit value represented in an IPv4 address format used by the specific OSPFv3 instance as the router ID. IPv4 addresses should be entered in dotted decimal notation. Valid IPv4 address range is **0.0.0.1** to **255.255.255.255**.
(config-ospfv3)#	**[no] timers spf** <*delay*> **[**<*hold time*>**]**	Specifies the OSPFv3 SPF calculation and hold interval timers. The <*delay*> parameter is the time between receipt of OSPFv3 topology changes and the beginning of SPF calculations. Valid range is **0** to **65535** seconds with a default value of **5** seconds. The optional <*hold time*> parameter specifies the time between consecutive SPF calculations. Valid range is **0** to **65535** seconds with a default value of **10** seconds.
(config-ospfv3)#	**address-family ipv6 unicast**	Specifies an address family for the OSPFv3 process and enters the address family's configuration mode. By default, no address family exists.

Table 5. OSPFv3 IPv6 Address Family Configuration Commands

Prompt	Command	Description
(config-ospfv3-ipv6)#	**[no] area** *<area id | ipv4 address>* **range** *<ipv6 address/prefix-length>* **[advertise | not-advertise]**	Controls the route summarization between OSPFv3 areas. The *<area id | ipv4 address>* parameter specifies the ID of the area of which the router is a member. Area IDs can be an integer, with a range of **1** to **4294967295**, or an IPv4 address. IPv4 addresses are entered in dotted decimal notation (**10.10.10.1**). The *<ipv6 address/prefix-length>* parameter is the IPv6 prefix and length to be advertised. Enter IPv6 addresses in colon hexadecimal format (**X:X:X:X::X/<Z>**), for example, **2001:DB8::1/64**. The prefix length (<**Z**>) is an integer with a value between **0** and **128**. The optional **advertise** and **not-advertise** parameters specify whether the summary will or will not be advertised to other areas. Summaries are advertised by default.
(config-ospfv3-ipv6)#	**[no] default-information originate [always] [metric** <*value*>**] [metric-type** <**1 | 2**>**]**	Specifies that a default external route is injected into this OSPFv3 process. The optional **always** parameter allows the ASBR to inject a default route, the optional **metric** <*value*> parameter assigns an OSPFv3 metric value to the route being injected. A metric of **10** is used if the **default-metric** command has not been used. Valid metric range is **0** to **16777214**. The optional **metric-type** parameter specifies an external metric type for the route being injected. Metric type **1** specifies that routes are assigned a metric beginning with the metric value specified by this command and the cost of the OSPF path is added as they are advertised throughout the AS. Metric type **2** is not affected by the OSPF path cost and retains the original metric value. By default, metric type **2** is used.
(config-ospfv3-ipv6)#	**[no] default-metric** <*value*>	Specifies the metric value used for redistributed routes when the value is not otherwise specified. The <*value*> parameter specifies the OSPFv3 metric value assigned to the route being injected into OSPFv3. Valid range is **0** to **6777214**. By default, the metric is set to **20**.

Table 5. OSPFv3 IPv6 Address Family Configuration Commands *(Continued)*

Prompt	Command	Description
(config-ospfv3-ipv6)#	**[no] distance [**<*distance*> | **ospf [external** <*distance*> | **inter-area** <*distance*> | **intra-area** <*distance*>**]]**	Specifies the administrative distance for OSPFv3 routes. The optional **ospf external**, **inter-area**, and **intra-area** parameters allow you to specify the administrative distance for routes in the specified area. The <*distance*> parameter specifies the administrative distance to be used by the route type. Valid distance range is **0** to **255**. By default, all routes have an administrative distance of **110**.
(config-ospfv3-ipv6)#	**[no] duplicate-routerid-detection**	Enables the ability to detect when a duplicate router ID is found in the OSPFv3 link-state database.
(config-ospfv3-ipv6)#	**[no] maximum-paths** <*value*>	Specifies the maximum number of equal cost routes to a given prefix that OSPFv3 can enter into the route table. Valid range is **1** to **6**, with a default value of **4**.
(config-ospfv3-ipv6)#	**[no] summary-prefix** <*ipv6 address/prefix-length*> **[not-advertise]**	Controls route summarization and router advertisement that is redistributed into this OSPFv3 process. By default, the cost applied to the summary route is that of the lowest cost route it summarizes. The <*ipv6 address/prefix-length*> parameter specifies the IPv6 prefix and length that is advertised. Enter IPv6 addresses in colon hexadecimal format (**X:X:X:X::X/<Z>**), for example, **2001:DB8::1/64**. The prefix length (<**Z**>) is an integer with a value between **0** and **128**. The optional **not-advertise** parameter specifies that the summary will not be advertised to other OSPFv3 areas.

Table 5. OSPFv3 IPv6 Address Family Configuration Commands *(Continued)*

Prompt	Command	Description
(config-ospfv3-ipv6)#	**[no] redistribute [bgp | connected | static | ospf** <*process id*> **[include-connected]] [metric** <*value*>**] [metric-type [1 | 2]] [route-map** <*name*>**]**	Redistributes routes from a specified source into the OSPFv3 process. The **bgp**, **connected**, **static**, and **ospf** parameters define the source from which routes are redistributed. If routes are redistributed from an OSPF process, the optional **include-connected** parameter allows the redistribution of corresponding connected routes. The optional **metric** parameter assigns an OSPFv3 metric value to the route. By default, the metric is **10**. Valid metric range is **0** to **16777214**. The optional **metric-type** parameter specifies the external metric type for the route. The optional **route-map** parameter specifies a route map to be applied to routes being redistributed by this command.

Table 6. OSPFv3 Interface Configuration Commands

Prompt	Command	Description
(config-eth 0/1)#	**[no] ospfv3 encryption encryption [null | ipsec spi** <*spi*> **esp [null |** <*encryption type*> <*encryption key*>**] [**<*authentication type*> <*authentication key*>**]]**	Specifies a symmetrical, bidirectional SA that uses ESP for encryption and authenticates all OSPFv3 messages that are sent and received on the interface. The **ipsec spi** <*spi*> parameter specifies the SPI for the SA. Valid range is **256** to **4294967295**. The **null** parameter specifies that OSPFv3 messages are note encrypted. Valid encryption values are **3des**, **aes-cbc 128**, **aes-cbc 192**, **aes-cbc 256**, and **des**. Valid authentication values are **md5** and **sha1**
(config-eth 0/1)#	**[no] ospfv3** <*process id*> **area** <*area id*> **ipv6 [instance** <*instance*>**]**	Adds an interface to an OSPFv3 process. The <*process id*> parameter is the OSPFv3 routing process that this interface is to join. Valid range is **1** to **65535**. By default, a process ID is not defined. The **area** <*area id*> parameter specifies the ID of the area to which this interface is assigned. Valid range is **0** to **4294967295**. The **ipv6** parameter identifies this OSPFv3 address family. The optional **instance** <*instance*> parameter specifies the value to use in the instance ID field of messages sent or received by this OSPFv3 process on the interface's link. Valid range is **0** to **31**.

Table 6. OSPFv3 Interface Configuration Commands *(Continued)*

Prompt	Command	Description
(config-eth 0/1)#	**[no] ospfv3** <*process id*> **cost** <*cost*>	Specifies the OSPFv3 cost of the interface. By default, the cost of the interface is automatically computed. The <*process id*> parameter specifies the OSPFv3 process that this interface is to join. Valid process ID range is **1** to **65535**. The <*cost*> parameter is the OSPFv3 cost of the interface. This value overrides any computed cost value. Valid range is **1** to **65535**.
(config-eth 0/1)#	**[no] ospfv3** <*process id*> **dead-interval** <*value*>	Specifies the maximum interval allowed between OSPFv3 Hello packets on this interface. This value must be the same across all interfaces on a link. The <*process id*> parameter specifies the OSPFv3 process that this interface is to join. Valid process ID range is **1** to **65535**. The <*value*> parameter is the maximum number of seconds allowed between Hello packets. It is recommended that this value be 4 times the Hello packet interval value. Valid range is **1** to **65535** seconds, with a default value of **40** seconds.
(config-eth 0/1)#	**[no] ospfv3** <*process id*> **hello-interval** <*value*>	Specifies the interval between OSPFv3 Hello packets sent on this interface. The <*process id*> parameter specifies the OSPFv3 process that this interface is to join. Valid process ID range is **1** to **65535**. The <*value*> parameter is the number of seconds allowed between OSPFv3 Hello packets. Valid range is **1** to **65535**, with a default value of **10** seconds.
(config-eth 0/1)#	**[no] osfpv3** <*process id*> **network [broadcast | point-to-point]**	Specifies the OSPFv3 network type for the interface. The <*process id*> parameter specifies the OSPFv3 process that this interface is to join. Valid process ID range is **1** to **65535**. The **broadcast** parameter sets the OSPFv3 network type to broadcast, and the **point-to-point** parameter sets the network type as point-to-point. By default, Ethernet interfaces are set to broadcast, and PPP, Frame Relay, and loopback interfaces are set to point-to-point.

Table 6. OSPFv3 Interface Configuration Commands *(Continued)*

Prompt	Command	Description
(config-eth 0/1)#	**[no] ospfv3** <*process id*> **priority** <*value*>	Specifies the interface OSPFv3 priority. The <*process id*> parameter specifies the OSPFv3 process that this interface is to join. Valid process ID range is **1** to **65535**. The <*value*> parameter specifies the OSPFv3 priority. Valid range is **0** to **255**, with a default value of **1**. A value of **0** indicates the router is ineligible to become either the designated or backup designated router.
(config-eth 0/1)#	**[no] ospfv3** <*process id*> **retransmit-interval** <*value*>	Specifies the interval between OSPFv3 LSAs sent on the interface. The <*process id*> parameter specifies the OSPFv3 process that this interface is to join. Valid process ID range is **1** to **65535**. The <*value*> parameter specifies the number of seconds between LSAs. Valid range is **1** to **65535** seconds, with a default value of **5** seconds.
(config-eth 0/1)#	**[no] ospfv3** <*process id*> **transmit-delay** <*value*>	Specifies the estimated time required to propagate an LSA on the interface. The <*process id*> parameter specifies the OSPFv3 process that this interface is to join. Valid process ID range is **1** to **65535**. The <*value*> parameter specifies the number of seconds required to send LSAs from the interface. Valid range is **1** to **65535**, with a default value of **1** second.
(config-eth 0/1)#	**[no] ospfv3** <*process id*> **shutdown**	Disables an OSPFv3 process on the interface. The <*process id*> parameter specifies the OSPFv3 process to shut down. Valid process ID range is **1** to **65535**.

Troubleshooting OSPFv3

Various commands can be used from the Enable mode to view OSPFv3 configurations, clear certain OSPFv3 settings, and create debug messages for OSPFv3 events. These commands are included in the following sections.

OSPFv3 Show Commands

Show commands can be used to view current OSPFv3 configurations, including OSPFv3 processes, link state database information, and router interfaces configured for OSPFv3. The following section describes the **show** commands available for OSPFv3 troubleshooting.

Use the **show ospfv3 [**<*process id*>**]** command to display general information regarding OSPFv3 processes. You can limit the output of this command to a single OSPFv3 process by entering the process ID. The following is sample output from this command:

```
>enable
#show ospfv3
  Summary of OSPFv3 Process 61 with ID: 5.5.5.5, VRF RED
  Supports IPv6 Address Family
  SPF delay timer: 5 seconds, Hold time between SPFs: 10 seconds
  LSA interval: 1800 seconds
  Number of external LSAs: 4, Checksum Sum: 0x22a04
  Number of AS scoped unknown LSAs: 0, Checksum Sum: 0x0
  Number of areas: 2, normal: 2, stub: 0, NSSA: 0
  Reference bandwidth unit is 100 Mbps
    Area (0) 5.5.5.5
      Number of interfaces in this area: 2
      Authentication type: 0
      SPF algorithm execution count: 2
      Number of LSAs: 8, Checksum Sum: 0x3f91a
    Area (1) 5.5.5.5
      Number of interfaces in this area: 1
      Authentication type: 0
      SPF algorithm execution count: 3
      Number of LSAs: 6, Checksum Sum: 0x39601
```

Use the **show ospfv3 [**<*process id*> **[**<*area id*>**]] database [adv-router** <*router id*>**]** command to display information contained in the OSPFv3 link state database. You can limit the output of this command to a specified OSPFv3 process by entering the process ID, a specified OSPFv3 area by entering the area ID, or a specified advertising router by entering the router ID. The following is sample output from this command:

```
>enable
#show ospfv3 database
OSPFv3 router with ID: 4.4.4.4 (Process ID 61, VRF RED)

  Router Link States, Area 0
   Adv Router      Age      Seq #       Fragment ID  Link count   Bits
    4.4.4.4        222      0x8000005A  0            1            B, E
    5.5.5.5        215      0x80000066  0            1            B, E

  Network Link States, Area 0
   Adv Router      Age      Seq #       Link ID     Rtr count
    4.4.4.4        225      0x80000001  8           2

  Inter Area Prefix Link States, Area 0
   Adv Router      Age      Seq #       Prefix
    4.4.4.4        595      0x80000057  2001:10:24:204::/64
    5.5.5.5        220      0x80000001  2001:10:24:205::/64

  Link (Type-8) Link States, Area 0
```



```
 Adv Router      Age      Seq#         Link ID    Interface
  4.4.4.4        595      0x80000057   8          eth 0/1.106
  5.5.5.5        220      0x80000062   13         eth 0/1.106

 Intra Area Prefix Link States, Area 0
 Adv Router      Age      Seq #        Link ID    Ref-lstype  Ref-LSID
  4.4.4.4        225      0x80000001   8192       0x2002      8

 Router Link States, Area 1
 Adv Router      Age      Seq #        Fragment ID  Link count   Bits
  5.5.5.5        183      0x80000003   0            0            B, E

 Inter Area Prefix Link States, Area 1
 Adv Router      Age      Seq #        Prefix
  5.5.5.5        220      0x80000001   2001:10:24:106::/64
  5.5.5.5        211      0x80000001   2001:10:24:204::/64

 Inter Area Router Link States, Area 1
 Adv Router      Age      Seq #        Ref-router
  5.5.5.5        211      0x80000001   4.4.4.4

 Link (Type-8) Link States, Area 1
 Adv Router      Age      Seq#         Link ID    Interface
  5.5.5.5        223      0x80000001   14         eth 0/2.1

 Intra Area Prefix Link States, Area 1
 Adv Router      Age      Seq #        Link ID    Ref-lstype  Ref-LSID
  5.5.5.5        183      0x80000003   0          0x2001      0

 External Link States
 Adv Router      Age      Seq #        Prefix
  4.4.4.4        595      0x80000057   2001:7:1::/64
  4.4.4.4        595      0x80000057   2001:10:24:202::/64
  5.5.5.5        223      0x80000001   2001:8:1::/64
  5.5.5.5        223      0x80000001   2001:8:2::/64
```

Use the **show ospfv3 [**<*process id*> **[**<*area id*>**]] database database-summary** command to display a summary of information about the OSPFv3 link state database. You can limit the output of this command to a specified OSPFv3 process by entering the process ID, or to a specified OSPFv3 area by entering the area ID. The following is sample output from this command:

```
>enable
#show ospfv3 database database-summary
OSPFv3 router with ID: 5.5.5.5 (Process ID 61, VRF RED)
Area 0 database summary
    LSA Type          Count
    Router              2
    Network             1
    Link                2
    Prefix              1
```


Inter-area Prefix	2
Inter-area Router	0
Unknown	0
Subtotal	8
External	4
AS Unknown	0

Area 1 database summary

LSA Type	Count
Router	1
Network	0
Link	1
Prefix	1
Inter-area Prefix	2
Inter-area Router	1
Unknown	0
Subtotal	6
External	4
AS Unknown	0

Use the **show ospfv3 [**<*process id*>**] database external [**<*ipv6 address/prefix-length*>**] [**<*link state id*>**] [adv-router** <*router id*>**]** to provide details from the OSPFv3 link state database of external LSAs. You can limit the output of this command to a specified OSPFv3 process by entering a process ID, to a specified IPv6 address by entering the IPv6 address and prefix length, to a specified LSA by entering the link state ID, or to a specified advertising router by entering a router ID. The following is sample output from this command:

```
>enable
#show ospfv3 database external
External Link States

  Link State age: 689
  Link State type: AS External-LSA (0x4005)
  Link State ID: 0
  Advertising Router: 4.4.4.4
  Sequence Number: 0x80000057
  Checksum: 0xBE61
  Length: 36
    Prefix Address: 2001:7:1::
    Prefix Length: 64, Options: None
    Metric Type: 1 (Comparable directly to link state metric)
    Metric: 22222

  Link State age: 689
  Link State type: AS External-LSA (0x4005)
  Link State ID: 1
  Advertising Router: 4.4.4.4
  Sequence Number: 0x80000057
  Checksum: 0xAB43
  Length: 36
```



```
    Prefix Address: 2001:10:24:202::
    Prefix Length: 64, Options: None
    Metric Type: 1 (Comparable directly to link state metric)
    Metric: 22222

  Link State age: 317
  Link State type: AS External-LSA (0x4005)
  Link State ID: 2
  Advertising Router: 5.5.5.5
  Sequence Number: 0x80000001
  Checksum: 0x5D34
  Length: 52
    Prefix Address: 2001:8:1::
    Prefix Length: 64, Options: None
    Metric Type: 1 (Comparable directly to link state metric)
    Metric: 11111
    Forwarding Address: 2001:10:24:205::2

  Link State age: 317
  Link State type: AS External-LSA (0x4005)
  Link State ID: 3
  Advertising Router: 5.5.5.5
  Sequence Number: 0x80000001
  Checksum: 0x632C
  Length: 52
    Prefix Address: 2001:8:2::
    Prefix Length: 64, Options: None
    Metric Type: 1 (Comparable directly to link state metric)
    Metric: 11111
    Forwarding Address: 2001:10:24:205::2
```

Use the **show ospfv3 [**<*process id*> **[**<*area id*>**]] database inter-area prefix [**<*ipv6 address/prefix-length*>**] [**<*link state id*>**] [adv-router** <*router id*>**]** command to provide details from the OSPFv3 link state database of inter-area prefix LSAs. You can limit the output of this command to a specified OSPFv3 process by entering a process ID, to an OSPFv3 area by entering the area ID, to a specified IPv6 address by entering the IPv6 address and prefix length, to a specified LSA by entering the link state ID, or to a specified advertising router by entering a router ID. The following is sample output from this command:

```
>enable
#show ospfv3 database inter-area prefix
OSPFv3 router with ID: 4.4.4.4 (Process ID 61, VRF RED)

  Inter Area Prefix Link States, Area 0

  Link State age: 728
  Link State type: Inter-Area-Prefix-LSA (0x2003)
  Link State ID: 0
  Advertising Router: 4.4.4.4
  Sequence Number: 0x80000057
```



```
  Checksum: 0x55DE
  Length: 36
    Metric: 1
    Prefix Address: 2001:10:24:204::
    Prefix Length: 64, Options: None

  Link State age: 353
  Link State type: Inter-Area-Prefix-LSA (0x2003)
  Link State ID: 1
  Advertising Router: 5.5.5.5
  Sequence Number: 0x80000001
  Checksum: 0xEB98
  Length: 36
    Metric: 1
    Prefix Address: 2001:10:24:205::
    Prefix Length: 64, Options: None

Inter Area Prefix Link States, Area 1

  Link State age: 353
  Link State type: Inter-Area-Prefix-LSA (0x2003)
  Link State ID: 2
  Advertising Router: 5.5.5.5
  Sequence Number: 0x80000001
  Checksum: 0xE2A0
  Length: 36
    Metric: 1
    Prefix Address: 2001:10:24:106::
    Prefix Length: 64, Options: None

  Link State age: 344
  Link State type: Inter-Area-Prefix-LSA (0x2003)
  Link State ID: 4
  Advertising Router: 5.5.5.5
  Sequence Number: 0x80000001
  Checksum: 0xBBC6
  Length: 36
    Metric: 1
    Prefix Address: 2001:10:24:204::
    Prefix Length: 64, Options: None
```

Use the **show ospfv3 [**<*process id*> **[**<*area id*>**]] database inter-area router [**<*router id*>**] [**<*link state id*>**] [adv-router** <*router id*>**] [internal]** command to provide details from the OSPFv3 link state database of inter-area router LSAs. You can limit the output of this command to a specified OSPFv3 process by entering a process ID, to an OSPFv3 area by entering the area ID, to a specified router by entering the router ID, to a specified LSA by entering the link state ID, or to a specified advertising router by entering a router ID. The optional **internal** parameter displays the SPF calculation results for this LSA, and whether this LSA was used in route calculation.

The following is sample output from this command:

```
>enable
#show ospfv3 database inter-area router

OSPFv3 router with ID: 5.5.5.5 (Process ID 61, VRF RED)

  Inter Area Router Link States, Area 1

    Link State age: 394
    Link State type: Inter-Area-Prefix-LSA (0x2004)
    Link State ID: 3
    Advertising Router: 5.5.5.5
    Sequence Number: 0x80000001
    Checksum: 0x37C6
    Length: 32
      Options: V6, E, R, AF
      Metric: 1
      Destination Router: 4.4.4.4
```

Use the **show ospfv3 [**<*process id*> **[**<*area id*>**]] database link [interface** <*interface*>**] [**<*link state id*>**] [adv-router** <*router id*>**]** command to provide details from the OSPFv3 link state database of link LSAs. You can limit the output of this command to a specified OSPFv3 process by entering a process ID, to an OSPFv3 area by entering the area ID, to a specified interface by entering the interface, to a specified LSA by entering the link state ID, or to a specified advertising router by entering a router ID. The following is sample output from this command:

```
>enable
#show ospfv3 database link
OSPFv3 router with ID: 4.4.4.4 (Process ID 61, VRF RED)

  Link (Type-8) Link States, Area 0

    Link State age: 813
    Link State type: Link-LSA (Interface: eth 0/1.106) (0x0008)
    Link State ID: 8
    Advertising Router: 4.4.4.4
    Sequence Number: 0x80000057
    Checksum: 0xEFB2
    Length: 56
      Options: V6, E, R, DC
      Router Priority: 1
      Link-Local Address: FE80::CA9C:1DFF:FED6:E0A0
      Number of Prefixes: 1
        Prefix Address: 2001:10:24:106::
        Prefix Length: 64, Options: None

    Link State age: 438
    Link State type: Link-LSA (Interface: eth 0/1.106) (0x0008)
    Link State ID: 13
```



```
  Advertising Router: 5.5.5.5
  Sequence Number: 0x80000062
  Checksum: 0xCD3
  Length: 56
    Options: V6, E, R, AF
    Router Priority: 1
    Link-Local Address: FE80::2A0:C8FF:FE1F:CC53
    Number of Prefixes: 1
      Prefix Address: 2001:10:24:106::
      Prefix Length: 64, Options: None

Link (Type-8) Link States, Area 1

  Link State age: 441
  Link State type: Link-LSA (Interface: eth 0/2.1) (0x0008)
  Link State ID: 14
  Advertising Router: 5.5.5.5
  Sequence Number: 0x80000001
  Checksum: 0xD965
  Length: 56
    Options: V6, E, R, AF
    Router Priority: 1
    Link-Local Address: FE80::2A0:C8FF:FE1F:CC54
    Number of Prefixes: 1
      Prefix Address: 2001:10:24:205::
      Prefix Length: 64, Options: None
```

Use the **show ospfv3 [**<*process id*> **[**<*area id*>**]] database network [**<*link state id*>**] [adv-router** <*router id*>**]** command to provide details from the OSPFv3 link state database of network LSAs. You can limit the output of this command to a specified OSPFv3 process by entering a process ID, to an OSPFv3 area by entering the area ID, to a specified LSA by entering the link state ID, or to a specified advertising router by entering a router ID. The following is sample output from this command:

```
>enable
#show ospfv3 database network
OSPFv3 router with ID: 4.4.4.4 (Process ID 61, VRF RED)

  Network Link States, Area 0

  Link State age: 471
  Link State type: Network-LSA (0x2002)
  Link State ID: 8
  Advertising Router: 4.4.4.4
  Sequence Number: 0x80000001
  Checksum: 0xB01B
  Length: 32
    Options: V6, E, R, DC, AF
    Number of Attached Routers: 2
      Attached Router: 4.4.4.4
      Attached Router: 5.5.5.5
```


Use the **show ospfv3 [**<*process id*> **[**<*area id*>**]] database prefix [ref-lsa [network | router]] [**<*link state id*>**] [adv-router** <*router id*>**]** command to provide details from the OSPFv3 link state database of intra-area prefix LSAs. You can limit the output of this command to a specified OSPFv3 process by entering a process ID, to an OSPFv3 area by entering the area ID, to a referenced network or router LSA, to a specified LSA by entering the link state ID, or to a specified advertising router by entering a router ID. The following is sample output from this command:

```
>enable
#show ospfv3 database prefix
OSPFv3 router with ID: 4.4.4.4 (Process ID 61, VRF RED)

  Intra Area Prefix Link States, Area 0

    Link State age: 539
    Link State type: Intra-Area-Prefix-LSA (0x2009)
    Link State ID: 8192
    Advertising Router: 4.4.4.4
    Sequence Number: 0x80000001
    Checksum: 0x1809
    Length: 44
      Referenced LSA Type: 0x2002
      Referenced Link State ID: 8
      Referenced Advertising Router: 4.4.4.4
      Number of Prefixes: 1
        Prefix Address: 2001:10:24:106::
        Prefix Length: 64, Options: None, Metric: 0

  Intra Area Prefix Link States, Area 1

    Link State age: 497
    Link State type: Intra-Area-Prefix-LSA (0x2009)
    Link State ID: 0
    Advertising Router: 5.5.5.5
    Sequence Number: 0x80000003
    Checksum: 0x44FA
    Length: 44
      Referenced LSA Type: 0x2001
      Referenced Link State ID: 0
      Referenced Advertising Router: 5.5.5.5
      Number of Prefixes: 1
        Prefix Address: 2001:10:24:205::
        Prefix Length: 64, Options: None, Metric: 1
```

Use the **show ospfv3 [**<*process id*> **[**<*area id*>**]] database router [**<*link state id*>**] [adv-router** <*router id*>**] [internal]** command to provide details from the OSPFv3 link state database of router LSAs. You can limit the output of this command to a specified OSPFv3 process by entering a process ID, to an OSPFv3 area by entering the area ID, to a specified LSA by entering the link state ID, or to a specified advertising router by entering a router ID. The optional **internal** parameter displays SPF calculation results for this LSA, and whether this LSA was used in route calculation.

The following is sample output from this command:

```
>enable
#show ospfv3 database router
OSPFv3 router with ID: 4.4.4.4 (Process ID 61, VRF RED)

  Router Link States, Area 0

    Link State age: 575
    Link State type: Router-LSA (0x2001)
    Link State ID: 0
    Advertising Router: 4.4.4.4
    Sequence Number: 0x8000005A
    Checksum: 0x4F23
    Length: 40
      Options: V6, E, R, DC
      Flags: Area Border Router, AS Boundary Router
      Number of Links: 1
        Link connected to: Transit network Link
          Link Metric: 1
          Local Interface ID: 8
          Neighbor (DR) Interface ID: 8
          Neighbor (DR) Router ID: 4.4.4.4

    Link State age: 568
    Link State type: Router-LSA (0x2001)
    Link State ID: 0
    Advertising Router: 5.5.5.5
    Sequence Number: 0x80000066
    Checksum: 0xA3D8
    Length: 40
      Options: V6, E, R, AF
      Flags: Area Border Router, AS Boundary Router
      Number of Links: 1
        Link connected to: Transit network Link
          Link Metric: 1
          Local Interface ID: 13
          Neighbor (DR) Interface ID: 8
          Neighbor (DR) Router ID: 4.4.4.4

  Router Link States, Area 1

    Link State age: 536
    Link State type: Router-LSA (0x2001)
    Link State ID: 0
    Advertising Router: 5.5.5.5
    Sequence Number: 0x80000003
    Checksum: 0xA176
    Length: 24
```



```
   Options: V6, E, R, AF
   Flags: Area Border Router, AS Boundary Router
   Number of Links: 0
```

Use the **show ospfv3 [**<*process id*> **[**<*area id*>**]] database unknown [**<*link state id*>**] [adv-router** <*router id*>**] [as** | **area** | **link]** command to provide details from the OSPFv3 link state database of unknown LSAs. You can limit the output of this command to a specified OSPFv3 process by entering a process ID, to an OSPFv3 area by entering the area ID, to a specified LSA by entering the link state ID, or to a specified advertising router by entering a router ID. The optional **as**, **area**, and **link** parameters filter specific scoped unknown LSAs by AS, area, or link. The following is sample output from this command:

```
>enable
#show ospfv3 database unknown
OSPFv3 router with ID: 4.4.4.4 (Process ID 61, VRF RED)

 Unknown Link States, Area 603979776

  Link State age: 273
  Link State type: Unknown (0x3FFF)
  Link State ID: 1
  Advertising Router: 10.80.52.106
  Sequence Number: 0x808DE8B1
  Checksum: 0xF1AE
  Length: 40
   Scope: Link Local
```

Use the **show ospfv3 [**<*process id*> **[**<*area id*>**]] interface [**<*interface*>**]** command to display OSPFv3 information related to router interfaces. You can limit the output of this command to a specified OSPFv3 process by entering a process ID, to an OSPFv3 area by entering the area ID, to a specified interface by entering the interface. The following is sample output from this command:

```
>enable
#show ospfv3 interface
eth 0/1.106 is UP
    Link Local Address FE80::2A0:C8FF:FE1F:CC53, Interface ID 13
    Area 0, Process ID 61, VRF RED, Instance ID 0, Router ID 5.5.5.5
     Area 1, Process ID 24, VRF RED, Instance ID 1, Router ID 5.5.5.5
    Network type Broadcast, Cost: 1
    Transmit delay is 1, State BDR, Priority 1
    Designated Router (ID) 4.4.4.4, local address FE80::CA9C:1DFF:FED6:E0A0
    Backup Designated Router (ID) 5.5.5.5, local address FE80::2A0:C8FF:FE1F:CC53
    Timer intervals configured, Hello 10, Dead 40, Retransmit 5
        Hello due in 00:00:04
     Neighbor Count is 1, Adjacent neighbors count is 1
        Adjacent with neighbor 4.4.4.4 (Designated Router)
    Suppress hello for 0 neighbor(s)
eth 0/2.1 is UP
    Link Local Address FE80::2A0:C8FF:FE1F:CC54, Interface ID 14
     Area 1, Process ID 61, Instance ID 0, Router ID 5.5.5.5
    Network type Broadcast, Cost: 1
    Transmit delay is 1, State DR, Priority 1
```



```
Designated Router (ID) 5.5.5.5, local address FE80::2A0:C8FF:FE1F:CC54
Backup Designated Router (ID) 0.0.0.0, local address ::
Timer intervals configured, Hello 10, Dead 40, Retransmit 5
    Hello due in 00:00:06
Neighbor Count is 0, Adjacent neighbors count is 0
Suppress hello for 0 neighbor(s)
```

Use the **show ospfv3 [**<*process id*> **[**<*area id*>**]] neighbor [**<*interface*>**] [**<*router id*> | **hostname] [detail]** command to display OSPFv3 information related to OSPFv3 neighbors. You can limit the output of this command to a specified OSPFv3 process by entering a process ID, to an OSPFv3 area by entering the area ID, to a specified neighbor by entering the connecting interface, or to a neighbor having the specific router ID or host name. The optional **detail** parameter specifies that more detailed information is displayed.

The following is sample output from this command:

```
>enable
#show ospfv3 neighbor
OSPFv3 Router with ID (5.5.5.5) (Process ID 61, VRF RED)

Neighbor ID     Pri State           Dead Time  Intf-ID  Interface
4.4.4.4           1 FULL/DR         00:00:34      8       eth 0/1.106
```

Use the **show ospfv3 [**<*process id*>**] summary-prefix** command to provide details about redistributed routes that have been summarized using the **summary-prefix** configuration command (refer to Step 7 on *page 15*). You can limit the output of this command to a specified OSPFv3 process by entering a process ID. The following is sample output of from this command:

```
>enable
#show ospfv3 61 summary-prefix
OSPFv3 Summary Addresses, Process ID 61, VRF RED:
  2001:8:7::/48 Metric 11111, Type 1, advertise
```

Use the **show ipv6 crypto ipsec sa ospfv3** command to display SAs created for the local stack, such as OSPFv3. The following is sample output from this command (only the part pertaining to OSPFv3 is shown):

```
>enable
#show ipv6 crypto ipsec sa ospfv3
Peer IP Address: ::
  Direction: Inbound
  Encapsulation: ESP transport
  SPI: 0x00000BB8 (3000)
  RX Bytes: 512

Peer IP Address: ::
  Direction: Outbound
  Encapsulation: ESP transport
  SPI: 0x00000BB8 (3000)
  TX Bytes: 512
```


OSPFv3 Clear Commands

Clear commands are used to reset and restart specific OSPFv3 processes, and to refresh routes distributed into OSPFv3 processes.

Use the **clear ospfv3 [**<*process id*>**] process** command to reset and restart all OSPFv3 processes. To only reset or restart one OSPFv3 process, enter the process ID. For example, to restart OSPFv3 process **5**, enter the command from the Enable mode as follows:

```
>enable
#clear ospfv3 5 process
```

Use the **clear ospfv3 [**<*process id*>**] redistribution** command to refresh routes distributed into the specified OSPFv3 process(es). To only refresh routes distributed to a single OSPFv3 process, enter the process ID. For example, to refresh routes distributed into OSPFv3 process **5**, enter the command from the Enable mode as follows:

```
>enable
#clear ospfv3 5 redistribution
```

OSPFv3 Debug Commands

Debug commands are used to enable debug messages for the OSPFv3 feature. Use the **debug ospfv3 [**<*process id*>**] [adj [errors] | database-timer | events | flood [errors] | hello [errors] | lsa-generation | packet errors | packet [rx | tx [summary]] | retransmission | spf [route-calculation]]** command to enable OSPFv3 debug messages. You can limit debug messages to a specific OSPFv3 process by entering a process ID. In addition, you can specify that OSPFv3 adjacency events are displayed (**adj**), that OSPFv3 database timer information is displayed (**database-timer**), that OSPFv3 events are displayed (**events**), that OSPFv3 flooding information is displayed (**flood**), that OSPFv3 Hello events are displayed (**hello**), that OSPFv3 LSA generation information is displayed (**lsa-generation**), that OSPFv3 packet details in either directions (**rx** and **tx**) are displayed, that OSPFv3 errors with received packets are displayed (**packet errors**), that OSPFv3 retransmission events are displayed (**retransmission**), or that OSPFv3 SPF events (**spf**) and route calculations are displayed (**spf route-calculation**). In addition, you can specify that errors about certain information are displayed using the **errors** parameters, and that summary information about packet details is displayed using the **summary** parameter (this option limits packet information to a single line). For example, to display debug event messaging for all of OSPFv3 on the unit, enter the command from the Enable mode as follows:

```
>enable
#debug ospfv3
Receiving OSPFv3 packet from 2001:db8:10:24::106.4 to FF02::5 on eth 0/1.6
    SysUpTime=1222577915 ms.
    Hello Packet from Router ID: 1.1.1.4; Ver:2 Length:48
    Area ID: 0.0.0.0 Checksum: )x8659;; Using Null Authentication: 0:0
    PrefixLenV4: /64; Hello Interval: 10 Options: 0x13 Router Priority: 1 Router Dead Interval: 40
    Designated Router: 10.24.106.4 Backup Designated Router: 10.24.106.5
    1 Neighbors:
        123.1.1.1
    16:35:24: OSPFv3: HELLO received form 1.1.1.4, neighbor state is FULL
```

