AlliedWare Plus™ OS

# How To | Configure Some Basic OSPF Routing Scenarios

## Introduction

OSPF is an Open Standards link-state routing protocol used to exchange routing information between devices dynamically. This document provides an example OSPF network diagram and describes some of the main configuration parameters for OSPF using the AlliedWare Plus Operating System (AW+).

## Contents

## Related How To Notes

You may also find the following AlliedWare Plus OS How To Notes useful:

- *How To Get Started With The AlliedWare PlusTM Operating System*
- *How To Configure Basic Switching Functionality*

## Which products and software version does it apply to?

This How To Note applies to the following Allied Telesis switches, running the AlliedWare Plus OS software version 5.2.1 or later:

- SwitchBlade x908
- x900 AW+ switches
- x600

**List of terms:**

**ACL**
Access Control List

**AD**
Administrative Distance

**LSA**
Link State Advertisement

**Cost**
An indication of the overhead required to send packets across a certain interface.

**Metric**
The sum of all the costs along the path to a given destination.

**Redistribute**
Advertise routes learnt from one routing protocol into another routing protocol

**Wildcard mask**
A subnet mask in which bits set to 0 indicate an exact match and bits set to 1 indicate 'don't care'

**Route-map**
A mechanism for filtering IP routes and changing their attributes

# Configuring OSPF on a network

This section describes how to configure the main parameters of OSPF on AW+. First you will need to enter OSPF configuration mode, then setup your active IP interfaces to run OSPF. Once this is complete you can import routes learned via other protocols, configure administrative distance, default routes, passive interfaces, filter routes, and apply route maps.

Throughout the document, our configurations will relate to the simple OSPF network diagram on page 4. The document culminates in complete configuration scripts for all routing devices in that network.

## Entering OSPF configuration mode

Enter commands one line at a time.

- To **enter** OSPF configuration mode, type in:

```
configure terminal
router ospf 1
```

This is what you would see at the SwitchBlade x908 command prompt:

```
sbx908-awplus#configure terminal
sbx908-awplus(config)#router ospf 1
sbx908-awplus(config-router)#
```

- To **exit** configuration mode - type in Ctrl+Z.

**Note:** The value **1** on the command **router ospf 1** represents the OSPF *instance number*. It is, in fact, possible to run multiple separate instances of OSPF on the same switch, but that is a rare situation that is beyond the scope of this document. For the purposes of this document, you can use any number you like for the instance number. The instance number is local to the switch and has no relation to other switches.

# Example OSPF network diagram

This simple example OSPF network diagram will help you visualise the configurations provided and discussed in this document.

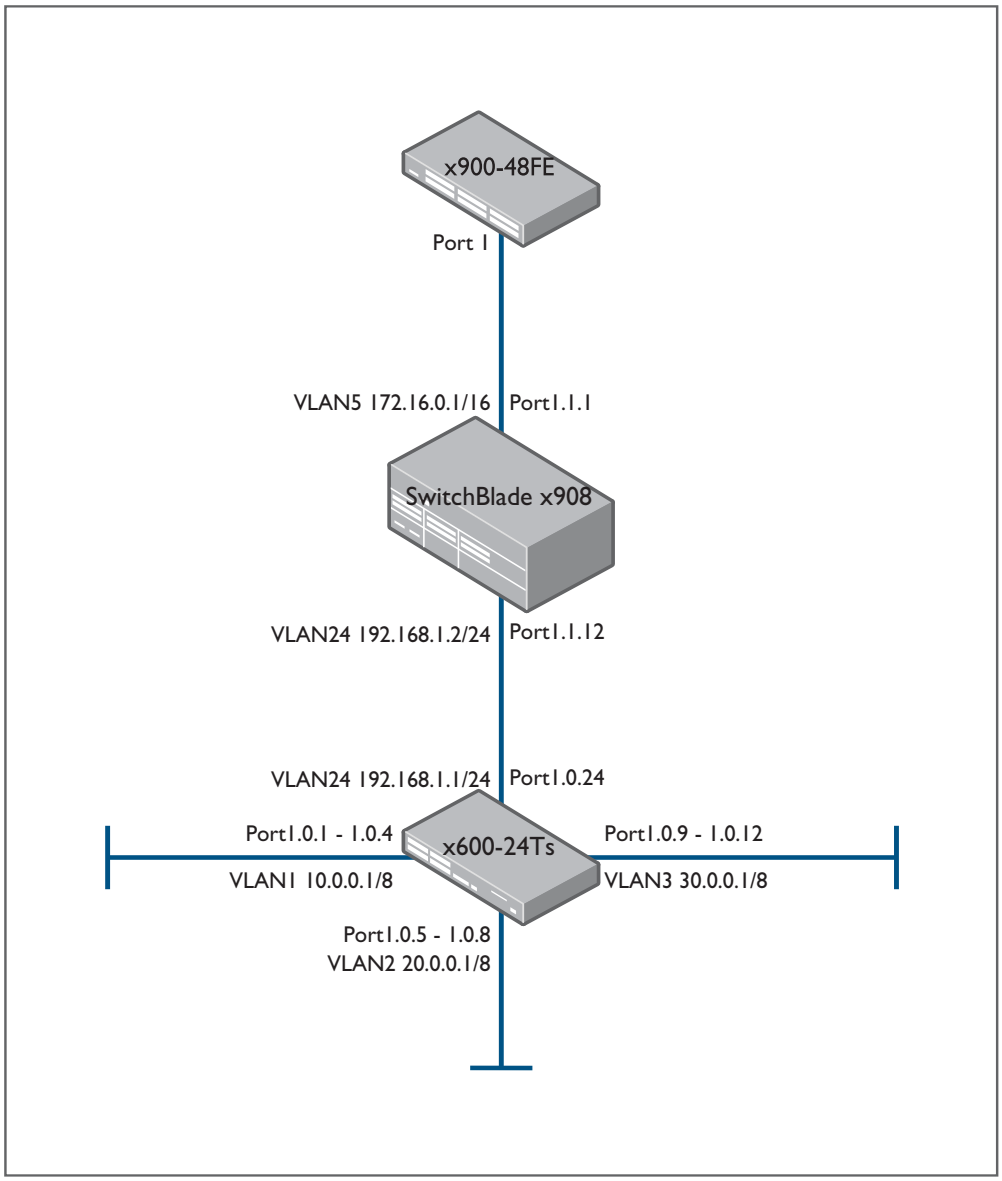


OSPF basic routing.eps

# Using the Network command to define active IP interfaces

The **network** command defines a range of IP addresses. Any IP interface on the switch that has an IP address within the range will run OSPF.

The command includes a **wild card mask** that specifies the range of IP addresses.

1. In the OSPF command, type in:

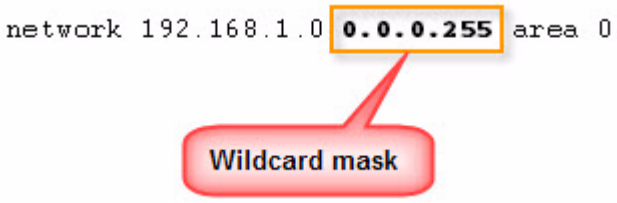


   The network address is 192.168.1.0. The *wildcard mask* specifies how much of the network address must match exactly. Where there is a 0 in the wildcard mask, the corresponding bits in the network address must match exactly, in this case all of the three bytes 192.168.1. Where there is a 255, it doesn't care what this part of the network address is - in this case the last part of the network address can be anything. Essentially this is the same as the more common notation 192.168.1.0/24.

   Note that the network command includes the text **area 0**. This defines which OSPF area the interfaces covered by the network command will belong to. The area number is a 32-bit integer that is typically written in either dotted decimal format a.b.c.d or as an integer like 0 or 1.

2. To enable OSPF on **all** IP interfaces configured on the switch, you can use the command:

```
network 0.0.0.0 255.255.255.255 area x
```

# Redistributing routes

The OSPF **redistribute** command allows routes learned via other protocols/methods to be imported into OSPF and advertised by it. The OSPF implementation in AW+ can redistribute BGP, Connected, RIP, and Static routes.

Let us configure the x600-24Ts, in the simple OSPF network diagram on page 4, to run OSPF on its VLAN24 interface, and redistribute the routes from its other interfaces into OSPF:

```
router ospf 1
redistribute connected
network 192.168.1.0 0.0.0.255 area 0
```

The **`redistribute connected`** command tells the switch to add routes for the directly connected subnets on all the other IP interfaces into OSPF, in this case they will be:

```
10.0.0.0/8
20.0.0.0/8
30.0.0.0/8
```

The x908 will see the following routes via OSPF:

```
sbx908-x600-awplus#show ip route
    O E2    10.0.0.0/8 [110/20] via 192.168.1.1, vlan24, 00:00:08
    O E2    20.0.0.0/8 [110/20] via 192.168.1.1, vlan24, 00:00:08
    O E2    30.0.0.0/8 [110/20] via 192.168.1.1, vlan24, 00:00:08
    B       40.0.0.0/8 [20/0] via 172.16.0.2, vlan5, 00:16:47
    C       172.16.0.0/16 is directly connected, vlan5
    C       192.168.1.0/24 is directly connected, vlan24
```

- The routes shown as O E2 mean these are OSPF external type 2 routes. By default, OSPF will advertise imported routes as external type 2.
- Note that the metric of the routes is 20 (see the [110/20]). This is because the default metric given to redistributed routes is 20 (unless they are redistributed from BGP, in which case the metric is 1).

The equivalent commands to import Static, BGP, or RIP routes would be, respectively:

- `redistribute static`
- `redistribute BGP`
- `redistribute RIP`

As will be seen on page 13 and page 15, it is possible to use filters and route maps to control which routes from a given protocol are redistributed.

## Altering metrics in redistributed routes

You will see that in the route table of the x908, it has a route to 40.0.0.0/8, learnt by BGP. If the command **redistribute BGP** is configured on the x908, then the x600 will learn that route as an external route with a metric 1. Because, routes redistributed from BGP are given metric 1.

## x600 IP route table

```
x600-awplus#show ip route
    C       10.0.0.0/8 is directly connected, vlan1
    C       20.0.0.0/8 is directly connected, vlan2
    C       30.0.0.0/8 is directly connected, vlan3
    O E2    40.0.0.0/8 [110/1] via 192.168.1.2, vlan24, 00:00:11
    O       50.0.0.0/8 [110/2] via 192.168.1.2, vlan24, 00:15:54
    O       172.16.0.0/16 [110/2] via 192.168.1.2, vlan24, 00:15:54
    C       192.168.1.0/24 is directly connected, vlan24
    O E2    200.0.0.0/24 [110/1] via 192.168.1.2, vlan24, 00:00:11
```

On the x908 the metric applied to the BGP routes redistributed into OSPF can be changed to 5 by configuring a metric value on the redistribute command

```
router ospf 1
 redistribute bgp metric 5
```

Now we can see the metric changes on the x600

### x600 IP route table

```
x600-awplus#show ip route
    C       10.0.0.0/8 is directly connected, vlan1
    C       20.0.0.0/8 is directly connected, vlan2
    C       30.0.0.0/8 is directly connected, vlan3
    O E2    40.0.0.0/8 [110/5] via 192.168.1.2, vlan24, 00:00:18
    O       50.0.0.0/8 [110/2] via 192.168.1.2, vlan24, 00:12:25
    O       172.16.0.0/16 [110/2] via 192.168.1.2, vlan24, 00:12:25
    C       192.168.1.0/24 is directly connected, vlan24
    O E2    200.0.0.0/24 [110/5] via 192.168.1.2, vlan24, 00:00:18
```

You will see in the **show IP route** outputs above that the routes are prefixed with codes like C, O, O E2, etc. The full set of such codes and their descriptions is shown in the following table:

### Route learning codes

| Code | Description |
|---|---|
| C | Connected |
| S | Static |
| R | RIP |
| B | BGP |
| O | OSPF |
| IA | OSPF inter area |
| NI | OSPF NSSA external type 1 |
| N2 | OSPF NSSA external type 2 |
| E1 | OSPF external type 1 |
| E2 | OSPF external type 2 |
| * | Candidate default |

# Administrative Distance

The [110/2] beside the route in an IP route table means the Administrative Distance (AD) is 110 (OSPF has an AD of 110) and the Metric of the route is 2.

Examples of the Administrative Distances for the various routing protocols are:

| AD | Protocol |
|---|---|
| 0 | Connected interface |
| 1 | Static route |
| 110 | OSPF |
| 120 | RIP |
| 20 | EBGP |

Administrative Distance is a measure of the trustworthiness of the source of the routing information.

- The lower the AD value, the more reliable the routing protocol is considered.
- The default AD of 110 can be changed with the OSPF **distance** command.
- Different ADs can be defined for different OSPF route types.

  For example:

```
Router OSPF 1
Distance OSPF external 150
Distance OSPF inter-area 120
Distance OSPF intra-area 90
```

- When deciding the best route to a destination, the first criterion is always prefix length. But if two different routing protocols put forward routes to the same destination, with the same prefix length, it will be the route with the lower AD that will be chosen to be installed into the route table.

**Note:** If the number of 255 is specified as a Distance value, the route information is not displayed when using the **show ip route** command. This is because the route is judged as a route which isn't trustworthy. When the Distance value is 255, the route cannot be routing.

# More on the network command

We have seen that the network command is used to define which interfaces OSPF runs on. It is also has the effect of originating routes.

If we change the configuration on the x600-24Ts so that we use the OSPF **network** command to advertise these connected routes, instead of the **redistribute connected** command:

```
router ospf 1
 network 10.0.0.0 0.0.0.255 area 0
 network 20.0.0.0 0.0.0.255 area 0
 network 30.0.0.0 0.0.0.255 area 0
 network 192.168.1.0 0.0.0.255 area 0
```

We get the following output on the x908:

```
sbx908-x600-awplus#show ip route
    O       10.0.0.0/8 [110/2] via 192.168.1.1, vlan24, 00:04:55
    O       20.0.0.0/8 [110/2] via 192.168.1.1, vlan24, 00:04:45
    O       30.0.0.0/8 [110/2] via 192.168.1.1, vlan24, 00:04:45
    B       40.0.0.0/8 [20/0] via 172.16.0.2, vlan5, 00:16:47
    C       172.16.0.0/16 is directly connected, vlan5
    C       192.168.1.0/24 is directly connected, vlan24
```

For code descriptions refer to: "Route learning codes" on page 7

So, we still see the same routes from the x600-24Ts via OSPF, but they show as a type 'O' route, and while the AD is still 110, the metric is now 2. The metric is lower because the routes have not been redistributed into OSPF, but are configured to be advertised by OSPF directly.

The metric of 2 is derived as follows:

- The network command causes the x600 to create a network LSA with metric 1
- When the network LSA is advertised to the x908, the x908 adds on the cost as 100,000,000/<link bandwidth>. In this scenario, the links have 1 Gigabit bandwidth, so the cost is calculated at 0.1, which rounds up to 1. So the x908 gives those routes a metric of 1+1 = 2.

If we only wish to advertise some of our networks, for instance network 10.0.0.0/8 and 20.0.0.0/8 but not 30.0.0.0/8, we would just remove this network from the OSPF configuration:

```
router ospf 1
 network 10.0.0.0 0.0.0.255 area 0
 network 20.0.0.0 0.0.0.255 area 0
 network 192.168.1.0 0.0.0.255 area 0
```

Therefore, you will not see network 30.0.0.0/8 in the x908 routing table:

```
sbx908-x600-awplus#show ip route
    O       10.0.0.0/8 [110/2] via 192.168.1.1, vlan24, 00:09:12
    O       20.0.0.0/8 [110/2] via 192.168.1.1, vlan24, 00:09:02
    B       40.0.0.0/8 [20/0] via 172.16.0.2, vlan5, 00:16:47
    C       172.16.0.0/16 is directly connected, vlan5
    C       192.168.1.0/24 is directly connected, vlan24
```

For code descriptions refer to: "Route learning codes" on page 7

**Note:** Configuring these **`network`** commands will also have the effect of enabling OSPF on the VLAN1, VLAN2, etc. interfaces of the x600, which may not be desirable. "Passive-interfaces" on page 12, explains how to advertise an interface's network as an OSPF internal route without enabling OSPF on that interface.

# Advertising a default route in OSPF

The OSPF configuration command:

```
default-information originate
```

will instruct the switch to advertise a default route 0.0.0.0 into the OSPF domain, if it exists in the IP route table.

For example, in the configuration below, the switch will advertise the default route into OSPF, even though it is not configured with **redistribute static**.

### x908 configuration

```
router ospf 1
 network 192.168.1.0 0.0.0.255 area 0
 redistribute bgp
 distribute-list 2 out bgp
 default-information originate
ip route 0.0.0.0/0 172.16.0.2
```

The x600 would receive the default route as an OSPF E2 route, with metric 1.

### x600 IP route table

x600-awplus#show ip route

```
O*E2    0.0.0.0/0 [110/1] via 192.168.1.2, vlan24, 00:00:29
C       10.0.0.0/8 is directly connected, vlan1
C       20.0.0.0/8 is directly connected, vlan2
C       30.0.0.0/8 is directly connected, vlan3
C       192.168.1.0/24 is directly connected, vlan24
O E2    200.0.0.0/24 [110/1] via 192.168.1.2, vlan24, 00:08:18
```

For code descriptions refer to: "Route learning codes" on page 7

The command:

```
default-information originate always
```

will advertise a default route 0.0.0.0 into the OSPF domain, even if this route does not exist in the IP route table.

# Passive-interfaces

An interface can be configured in OSPF as a passive-interface:

```
passive-interface vlan50
```

This means that the network associated with this interface can be advertised into OSPF using the **network** command, but no OSPF protocol packets will be sent out of this interface.

This command can be used when an interface has no OSPF neighbours, but the network is required to be advertised in OSPF.

### x908 configuration

```
interface vlan50
 ip address 50.0.0.5/8
router ospf 1
 passive-interface vlan50
 network 50.0.0.0 0.255.255.255 area 0
 network 172.16.0.0 0.0.255.255 area 0
 network 192.168.1.0 0.0.0.255 area 0
```

### x600 IP route table

```
x600-awplus#show ip route
    C       10.0.0.0/8 is directly connected, vlan1
    C       20.0.0.0/8 is directly connected, vlan2
    C       30.0.0.0/8 is directly connected, vlan3
    O       40.0.0.0/8 [110/3] via 192.168.1.2, vlan24, 00:00:52
    O       50.0.0.0/8 [110/2] via 192.168.1.2, vlan24, 00:00:32
    O       172.16.0.0/16 [110/2] via 192.168.1.2, vlan24, 00:00:52
    C       192.168.1.0/24 is directly connected, vlan24
```

# Filtering routes in OSPF

On the x908 we want to prevent network 20.0.0.0/8, which is received from the x600 via OSPF, from being installed in the IP route table.

To achieve this, we first create an access list that matches the route 20.0.0.0/8

**On the x908**

```
access-list 1 deny 20.0.0.0 0.255.255.255
access-list 1 permit any
```

Configure an access-list in the range 1-99 (IP standard access-list) or a named IP standard access-list. The access-list 1 above tells the switch to deny route 20.0.0.0/8, but permit all others.

Then apply this ACL as the inbound distribute list on OSPF.

```
router ospf 1
 network 172.16.0.0 255.255.0.0 area 0
 network 192.168.1.0 255.255.255.0 area 0
 distribute-list 1 in
```

The **`distribute-list in`** command instructs the switch to filter routes through access-list 1 before installing them into the IP route table. If the access list denies a route, do not install it into the route table.

As you can see in the **`show ip route`** output below, the x908 does not now have the learned route to network 20.0.0.0/8 in its routing table:

```
sbx908-x600-awplus#show ip route
    O       10.0.0.0/8 [110/2] via 192.168.1.1, vlan24, 00:10:05
    O       30.0.0.0/8 [110/2] via 192.168.1.1, vlan24, 00:10:05
    B       40.0.0.0/8 [20/0] via 172.16.0.2, vlan5, 00:16:47
    C       50.0.0.0/8 is directly connected, vlan50
    C       172.16.0.0/16 is directly connected, vlan5
    C       192.168.1.0/24 is directly connected, vlan24
```

For code descriptions refer to: "Route learning codes" on page 7

**Note:** The network LSA for 20.0.0.0/8 will still be in the OSPF LSA database and advertised to OSPF neighbours.

### Distribute-list out command

The **distribute-list access-list** out command has no effect when used with OSPF. OSPF is a link-state protocol and does not advertise routes to a neighbour but advertises a link-state database. The neighbour determines the routes from the information in the link-state database.

## Filtering routes learned via other methods

Route information learned from BGP, RIP, static routes and/or connected interfaces can be redistributed into OSPF and advertised to OSPF neighbours.

Route information from all of these sources can also be filtered if required.

In this example a route (40.0.0.0/8) learned via BGP from the x900-48FE will be removed before BGP is redistributed into OSPF on the x908. The x908 will not add this route to its LSA database, and so will not advertise it in OSPF to its neighbours.

### x908 configuration

```
access-list 2 deny 40.0.0.0 0.255.255.255
access-list 2 permit any
router ospf 1
 network 192.168.1.0 255.255.255.0 area 0
 redistribute bgp
 distribute-list 2 out bgp
router bgp 65000
 network 172.16.0.0/16
 neighbor 172.16.0.2 remote-as 62000
```

The command **distribute-list 2 out bgp** instructs the switch to apply a filter (access-list 2 which denies network 40.0.0.0/8) to the routes that will be redistributed (imported) from BGP into OSPF. The command **redistribute bgp** instructs the switch to add routes learned via BGP into OSPF.

If the x908 is redistributing the routes 200.0.0.0/8 and 40.0.0.0/8 into OSPF from BGP, then this filter will mean that the only BGP-sourced route learned via OSPF on the x600 will be 200.0.0.0/24.

### The x600 IP route table

```
x600-awplus#show ip route
    C       10.0.0.0/8 is directly connected, vlan1
    C       20.0.0.0/8 is directly connected, vlan2
    C       30.0.0.0/8 is directly connected, vlan3
    C       192.168.1.0/24 is directly connected, vlan24
    O E2    200.0.0.0/24 [110/1] via 192.168.1.2, vlan24, 00:00:10
```

For code descriptions refer to: "Route learning codes" on page 7

The same filtering can be applied in the same way to connected, static or RIP routes. Connected, static and RIP routes can also be redistributed into OSPF with the **redistribute** command.

# Route Maps

Route-maps can match on and modify (set), a number of different attributes of routing information.

We will be looking at the following parameters: **interface**, **ip**, and **metric**.

Route-maps are applied in OSPF using the **redistribute** command and apply to:

- BGP
- RIP
- Static routes
- Connected interfaces

A route-map can be configured to **permit** or **deny**.

- If the match criteria are met for the route map, and the permit keyword is used, the route is redistributed and may be modified by set actions.
- If the match criteria are met for the route map and the **deny** keyword is used, the route is not redistributed.
- If a route passes none of the match criteria in the route map, it is not redistributed as the route map has an implicit **deny all** entry at the end.

## Example 1: Matching on a BGP metric of 0 and modifying it to a metric of 5

**x908**

The route-map **bgp-metric** has just one entry. This entry has a **permit** action and a sequence number of 10. If the BGP metric matches 0, then it sets the metric to 5.

```
route-map bgp-metric permit 10
 match metric 0
 set metric 5
```

In the OSPF configuration, apply route map **bgp-metric** to routes redistributed from BGP.

```
router ospf 1
 redistribute bgp route-map bgp-metric
 passive-interface vlan50
 network 50.0.0.0 0.255.255.255 area 0
 network 172.16.0.0 0.0.255.255 area 0
 network 192.168.1.0 0.0.0.255 area 0
```

**x600**

The routes to 40.0.0.0/8 and 200.0.0.0/24, which have been redistributed from BGP into OSPF on the x908 are now advertised to the x600 with a metric of **5**:

```
x600-awplus#show ip route
    C       10.0.0.0/8 is directly connected, vlan1
    C       20.0.0.0/8 is directly connected, vlan2
    C       30.0.0.0/8 is directly connected, vlan3
    O E2    40.0.0.0/8 [110/5] via 192.168.1.2, vlan24, 00:00:24
    O       50.0.0.0/8 [110/2] via 192.168.1.2, vlan24, 00:05:42
    O       172.16.0.0/16 [110/2] via 192.168.1.2, vlan24, 00:05:42
    C       192.168.1.0/24 is directly connected, vlan24
    O E2    200.0.0.0/24 [110/5] via 192.168.1.2, vlan24, 00:00:24
```

## Example 2: Modifying the metric on BGP routes received from a particular interface

The route-map **bgp-interface** has just one entry. This entry has a **permit** action and a sequence number of 10. If the BGP routes are received on interface VLAN5, then set the metric to **25**.

```
route-map bgp-interface permit 10
 match interface vlan5
 set metric 25
```

In the OSPF configuration, apply route map **bgp-interface** to routes redistributed from BGP:

```
router ospf 1
 redistribute bgp route-map bgp-interface
 passive-interface vlan50
 network 50.0.0.0 0.255.255.255 area 0
 network 172.16.0.0 0.0.255.255 area 0
 network 192.168.1.0 0.0.0.255 area 0
```

**x600**

The routes to 40.0.0.0/8 and 200.0.0.0/24, which have been learnt via VLAN5, and redistributed from BGP into OSPF, on the x908 are now advertised to the x600 with a metric of 25:

```
x600-awplus#show ip route
    C       10.0.0.0/8 is directly connected, vlan1
    C       20.0.0.0/8 is directly connected, vlan2
    C       30.0.0.0/8 is directly connected, vlan3
    O E2    40.0.0.0/8 [110/25] via 192.168.1.2, vlan24, 00:01:44
    O       50.0.0.0/8 [110/2] via 192.168.1.2, vlan24, 00:24:03
    O       172.16.0.0/16 [110/2] via 192.168.1.2, vlan24, 00:24:03
    C       192.168.1.0/24 is directly connected, vlan24
    O E2    200.0.0.0/24 [110/25] via 192.168.1.2, vlan24, 00:01:44
```

**Note:** The metric will still show as 0 for networks 40.0.0.0/8 and 200.0.0.0/24 on the x908 itself. They will be advertised to it's OSPF neighbours with a metric of 25.

**x908**

```
sbx908-awplus#show ip route
    O       10.0.0.0/8 [110/2] via 192.168.1.1, vlan24, 00:45:13
    O       20.0.0.0/8 [110/2] via 192.168.1.1, vlan24, 00:45:13
    O       30.0.0.0/8 [110/2] via 192.168.1.1, vlan24, 00:45:13
    B       40.0.0.0/8 [20/0] via 172.16.0.2, vlan5, 00:41:43
    C       50.0.0.0/8 is directly connected, vlan50
    C       172.16.0.0/16 is directly connected, vlan5
    C       192.168.1.0/24 is directly connected, vlan24
    B       200.0.0.0/24 [20/0] via 172.16.0.2, vlan5, 00:44:29
```

For code descriptions refer to: "Route learning codes" on page 7

## Example 3: Modifying the metric on a specific BGP network

**x908**

Access-list 5 classifies on network 40.0.0.0/8 and access-list 6 classifies on any networks

```
access-list 5 permit 40.0.0.0 0.255.255.255
access-list 6 permit any
```

Route-map **bgp-network** has two entries. The first entry checks for a match on access-list 5 (network 40.0.0.0/8). If a match is found, the metric of this route is set to 35. For all other networks (matching access-list 6), the metric is set to 30 by the second entry of the route-map.

```
route-map bgp-network permit 10
 match ip address 5
 set metric 35


route-map bgp-network permit 20
 match ip address 6
 set metric 30
```

Apply this route-map to the BGP redistribution.

```
router ospf 1
redistribute bgp route-map bgp-network
 passive-interface vlan50
 network 50.0.0.0 0.255.255.255 area 0
 network 172.16.0.0 0.0.255.255 area 0
 network 192.168.1.0 0.0.0.255 area 0
```

**x600**

In the route table on the x600, we see that the route 40.0.0.0/8 has metric 35, and the route 200.0.0.0/24 has metric 30.

```
x600-awplus#show ip route

    C       10.0.0.0/8 is directly connected, vlan1
    C       20.0.0.0/8 is directly connected, vlan2
    C       30.0.0.0/8 is directly connected, vlan3
    O E2    40.0.0.0/8 [110/35] via 192.168.1.2, vlan24, 00:09:39
    O       50.0.0.0/8 [110/2] via 192.168.1.2, vlan24, 01:16:06
    O       172.16.0.0/16 [110/2] via 192.168.1.2, vlan24, 01:16:06
    C       192.168.1.0/24 is directly connected, vlan24
    O E2    200.0.0.0/24 [110/30] via 192.168.1.2, vlan24, 00:00:12
```

For code descriptions refer to: "Route learning codes" on page 7

# Full example configurations

Here are the complete configurations for the switches used in our example network.

**x908**

```
vlan database
 vlan 5,24 state enable
interface port1.1.1
 switchport access vlan 5
interface port1.1.12
 switchport access vlan 24
interface vlan5
 ip address 172.16.0.1/16
interface vlan24
 ip address 192.168.1.2/24
access-list 5 permit 40.0.0.0 0.255.255.255
access-list 6 permit any
route-map bgp-network permit 10
 match ip address 5
 set metric 35
route-map bgp-network permit 20
 match ip address 6
 set metric 30
router ospf 1
 redistribute BGP route-map bgp-network
 network 172.16.0.0 0.0.255.255 area 0
 network 192.168.1.0 0.0.0.255 area 0
 passive-interface vlan5
router BGP 65000
neighbor 172.16.0.2 remote-as 62000
```

**x600-24Ts**

```
vlan database
 vlan 2-3,24 state enable
interface port1.0.5-1.0.8
 switchport access vlan 2
interface port1.0.9-1.0.12
 switchport access vlan 3
interface port1.0.24
```

```
 switchport access vlan 24
interface vlan1
 ip address 10.0.0.1/8
interface vlan2
 ip address 20.0.0.1/8
interface vlan3
 ip address 30.0.0.1/8
interface vlan24
 ip address 192.168.1.1/24
router ospf 1
 redistribute connected
 network 192.168.1.0 0.0.0.255 area 0
```

### x900-48FE (AlliedWare)

```
create vlan="vlan5" vid=5
add vlan="5" port=1

enable ip
set ip asnumber=62000
add ip int=vlan5 ip=172.16.0.2
add ip int=vlan1 ip=40.0.0.1
add ip route=200.0.0.0 int=vlan1 next=40.0.0.2
add bgp peer=172.16.0.1 remote=65000
enable bgp peer=172.16.0.1
add bgp network=40.0.0.0
add bgp network=200.0.0.0
```

USA Headquarters | 19800 North Creek Parkway | Suite 100 | Bothell | WA 98011 | USA | T: +1 800 424 4284 | F: +1 425 481 3895
European Headquarters | Via Motta 24 | 6830 Chiasso | Switzerland | T: +41 91 69769.00 | F: +41 91 69769.11
Asia-Pacific Headquarters | 11 Tai Seng Link | Singapore | 534182 | T: +65 6383 3832 | F: +65 6383 3830
www.alliedtelesis.com



C613-16143-00 REV A