

# How To | Configure Basic OSPFv3 Scenarios

## Contents

---

1.	Contents .....	1
2.	Introduction .....	2
3.	OSPFv3 Basic Configuration .....	3
	Commands for basic OSPFv3 configuration .....	4
	OSPFv3 basic configuration explained .....	5
4.	The show commands .....	7
	show ipv6 route .....	7
	show ipv6 ospf .....	8
	show ipv6 ospf route .....	10
	show ipv6 ospf database.....	11
	show ipv6 ospf database <LSA type>.....	13
5.	OSPFv3 Configuration - Multiple Areas .....	17
	OSPFv3 multiple area configuration explained .....	20
6.	OSPFv3 summarization .....	22
7.	Using RIPng and OSPFv3 together .....	23
8.	Route-maps and OSPFv3 .....	25
	Matching on a prefix-list .....	25
	Matching on an access-list .....	26
	Matching on route-type .....	27
	Other possibilities .....	28
9.	Differences between OSPFv2 and OSPFv3 .....	29

# Introduction

---

This How To Note provides an introduction to Open Shortest Path First Version 3 (OSPFv3) operation in AlliedWare Plus. It is not a complete explanation of the protocol, rather it is an explanation of how to configure some introductory OSPFv3 scenarios, and how to interpret the output of the AlliedWare Plus OSPFv3 **show** commands.

This note will teach you the basics of using OSPFv3 on AlliedWare Plus, and how to view and monitor your OSPFv3 set up and behaviour.

This note contains:

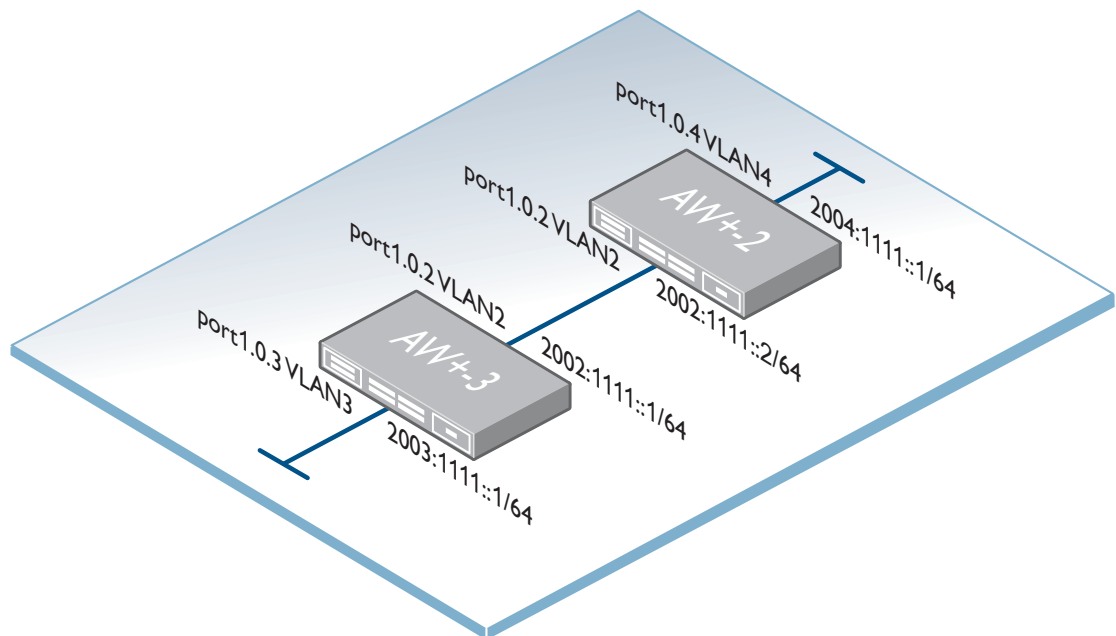
- a work through of a simple 2-router scenario
- an analysis of a basic inter-area configuration
- a discussion on OSPFv3 route summarization
- a summary of the key differences between OSPFv3 and OSPFv2.

In addition, this How To Note also briefly looks at redistribution of connected and RIPng routes.

# OSPFv3 Basic Configuration

This section outlines a basic OSPFv3 configuration.

As a starting point for an introduction to OSPFv3 commands and concepts, consider a minimum single area configuration as shown below:



## Commands for basic OSPFv3 configuration

The following commands are required to configure the basic OSPFv3 operation. Commands in **bold** are explained in further detail in the next section.

```
awplus-2: vlan database
            vlan 2
            vlan 4
            !
            interface port1.0.2
            switchport access vlan 2
            !
            interface port1.0.4
            switchport access vlan 4
            !
            interface vlan2
ipv6 address 2002:1111::2/64
ipv6 router ospf area 0 tag 10
            !
            interface vlan4
ipv6 address 2004:1111::1/64
            !
ipv6 forwarding
            !
router ipv6 ospf 10
router-id 192.168.1.2
redistribute connected
```

```
awplus-3: vlan database
            vlan 2
            vlan 3
            !
            interface port1.0.2
            switchport access vlan 2
            !
            interface port1.0.3
            switchport access vlan 3
            !
            interface vlan2
ipv6 address 2002:1111::1/64
ipv6 router ospf area 0 tag 10
```

```
!  
interface vlan3  
  ipv6 address 2003:1111::1/64  
!  
ipv6 forwarding  
!  
router ipv6 ospf 10  
  router-id 192.168.1.1  
  redistribute connected
```

## OSPFv3 basic configuration explained

This section explains the configuration commands that are shown in **bold** above:

### IPv6 interface commands

```
awplus-2(config-if)#ipv6 address 2002:1111::2/64
```

This is the IPv6 prefix (address) and prefix length (mask) assigned to the interface. A /64 netmask is used, as this is the standard choice of netmask for IPv6 subnets

```
awplus-2(config-if)#ipv6 router ospf area 0 tag 10
```

This command tells the switch to advertise this OSPFv3 interface's IPv6 network in OSPFv3. It also configures the switch to seek OSPFv3 neighbors on the subnet accessed via this interface.

The **area** parameter configures the OSPF area this interface is in. Area 0 is the backbone. The **tag** parameter references the process-id number used in the global configuration **router ipv6 ospf x** command, explained below.

### IPv6 forwarding commands

```
awplus-2(config)#ipv6 forwarding
```

This command turns on IPv6 forwarding, and is necessary for the operation of OSPFv3. By default, AlliedWare Plus switches do not L3 forward IPv6.

## IPv6 router commands

```
awplus-2 (config) #router ipv6 ospf 10
```

This command accesses the OSPFv3 configuration mode. The OSPF process number, in this case 10, is a user-configurable number which will match at least one of the numbers configured with **ipv6 router ospf area tag** on an interface. You can have multiple OSPFv3 processes, allowing different OSPFv3 processes to operate on different interfaces.

```
awplus-2 (config-router) #router-id 192.168.1.2
```

In OSPFv3 for IPv6, it is possible that no IPv4 addresses will be configured on any interface. In this case, you must use the **router-id** command to configure a dotted decimal address as the router ID, before the OSPF process begins. This router-id address does not have to actually be configured on any interface on the switch. Although the router-id is written in the form of an IPv4 address, it is actually just a 32 bit number, so there is no need for it to be an IP address that exists on the router.

```
awplus-2 (config-router) #redistribute connected
```

This command tells the switch to advertise any IPv6 networks configured on the switch that are not running OSPFv3. The interface must be in the "up" state to be advertised.

---

**Note:** *OSPFv3 automatically advertises the networks that are attached to the switch's other OSPFv3 interfaces and configured for the same process ID.*

---

# The show commands

Once you have connected and configured OSPFv3, you can view and monitor its operation with a set of **show** commands. This section outlines the OSPFv3 **show** commands available in AlliedWare Plus. Each command's output is described in detail, along with how you can use it to monitor the operation of your OSPFv3 network.

## show ipv6 route

This command uses the switch's IPv6 route table to discover what routes have been learnt via OSPFv3.

### Show output

```
awplus-2#sh ipv6 route
IPv6 Routing Table
Codes: C - connected, S - static, R - RIP, O - OSPF, B - BGP
Timers: Uptime
C   2002:1111::/64 via ::, vlan2, 17:58:51
O   2003:1111::/64 [110/20] via fe80::200:cdff:fe27:c1ff, vlan2, 00:00:17
C   2004:1111::/64 via ::, vlan4, 00:00:26
C   fe80::/64 via ::, vlan4, 00:00:26
C   fe80::/64 via ::, vlan2, 17:58:51
```

### Output explained

On the switch **awplus-2**, you can see that the route to network **2003:1111::/64** on 'awplus-3' has been learnt via OSPF.

The table shows each IPv6 route and its prefix, as well as how the route was learned, which interface it was learned on, and the length of time for which the route has been in the route table.

The route that has been learnt via OSPFv3 from awplus-3 is shown as **[110/20]** These numbers are, in two parts:

#### 110 - the Administrative Distance (AD)

The Administrative Distance (AD) defines the reliability of a routing protocol. OSPF = 110, Connected routes = 0, static routes = 1, and RIP = 120. The lower the AD, the more reliable the protocol is considered.

If the switch has learnt multiple routes to the same network with different ADs, it uses the route with the lowest AD.

#### 20 - the Metric

The nexthop IPv6 address **fe80::200:cdff:fe27:c1ff** may seem surprising, given that the IPv6 address on the VLAN2 interface of 'awplus-3' is **2002:1111::1**. In fact, **fe80::200:cdff:fe27:c1ff** is the link-local IPv6 address of awplus-3's VLAN2 interface.

Every IPv6 interface, on any device, has an automatically calculated link-local address. It is quite common to use the link-local address as the next-hop in IPv6 routes. Further explanation of link-local IPv6 addresses is beyond the scope of this note,

On switch awplus-3 you can see the route to network **2004:1111::/64** on awplus-2 via OSPF

```
awplus-3#sh ipv6 route
IPv6 Routing Table
Codes: C - connected, S - static, R - RIP, O - OSPF, B - BGP
Timers: Uptime
C   2002:1111::/64 via ::, vlan2, 17:59:25
C   2003:1111::/64 via ::, vlan3, 00:00:54
O   2004:1111::/64 [110/20] via fe80::200:cdf:fe29:8ed3, vlan2, 00:00:57
C   fe80::/64 via ::, vlan3, 00:00:54
C   fe80::/64 via ::, vlan2, 17:59:25
```

## show ipv6 ospf

This command gives general information about the current OSPFv3 operation.

### Show output

```
awplus-2#sh ipv6 ospf
Routing Process "OSPFv3 (10)" with ID 192.168.1.2
Route Licence: Route : Limit=Unlimited, Allocated=3, Visible=2, Internal=1
OSPFv3-FULL Route Licence: Breach: Current=0, Watermark=0
Process uptime is 18 hours 2 minutes
Current grace period is 0 secs
SPF schedule delay min 0.500 secs, SPF schedule delay max 50.0 secs
```

**Output explained** **Routing Process "OSPFv3 (10)" with ID 192.168.1.2** identifies OSPFv3 process-id 10 with router-id 192.168.1.2

#### Route Licence:

- **Limit=Unlimited** refers to the OSPFv3 route-limit licence status of the switch. There are 2 OSPFv3 licenses available, providing either a maximum of 64 IPv6 routes or unlimited routes. This switch has an unlimited route OSPFv3 license, which you can verify with the **show license** command.
- The number of **Allocated** OSPFv3 routes is 3, of which 2 are **Visible** OSPFv3 routes, and 1 is **Internal**. You can view the visible routes with the **show ipv6 ospf route** command. Internal routes refer to internal OSPFv3 routes to Autonomous System Boundary Routers (ASBRs), which are only used for calculating nexthops for AS-EXT routes, and which are not usable as user routes.

**The OSPFv3-FULL Route Licence: Breach:** output shows you whether the number of OSPFv3 routes learned by the switch has exceeded the number allowed by the license.

**Current=** shows the number of OSPFv3 routes that currently exceed the license.

**Watermark=** shows the highest number of OSPFv3 routes that have ever exceeded the license.

**Process uptime** is the time elapsed since this OSPFv3 process was configured and started.



**Current grace period** shows the grace restart period. The default is 120 seconds.

### **SPF timers**

You can configure **SPF schedule delay** timers in **router ipv6 ospf x** configuration mode.

You cannot configure the **Minimum LSA interval** and **arrival** times. The minimum interval is always 5 seconds, and the minimum arrival is always 1 second.

The **Number of incoming current DD exchange neighbors** and the **Number of outgoing current DD exchange neighbors** counters are not incremented.

You can view the **Number of external LSA** with the **show ipv6 ospf database external** command. This is explained further in "[show ipv6 ospf database <LSA type>](#)" on page 13.

**SPF algorithm executed** shows the number of times that SPF has executed due to a network topology change

**Number of LSA** refers to the total number of LSAs in the LSDB.

**Number of Unknown LSA** refers to the total number of unknown LSAs in the LSDB. OSPFv3 can either treat Unknown LSAs as having link-local flooding scope, or store and flood them as if they are understood, while ignoring them in their own SPF algorithms.

## show ipv6 ospf route

This command shows the contents of the OSPFv3 route table. These are the routes that OSPFv3 has calculated from its LSA database. Not necessarily all the routes in the OSPFv3 route table are in the router's main IPv6 route table. If any routes in the table clash with routes that are learnt from another routing protocol with a lower Administrative Distance, then the latter protocol's routes are the ones that go into the main IPv6 route table.

### Show output

```
awplus-2#sh ipv6 ospf route
OSPFv3 Process (10)
Codes: C - connected, D - Discard, O - OSPF, IA - OSPF inter area
       E1 - OSPF external type 1, E2 - OSPF external type 2

  Destination                                Metric
  Next-hop
C  2002:1111::/64                             1
   directly connected, vlan2, Area 0.0.0.0
E2 2003:1111::/64                             1/20
   via fe80::200:cdff:fe27:c1ff, vlan2
```

### Output explained

Although the route to network **2002:1111::/64** is shown as **C**, which means connected, it is also learned via OSPFv3 from its neighbor.

The route to **2003:1111::/64** on VLAN3 of awplus-3 is an AS-EXT route. This is because OSPFv3 is not configured on VLAN3 of awplus-3, so it is not an OSPFv3 interface. Its subnet is only advertised into OSPFv3 because awplus-3 is configured to redistribute connected routes. The route type is **E2**, which is the default for redistributed routes.

In the E2 route entry **1/20**, **1** shows the total internal cost for the route, and **20** shows the external cost. Since the route is E2, the route calculation uses only the external cost.

# show ipv6 ospf database

## Show output

```

awplus-2#sh ipv6 ospf database
OSPFv3 Router with ID (192.168.1.2) (Process 10)

Link-LSA (Interface vlan2)
Link State ID  ADV Router      Age  Seq#          CkSum  Prefix
0.0.0.202     192.168.1.1    1672 0x8000002d   0xb654  1
0.0.0.202     192.168.1.2    1676 0x8000002d   0xd98d  1

Router-LSA (Area 0.0.0.0)
Link State ID  ADV Router      Age  Seq#          CkSum  Link
0.0.0.0        192.168.1.1    47   0x80000030   0x0b75  1
0.0.0.0        192.168.1.2    236  0x80000030   0x057a  1

Network-LSA (Area 0.0.0.0)
Link State ID  ADV Router      Age  Seq#          CkSum
0.0.0.202     192.168.1.2    1631 0x8000002d   0xbf2e

Intra-Area-Prefix-LSA (Area 0.0.0.0)
Link State ID  ADV Router      Age  Seq#          CkSum  Prefix  Reference
0.0.0.2        192.168.1.2    1626 0x8000002d   0x0aa5  1       Network-LSA

AS-external-LSA
Link State ID  ADV Router      Age  Seq#          CkSum
0.0.0.2        192.168.1.1    1792 0x80000009   0x15ec  E2
0.0.0.3        192.168.1.2    616  0x80000004   0x1be8  E2

```

## Output explained

Central to the operation of OSPFv3 is its LSA database. Each LSA in the database consists of a set of fields, as explained below:

Table 1:

Field	Description
Link State ID	For a: Network LSA: the interface ID of the designated router's (DR) interface Link LSA: the interface ID of the egress interface on the router that sent the LSA Grace LSA: the interface ID of the egress interface on the router that sent the LSA Other type of LSA: an index number that increments with each new LSA
ADV Router	Advertising Router ID
Age	Link-state age in seconds
Seq#	Link-state sequence number (detects old or duplicate LSAs)
CkSum	Checksum of the LSA
Prefix	The number of different IPv6 prefixes in the LSA. This is specific to Link and Prefix LSAs.
Link	The number of links in the LSA. This is specific to router LSAs.
Reference	The Intra-Area Prefix LSA tells us that the IPv6 prefixes associated with the Router ID appear in either a Router or Network LSA. The Reference field will therefore show either 'Router' or 'Network'.

The **Link-LSA** communicates information that is significant only to two directly connected neighbors. A separate Link-LSA originates on each of a router's attached links belonging to an OSPFv3 domain, and a receiving router.

VLAN interface indices starts at 201, so **Link State ID 0.0.0.202** is that of VLAN 2.

The **Router-LSA** describes the originating router and its links to neighbors, for use in SPF calculations. Note that the ID is 0.0.0.0, which indicates that this is the first LSA, other than a Link, Network or Grace LSA, originated by each router.

For broadcast networks only: the DR floods the **Network-LSA** out of all its interfaces. This LSA lists OSPF-speaking routers on the network.

When a link or its prefix changes, the connected router originates an **Intra-Area Prefix LSA** to flood the information throughout the area. This LSA does not trigger an SPF calculation.

The **AS-external-LSA** advertises prefixes that are external to the OSPF routing domain. One LSA is required for each external prefix advertised.

You can find more detailed information about LSAs in the following section.

## show ipv6 ospf database <LSA type>

The **show ipv6 ospf database** command shows a summary of each of the LSAs in the database. To see the full content of each LSA, you must use one of the **show ipv6 ospf database <LSA-type>** commands. For example, to see the full detail of all the network LSAs, you would use the command **show ipv6 ospf database network**.

### show ipv6 ospf database external

#### Show output

```
awplus-2#sh ipv6 ospf database external
OSPFv3 Router with ID (192.168.1.2) (Process 10)
AS-external-LSA
  LS age: 838
  LS Type: AS-External-LSA
  Link State ID: 0.0.0.2
  Advertising Router: 192.168.1.1
  LS Seq Number: 0x80000006
  Checksum: 0x1BE9
  Length: 36
    Metric Type: 2 (Larger than any link state path)
    Metric: 20
    Prefix: 2003:1111::/64
    Prefix Options: 0 (-|-|-)
  LS age: 842
  LS Type: AS-External-LSA
  Link State ID: 0.0.0.1
  Advertising Router: 192.168.1.2
  LS Seq Number: 0x80000006
  Checksum: 0x2BD8
  Length: 36
    Metric Type: 2 (Larger than any link state path)
    Metric: 20
    Prefix: 2004:1111::/64
    Prefix Options: 0 (-|-|-)
```

**Output explained** **LS age** is the LSA age in seconds.

**AS-External LSA** means that this route was learned by a method other than OSPF and imported (redistributed) into OSPFv3. In this case, these routes were directly connected interfaces.

There are 2 types of External LSAs - **Metric Type 1** and 2

- Metric Type 1 includes the metric to get to the ASBR plus the external metric
- Metric Type 2 only includes the external metric

By default, an imported route is given **Metric: 20**.

## show ipv6 ospf database network

### Show output

```
awplus-2#show ipv6 ospf database network
      OSPFv3 Router with ID (192.168.1.2) (Process 10)
      Network-LSA (Area 0.0.0.0)
LS age: 789
LS Type: Network-LSA
Link State ID: 0.0.0.202
Advertising Router: 192.168.1.2
LS Seq Number: 0x80000189
Checksum: 0x048C
Length: 32
Options: 0x000013 (-|R|-|-|E|V6)
  Attached Router: 192.168.1.2
  Attached Router: 192.168.1.1
```

### Output explained

Compared with OSPFv2, OSPFv3 has the additional **R** and **V6** bits in the **Options:** field:

- The R bit indicates that the originator is an active router. If the R bit is clear, the switch cannot process routes that transit the advertising node. Clearing the R bit would be appropriate for a multi-homed host that wants to participate in routing, but does not want to forward non-locally addressed packets.
- The V6 bit indicates that the router/link should be included in IPv6 routing calculations.

The OSPFv3 **Options** are a 24 bit field with the following options:

Table 2:

Value	Description	Reference
0x000001	V6-bit	[RFC5340]
0x000002	E-bit	[RFC2328]
0x000004	Deprecated	[RFC5340]
0x000008	N-bit	[RFC3101]
0x000010	R-Bit	[RFC5340]
0x000020	DC-bit	[RFC1793]
0x000040	Reserved for OSPFv2 migrated options	[RFC5340]
0x000080	Reserved for OSPFv2 migrated options	[RFC5340]
0x000100	AF-bit	[RFC5838]
0x000200	L-bit	[RFC5613]

So, **Options: 0x000013** means:

0x000010 - R-bit

0x000002 - E-bit

0x000001 - V6-bit

The **Attached Router** fields are the Router-IDs of the two OSPFv3 routers connected to the 2002:1111::/64 subnet.

## show ipv6 ospf database router

### Show output

```
awplus-2#show ipv6 ospf database router
      OSPFv3 Router with ID (192.168.1.2) (Process 10)
      Router-LSA (Area 0.0.0.0)

LS age: 1095
LS Type: Router-LSA
Link State ID: 0.0.0.0
Advertising Router: 192.168.1.1
LS Seq Number: 0x8000018C
Checksum: 0x4FD3
Length: 40
Flags: 0x02 (-|-|E|-)

Options: 0x000013 (-|R|-|-|E|V6)
  Link connected to: a Transit Network
  Metric: 1
  Interface ID: 202
  Neighbor Interface ID: 202
  Neighbor Router ID: 192.168.1.2
LS age: 783
LS Type: Router-LSA
Link State ID: 0.0.0.0
Advertising Router: 192.168.1.2
LS Seq Number: 0x8000018B
Checksum: 0x4BD7
Length: 40
Flags: 0x02 (-|-|E|-)
Options: 0x000013 (-|R|-|-|E|V6)
  Link connected to: a Transit Network
  Metric: 1
  Interface ID: 202
  Neighbor Interface ID: 202
  Neighbor Router ID: 192.168.1.1
```

### Output explained

The **Flags** parameter shows the values set in the OSPFv3 Router Properties Registry:

Table 3:

Value	Description	Reference
0x01	B-bit	[RFC5340]
0x02	E-bit	[RFC5340]
0x04	V-bit	[RFC5340]
0x08	Deprecated	[RFC5340]
0x10	Nt-bit	[RFC5340]

- **Bit V**  
When set, shows the router is an endpoint of one or more fully adjacent virtual links having the described area as a transit area. V stands for virtual link endpoint.
- **Bit E**  
When set, shows the router is an AS boundary router. E stands for external.

- **Bit B**  
When set, shows the router is an area border router. B stands for border.
- **Bit x**  
This bit was previously used by Multicast OSPF (MOSPF) and has been deprecated for OSPFv3. The bit should be set to 0 and ignored when received. It may be reassigned in the future.
- **Bit Nt**  
When set, shows the router is a Not So Stubby Area (NSSA) border router that is unconditionally translating NSSA-LSAs into AS-external-LSAs. Nt stands for NSSA translation. Note that such routers have their NSSATranslatorRole area configuration parameter set to Always.

The router lists the one OSPF link that is attached to it. The **Interface ID** is the ifindex of the interface (VLAN2). The router LSA will also list all the OSPFv3 routers that this router has neighbored with via that interface. In this case, it has just the one neighbor on VLAN2.

### show ipv6 ospf database link

#### Show output

```
awplus-2#show ipv6 ospf database link
      OSPFv3 Router with ID (192.168.1.2) (Process 10)
      Link-LSA (Interface vlan2)

LS age: 850
LS Type: Link-LSA
Link State ID: 0.0.0.202
Advertising Router: 192.168.1.1
LS Seq Number: 0x80000189
Checksum: 0xFAB2
Length: 56
Priority: 1
Options: 0x000013 (-|R|-|-|E|V6)
Link-Local Address: fe80::200:cdf:fe27:c1ff
Number of Prefixes: 1
  Prefix: 2002:1111::/64
  Prefix Options: 0 (-|-|-|-)
LS age: 834
LS Type: Link-LSA
Link State ID: 0.0.0.202
Advertising Router: 192.168.1.2
LS Seq Number: 0x80000189
Checksum: 0x1EEB
Length: 56
Priority: 1
Options: 0x000013 (-|R|-|-|E|V6)
Link-Local Address: fe80::200:cdf:fe29:8ed3
Number of Prefixes: 1
  Prefix: 2002:1111::/64
  Prefix Options: 0 (-|-|-|-)
```

#### Output explained

Every IPv6 interface has a unique **Link-Local Address**.


There may be multiple IPv6 secondary addresses configured on the interface, in which case they would all be listed here.



# OSPFv3 Configuration - Multiple Areas

This section looks at a slightly larger scenario, involving three routers spread over two areas. This enables us to introduce some aspects of OSPFv3 inter-area routing.

In this scenario, there are two areas - the backbone and area 1. There are two routes that are each fully within one area, and one router that is the border router between the 2 areas.



```
awplus-1: vlan database
    vlan 3
    vlan 5
!
interface port1.0.3
    switchport access vlan 3
!
interface port1.0.5
    switchport access vlan 5
!
interface vlan3
    ipv6 address 2003:1111::2/64
    ipv6 router ospf area 1 tag 10
!
interface vlan5
    ipv6 address 2005:1111::1/64
!
ipv6 forwarding
!
router ipv6 ospf 10
    router-id 192.168.1.3
    redistribute connected
```

```
awplus-2: vlan database
    vlan 2
    vlan 4
!
interface port1.0.2
    switchport access vlan 2
!
interface port1.0.4
    switchport access vlan 4
!
interface vlan2
    ipv6 address 2002:1111::2/64
    ipv6 router ospf area 0 tag 10
!
interface vlan4
    ipv6 address 2004:1111::1/64
```

```
!  
ipv6 forwarding  
!  
router ipv6 ospf 10  
  router-id 192.168.1.2  
  redistribute connected
```

**awplus-3:** vlan database

```
  vlan 2  
  vlan 3  
!  
interface port2.0.2  
  switchport access vlan 2  
!  
interface port2.0.3  
  switchport access vlan 3  
!  
interface vlan2  
  ipv6 address 2002:1111::1/64  
  ipv6 router ospf area 0 tag 10  
!  
interface vlan3  
  ipv6 address 2003:1111::1/64  
  ipv6 router ospf area 1 tag 10  
!  
ipv6 forwarding  
!  
router ipv6 ospf 10  
  router-id 192.168.1.1
```

## OSPFv3 multiple area configuration explained

Using **show ipv6 ospf route**, we can now see that there are Inter-Area routes:

```
awplus-3#sh ipv6 ospf route
OSPFv3 Process (10)
Codes: C - connected, D - Discard, O - OSPF, IA - OSPF inter area
       E1 - OSPF external type 1, E2 - OSPF external type 2

  Destination                                Metric
  Next-hop
C  2002:1111::/64                             1
   directly connected, vlan2, Area 0.0.0.0
C  2003:1111::/64                             1
   directly connected, vlan3, Area 0.0.0.1
E2 2004:1111::/64                             1/20
   via fe80::200:cdff:fe29:8ed3, vlan2
E2 2005:1111::/64                             1/20
   via fe80::215:77ff:fec2:4d53, vlan3
```

Now that there are 2 OSPF Areas, you can see some extra fields in the **show ipv6 ospf database** output for the Inter-Area LSAs:

```
awplus-3#show ipv6 ospf database
      OSPFv3 Router with ID (192.168.1.1) (Process 10)
      Inter-Area-Prefix-LSA (Area 0.0.0.0)
Link State ID  ADV Router      Age  Seq#      CkSum
0.0.0.7        192.168.1.1    231 0x80000001 0x72cb
      Inter-Area-Router-LSA (Area 0.0.0.0)
Link State ID  ADV Router      Age  Seq#      CkSum
0.0.0.2        192.168.1.1    44  0x80000001 0x58f2
      Intra-Area-Prefix-LSA (Area 0.0.0.0)
Link State ID  ADV Router      Age  Seq#      CkSum  Prefix  Reference
0.0.0.2        192.168.1.2    961 0x800001ad 0x0628  1      Network-LSA
      Inter-Area-Prefix-LSA (Area 0.0.0.1)
Link State ID  ADV Router      Age  Seq#      CkSum
0.0.0.3        192.168.1.1    226 0x80000002 0x8cb5
      Inter-Area-Router-LSA (Area 0.0.0.1)
Link State ID  ADV Router      Age  Seq#      CkSum
0.0.0.3        192.168.1.1    226 0x80000002 0x3e0c
```

## show ipv6 ospf database inter-prefix

### Show output

```
awplus-3#show ipv6 ospf database inter-prefix
      OSPFv3 Router with ID (192.168.1.1) (Process 10)
      Inter-Area-Prefix-LSA (Area 0.0.0.0)

LS age: 457
LS Type: Inter-Area-Prefix-LSA
Link State ID: 0.0.0.7
Advertising Router: 192.168.1.1

LS Seq Number: 0x80000001
Checksum: 0x72CB
Length: 36
  Metric: 1
  Prefix: 2003:1111::/64

Prefix Options: 0
      Inter-Area-Prefix-LSA (Area 0.0.0.1)

LS age: 452
LS Type: Inter-Area-Prefix-LSA
Link State ID: 0.0.0.3
Advertising Router: 192.168.1.1

LS Seq Number: 0x80000002
Checksum: 0x8CB5
Length: 36
  Metric: 1
  Prefix: 2002:1111::/64

Prefix Options: 0
```

### Output explained

#### The Inter-Area Prefix LSA:

- describes the destination outside the area but still in the Autonomous System
- the summary is created for one area, but flooded in all other areas
- originated by an Area Border Router
- only intra-area routes are advertised into the backbone (Area 0)
- external routes are also advertised into the backbone from non-stub areas

**Advertising Router: 192.168.1.1** shows the Router ID of this switch (awplus-3).

**Prefix: 2003:1111::/64** is configured on an interface on this switch in Area 1.

**Inter-Area-Prefix-LSA (Area 0.0.0.1)** is for Area 1.

# OSPFv3 summarization

---

OSPFv2 can summarize both AS-External routes with the **summary-address** command, and inter-area routes with the **area range** command. OSPFv3 can only summarize inter-area routes.

The inter-area route summarization in OSPFv3 operates similarly to that of OSPFv2.

The **area range** command defines summarization routes that will be advertised via inter-area summarization LSAs. The component routes within that range will not be advertised.

At least one component route has to be in the route table before it will advertise the summarization route.

To illustrate the operation of the inter-area summarization, we can send a summarization for networks 2003:1111::/64 and 2003:2222::/64 from switch awplus-3 to awplus-2 with the following configuration:

```
awplus-3:
router ipv6 ospf 10
  router-id 192.168.1.1
  area 1 range 2003::/16
```

Now on switch awplus-2 we see the summarization 2003::/16:

```
awplus-2#sh ipv6 route
IPv6 Routing Table
Codes: C - connected, S - static, R - RIP, O - OSPF, B - BGP
Timers: Uptime
C   2002:1111::/64 via ::, vlan2, 06:12:46
O   2003::/16 [110/3] via fe80::200:cdff:fe28:801, vlan2, 00:00:19
C   2004:1111::/64 via ::, vlan4, 05:55:28
C   fe80::/64 via ::, vlan4, 05:55:28
C   fe80::/64 via ::, vlan2, 06:12:46
```

Rather than the two individual routes:

```
awplus-2#sh ipv6 route
IPv6 Routing Table
Codes: C - connected, S - static, R - RIP, O - OSPF, B - BGP
Timers: Uptime
C   2002:1111::/64 via ::, vlan2, 06:17:59
O   2003:1111::/64 [110/2] via fe80::200:cdff:fe28:801, vlan2, 00:00:07
O   2003:2222::/64 [110/3] via fe80::200:cdff:fe28:801, vlan2, 00:00:07
C   2004:1111::/64 via ::, vlan4, 06:00:41
C   fe80::/64 via ::, vlan4, 06:00:41
C   fe80::/64 via ::, vlan2, 06:17:59
```

# Using RIPng and OSPFv3 together

OSPFv3 can redistribute routes that are learnt via RIPng. In this example we will learn an IPv6 route via RIPng and redistribute it into OSPFv3.

**Example** An IPv6 host running RIPng was connected to port1.0.10 (VLAN 10) on switch awplus-3.

This host is advertising a route to 2010:7777::/64 via RIPng

The extra configuration added on awplus-3 is:

```
vlan database
  vlan 10
!
interface port1.0.10
switchport access vlan 10
!
interface vlan10
  ipv6 address 2010:5555::1/64
  ipv6 router rip
!
router ipv6 ospf 10
  router-id 192.168.1.1
  redistribute connected
  redistribute rip
!
```

On switch awplus-3 we now see the route 2010:7777::/64 via RIP:

```
awplus-3#sh ipv6 route
IPv6 Routing Table
Codes: C - connected, S - static, R - RIP, O - OSPF, B - BGP
Timers: Uptime

C   2002:1111::/64 via ::, vlan2, 03:40:32
C   2003:1111::/64 via ::, vlan3, 03:40:32
O   2004:1111::/64 [110/20] via fe80::200:cdf:fe29:8ed3, vlan2, 03:39:33
O   2005:1111::/64 [110/20] via fe80::215:77ff:fec9:745e, vlan3, 03:36:29
C   2010:5555::/64 via ::, vlan10, 00:30:31
R   2010:7777::/64 [120/2] via fe80::200:cdf:fe20:f101, vlan10, 00:23:21
C   fe80::/64 via ::, vlan10, 00:30:31
C   fe80::/64 via ::, vlan3, 03:40:32
C   fe80::/64 via ::, vlan2, 03:40:32
```

And the neighboring switches will see this route via OSPFv3:

```
awplus-2#sh ipv6 route
IPv6 Routing Table
Codes: C - connected, S - static, R - RIP, O - OSPF, B - BGP
Timers: Uptime

C   2002:1111::/64 via ::, vlan2, 03:42:41
O   2003:1111::/64 [110/2] via fe80::200:cdf:fe27:c1ff, vlan2, 03:41:50
C   2004:1111::/64 via ::, vlan4, 03:53:41
O   2005:1111::/64 [110/20] via fe80::200:cdf:fe27:c1ff, vlan2, 03:38:38
O   2010:5555::/64 [110/20] via fe80::200:cdf:fe27:c1ff, vlan2, 00:32:37
O   2010:7777::/64 [110/20] via fe80::200:cdf:fe27:c1ff, vlan2, 00:25:27
C   fe80::/64 via ::, vlan2, 03:42:41
C   fe80::/64 via ::, vlan4, 03:53:41
```

The route learnt by RIP is advertised in OSPFv3 as an AS-external LSA.

### **Distributing OSPFv3 routes into RIPng**

You can also redistribute OSPFv3 learned routes into RIPng (so the IPv6 host will see all of the OSPFv3 routes in the network via RIPng):

```
router ipv6 rip
 redistribute ospf
!
```



# Route-maps and OSPFv3

---

At present, route-maps can match on the following:

- metric type
- metric value
- nexthop address
- prefix-list
- access-list

In OSPFv3 at present AlliedWare Plus only supports applying a route-map to redistributed routes - RIPng, static or connected routes. We do not yet support Distribute-lists, which can be applied when bringing routes from the SPF calculation into the OSPFv3 route table.

If you were learning IPv6 routes from RIPng, and wanted to modify the characteristics of a particular RIPng route and advertise it in OSPFv3, you could create a route-map to do this.

There are multiple options available for what the route-map can match on. The next section considers a few different cases.

## Matching on a prefix-list

1. Configure a prefix-list, in this case **test1**, which matches on IPv6 route **2010:7777::/64**, which is the prefix being learned via RIPng:

```
ipv6 prefix-list test1 seq 5 permit 2010:7777::/64

!
```

2. Create a route-map, in this case **rm1**, and match on the prefix-list. Use the **set** parameters to modify the metric-type and metric:

```
route-map rm1 permit 10

match ipv6 address prefix-list test1

set metric-type type-1

set metric 100

!
```

3. Apply the **rm1** route-map to the **redistribute rip** command

```
router ipv6 ospf 10

router-id 192.168.1.3

redistribute rip route-map rm1

!
```

4. Now in the LSA database, you can see that the AS-External LSA for this imported route to **2010:7777::/64** via OSPFv3 has metric-type of **1** and a metric of **100**:

```
LS age: 57
LS Type: AS-External-LSA
Link State ID: 0.0.45.228
Advertising Router: 192.168.1.3
LS Seq Number: 0x80000001
Checksum: 0x8A47
Length: 36
  Metric Type: 1 (Comparable directly to link state metric)
  Metric: 100
  Prefix: 2010:7777::/64
  Prefix Options: 0 (-|-|-)
```

As such, when other routes receive this LSA, they put this route into their OSPFv3 route table as an E1 route with metric 100.

## Matching on an access-list

1. Configure an access-list, in this case **test3**, which matches on IPv6 route **2010:7777::/64**, which is the prefix that is being learned via RIPng:

```
ipv6 access-list standard test3 permit 2010:7777::/64
!
```

2. Create a route-map, in this case **rm5**, and use the **set** parameters to modify the metric:

```
route-map rm5 permit 10
  match ipv6 address test3
  set metric 35
!
```

3. Apply this route-map (rm5) to the **redistribute rip** command:

```
router ipv6 ospf 10
  router-id 192.168.1.3
  redistribute connected
  redistribute rip route-map rm5
!
```

4. The neighboring switch sees the route to **2010:7777::/64** via OSPFv3 with a metric of **35**:

```
sh ipv6 route 2010:7777::/64
```

```
Routing entry for 2010:7777::/64
```

```
Known via "ospf", distance 110, metric 35, best
```

```
Last update 00:00:36 ago
```

```
* via fe80::200:cdf:fe29:8ed3, vlan5
```

## Matching on route-type

1. Create a route-map, in this case **rm3**, and match on **external type-2**. Use the **set** parameters to set the metric of any external type-2 routes to **50**:

```
route-map rm3 permit 10  
  
match route-type external type-2  
  
set metric 50
```

2. Apply the **rm3** route-map to the **redistribute rip** command:

```
router ipv6 ospf 10  
  
router-id 192.168.1.3  
  
redistribute rip route-map rm3  
  
!
```

The neighboring switch sees any external type-2 routes learned via OSPFv3 advertised with a metric of 50.

## Other possibilities

The other possible route-map matches are as follows:

- Metric:

```
awplus-2(config-route-map)#match metric ?  
<0-4294967295> Metric value
```

- Next-hop address:

```
awplus-2(config-route-map)#match ip next-hop ?  
  
<1-199> IP access-list number  
  
<1300-2699> IP access-list number (expanded range)  
  
WORD IP Access-list name  
  
prefix-list Match entries of prefix-lists
```

# Differences between OSPFv2 and OSPFv3

---

## OSPFv3 runs over a link, rather than a subnet

This allows the exchange of packets between two neighbors who share the same link but belong to different IPv6 subnets. A single link can belong to multiple IPv6 subnets.

You can configure multiple IPv6 addresses on an OSPFv3 link. None of the addresses are defined as secondary addresses, as is the case with IPv4. Two routers are adjacent even if no IPv6 prefix is common between the neighbors, except the link-local address. This is different from OSPFv2 for IPv4, in which OSPFv2 neighbors only become adjacent if the neighbors belong to the same IP subnet.

## Multiple instances per link

There are circumstances in which multiple OSPFv3 routers attached to a single broadcast link do not form a single adjacency among them.

For example: four routers are attached to an Ethernet link. Routers 1 and 2 belong to one OSPF domain, and routers 3 and 4 belong to a different OSPF domain. There are adjacencies between 1 and 2 and between 3 and 4, but not between 1 and 3.

## Neighbors identified by their Router ID

OSPFv2 neighbors on broadcast and NBMA links are identified by their interface addresses, whereas neighbors on other types of links are identified by Router ID. OSPFv3 removes this inconsistency, so that all neighbors on all link types are identified by Router ID.

## Standard authentication mechanisms

IPv6 uses the Authentication extension header as a standard authentication method, so there is no need for a separate OSPFv3 authentication mechanism.

## A link-local flooding scope

OSPFv3 adds the Link LSA for carrying information that is only relevant to neighbors on a single link.

## Uses link-local addresses

OSPFv3 uses a router's link-local IPv6 address as the source address and as next-hop address. This address always begins with FF80::/10.

## Addressing

OSPFv3 Router and Network LSAs do not carry IPv6 addresses. These are carried in the new Inter-Area Prefix LSA. 32 bit Router IDs and LSA IDs are maintained in IPv6. Though these are written in dotted decimal format, they are not IPv4 addresses.

The Router LSA describes the originating router and its links to neighbors. Below is a Router LSA with Router ID 192.168.1.2:

```
LS age: 805
  LS Type: Router-LSA
  Link State ID: 0.0.0.0
  Advertising Router: 192.168.1.2
  LS Seq Number: 0x80000007
  Checksum: 0x5751
  Length: 40
  Flags: 0x02 (-|-|E|-)
  Options: 0x000013 (-|R|-|-|E|V6)
```

```
  Link connected to: a Transit Network
    Metric: 1
    Interface ID: 202
    Neighbor Interface ID: 202
    Neighbor Router ID: 192.168.1.2
```

If we now look at the **show ipv6 ospf database intra-prefix** command output, we can see that the Intra-Prefix LSA for Router ID **192.168.1.2** is advertising prefix **2002:1111::64**:

```
awplus-2#sh ipv6 ospf database intra-prefix
      OSPFv3 Router with ID (192.168.1.2) (Process 10)
          Intra-Area-Prefix-LSA (Area 0.0.0.0)

  LS age: 864
  LS Type: Intra-Area-Prefix-LSA
  Link State ID: 0.0.0.4
  Advertising Router: 192.168.1.2
  LS Seq Number: 0x80000005
  Checksum: 0x468F
  Length: 44
  Number of Prefixes: 1
  Referenced LS Type: 0x2002
  Referenced Link State ID: 0.0.0.202
  Referenced Advertising Router: 192.168.1.2

  Prefix: 2002:1111::/64
  Prefix Options: 0 (-|-|-|-)
  Metric: 0
```

## Two new LSA types

OSPFv3 carries over the seven basic LSA types from OSPFv2. OSPFv3 also introduces two new LSA types: **Link** and **Intra-area Prefix**.

### OSPFv3

0x2001 Router LSA

0x2002 Network LSA

0x2003 Inter-area Prefix LSA

0x2004 Inter-area Router LSA

0x4005 AS-External LSA

0x2006 Group Membership LSA

0x2007 Type-7 LSA

0x0008 Link LSA

0x2009 Intra-area Prefix LSA

### OSPFv2

1 Router LSA

2 Network LSA

3 Network summarization LSA

4 ASBR summarization LSA

5 AS-External LSA

6 Group Membership LSA

7 NSSA External LSA

