

Open Shortest Path First (OSPF)

FEATURE OVERVIEW AND CONFIGURATION GUIDE

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™.

Contents

Introduction	1
Products and software version that apply to this guide	2
Configuring OSPF on a network.....	2
Entering OSPF configuration mode	2
Example OSPF network diagram.....	3
Using the Network command to define active IP interfaces	4
Redistributing routes.....	4
Further examples of altering the metric applied to redistributed routes.....	6
Administrative Distance	11
More on the network command.....	12
Advertising a default route in OSPF.....	13
Passive-interfaces.....	14
Filtering routes in OSPF	15
Route-maps	18
Full example configurations	22



Products and software version that apply to this guide

This guide applies to AlliedWare Plus™ products that support OSPF, running version **5.2.1** or later.

To see whether your product supports OSPF, see the following documents:

- The product's [Datasheet](#)
- The product's [Command Reference](#)

These documents are available from the above links on our website at alliedtelesis.com.

Feature support may change in later software versions. For the latest information, see the above documents.

Configuring OSPF on a network

This section describes how to configure the main parameters of OSPF on AlliedWare Plus. 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 in "[Example OSPF network diagram](#)" on page 3. 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
service ospf
router ospf 1
```

This is what you would see at the SwitchBlade SBx908 GEN3 command prompt:

```
sbx908gen3#configure terminal
sbx908gen3 (config) #service ospf
sbx908gen3 (config) #router ospf 1
sbx908gen3 (config-router) #
```

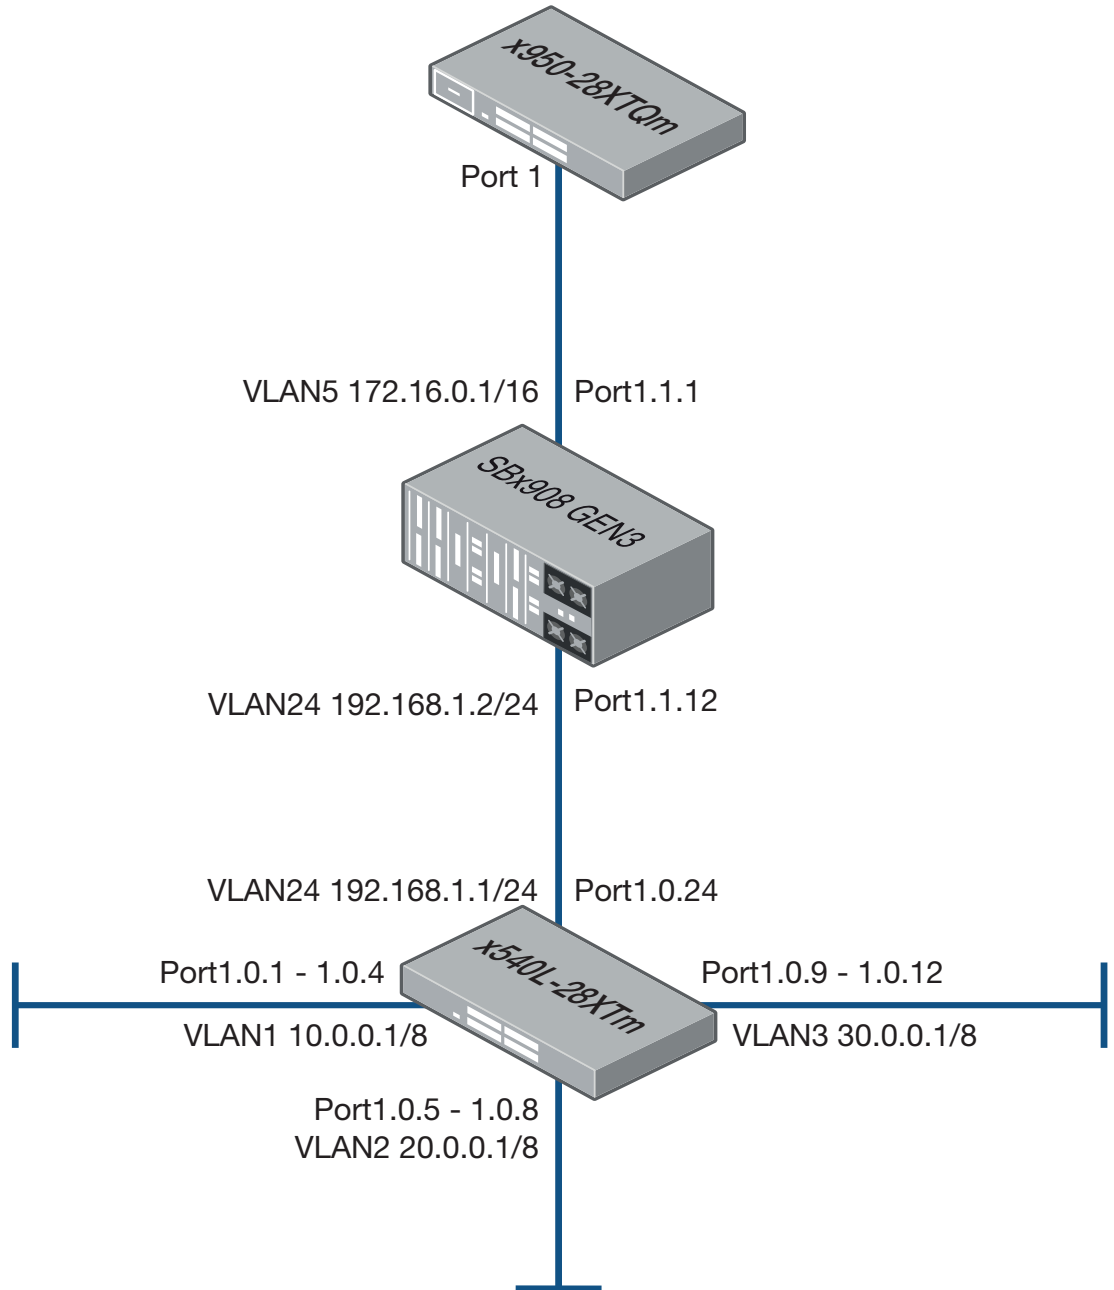
- To exit configuration mode, type in Ctrl+Z.

Note: The value 1 on the command **router ospf 1** represents an OSPF process reference ID. It is in fact possible to run multiple separate OSPF processes on the same switch, but this is a rare situation that is beyond the scope of this document. For the purposes of

this document, you can use any number between 1 and 255 for the process ID. The process ID 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.




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:

```
network 192.168.1.0 0.0.0.255 area 0
```



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 can redistribute BGP, Connected, RIP, and Static routes.

Let us configure the x540L-28XTm, in the simple OSPF network diagram on [page 3](#). We will configure it to run OSPF on its VLAN24 interface and redistribute the routes from its other interfaces into OSPF:

```
service 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 SBx908 GEN3 will see the following routes via OSPF:

```
sbx908gen3#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 in "[Filtering routes in OSPF](#)" on page 15 and "[Route-maps](#)" on page 18, 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 SBx908 GEN3, it has a route to 40.0.0.0/8, learnt by BGP. If the command **redistribute BGP** is configured on the SBx908 GEN3, then the x540L will learn that route as an external route with a metric 1, as routes redistributed from BGP are given metric 1.

x540L IP route table:

```
x5401#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 SBx908 GEN3 the metric applied to the BGP routes redistributed into OSPF can be changed to 5 by configuring a metric value on the redistribute command:

```
service ospf
router ospf 1
```

```
redistribute bgp metric 5
```

Now we can see the metric changes on the x540L.

```
x540L#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:

Table 1: 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

Further examples of altering the metric applied to redistributed routes

No seed metric defined in the redistribute command

When no metric has been defined, the default metric for routes redistributed (from sources other than BGP) into the OSPFv2 process is 20.

In this example, static routes are being redistributed into OSPF, and no metric is configured for the redistributed routes.

```
awplus#show running-config router ospf
!
service ospf
router ospf
  ospf router-id 1.1.1.1
  network 192.168.10.0/24 area 0
  redistribute static
!
```

The metric applied to these routes can be seen by looking at the External LSAs that OSPF creates for the redistributed routes.

```
awplus#show ip ospf database external adv-router 1.1.1.1

          AS External Link States

LS age: 273
Options: 0x2 (-|-|-|-|-|E|-)
LS Type: AS-external-LSA
Link State ID: 172.16.0.0 (External Network Number)
Advertising Router: 1.1.1.1
LS Seq Number: 80000004
Checksum: 0x02f8
Length: 36
Network Mask: /24
    Metric Type: 2 (Larger than any link state path)
    TOS: 0
    Metric: 20
    Forward Address: 0.0.0.0
    External Route Tag: 0
```

Seed metric defined in the redistribute command

In this next example, static routes are being redistributed into OSPF, and a metric value is configured on the **redistribute** command.

```
awplus#show running-config router ospf
!
service ospf
router ospf
  ospf router-id 1.1.1.1
  network 192.168.10.0/24 area 0
  redistribute static metric 100
!
```

The metric in the resulting External LSAs is now 100, which is the value configured in the **redistribute** command.

```
awplus#show ip ospf database external adv-router 1.1.1.1
```

```
AS External Link States
```

```
LS age: 29
Options: 0x2 (-|-|-|-|-|E|-)
LS Type: AS-external-LSA
Link State ID: 172.16.0.0 (External Network Number)
Advertising Router: 1.1.1.1
LS Seq Number: 80000005
Checksum: 0x2386
Length: 36
Network Mask: /24
    Metric Type: 2 (Larger than any link state path)
    TOS: 0
    Metric: 100
    Forward Address: 0.0.0.0
    External Route Tag: 0
```

No seed metric defined in the redistribute command and a route map is used for route filtering

If a route map is used and no seed metric is defined in either the **redistribute** command or set in the route map, the metric defaults to 20.

```
awplus#show running-config router ospf
!
service ospf
router ospf
  ospf router-id 1.1.1.1
  network 192.168.10.0/24 area 0
  redistribute static route-map REDISTRIBUTE_MAP
!
```

```
awplus# show running-config route-map
!
route-map REDISTRIBUTE_MAP permit 10
  match ip address REDISTRIBUTE_LIST
!
```

```
awplus# show ip ospf database external adv-router 1.1.1.1
```

```
AS External Link States
LS age: 53
Options: 0x2 (-|-|-|-|-|E|-)
LS Type: AS-external-LSA
Link State ID: 172.16.0.0 (External Network Number)
Advertising Router: 1.1.1.1
LS Seq Number: 80000008
Checksum: 0xf9fc
Length: 36
Network Mask: /24
    Metric Type: 2 (Larger than any link state path)
    TOS: 0
    Metric: 20
    Forward Address: 0.0.0.0
    External Route Tag: 0
```

Seed metric defined in the redistribute command but not set in route map

If a route map is used and a seed metric is defined in the **redistribute** command only, the metric defaults to the value defined in the **redistribute** command.

```
awplus#show running-config router ospf
!
router ospf
  ospf router-id 1.1.1.1
  network 192.168.10.0/24 area 0
  redistribute static metric 100 route-map REDISTRIBUTE_MAP
!
```

```
awplus#show running-config route-map
!
route-map REDISTRIBUTE_MAP permit 10
  match ip address REDISTRIBUTE_LIST
!
```

Here, we see the value of the metric in the resulting External LSAs is the value configured on the **redistribute** command (100).

```

awplus#show ip ospf database external adv-router 1.1.1.1
AS External Link States
LS age: 61
Options: 0x2 (-|-|-|-|-|E|-)
LS Type: AS-external-LSA
Link State ID: 172.16.0.0 (External Network Number)
Advertising Router: 1.1.1.1
LS Seq Number: 80000009
Checksum: 0x1b8a
Length: 36
Network Mask: /24
    Metric Type: 2 (Larger than any link state path)
    TOS: 0
    Metric: 100
    Forward Address: 0.0.0.0
    External Route Tag: 0

```

Seed metric defined in the redistribute command and also set in the route map

If a route map is used and a seed metric is defined in the **redistribute** command as well as in a route map set command, the metric defaults to the value defined in the **route map set** command.

The route map setting of the metric supersedes the value configured in the **redistribute** command.

In this example, the **redistribute** command is set to redistribute connected routes matching route map REDISTRIBUTE_MAP with a metric of 100. The route map is also set to use metric 200 for the routes matching the REDISTRIBUTE_LIST.

The result is the route map-set metric value of 200 is used.

```

awplus#show running-config router ospf
!
service ospf
router ospf
  ospf router-id 1.1.1.1
  network 192.168.10.0/24 area 0
  redistribute static metric 100 route-map REDISTRIBUTE_MAP
!

```

```

awplus#show running-config route-map
!
route-map REDISTRIBUTE_MAP permit 10
  match ip address REDISTRIBUTE_LIST
  set metric 200
!

```

The metric value in the external LSAs is the value specified in the route map (200).

```

awplus#show ip ospf database external adv-router 1.1.1.1
AS External Link States
LS age: 18
Options: 0x2 (-|-|-|-|-|E|-)
LS Type: AS-external-LSA
Link State ID: 172.16.0.0 (External Network Number)
Advertising Router: 1.1.1.1
LS Seq Number: 8000000a
Checksum: 0x053b
Length: 36
Network Mask: /24
    Metric Type: 2 (Larger than any link state path)
    TOS: 0
    Metric: 200
    Forward Address: 0.0.0.0
    External Route Tag: 0

```

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 listed in [Table 2:](#)

Table 2:

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 also has the effect of originating routes.

If we change the configuration on the x540L-28XTm so that we use the OSPF **network** command to advertise these connected routes, instead of the **redistribute connected** command:

```
service ospf
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 SBx908 GEN3:

```
sbx908gen3-x5401#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: [Table 1 on page 6](#).

So, we still see the same routes from the x540L-28XTm 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 x540L to create a network LSA with metric 1
- When the network LSA is advertised to the SBx908 GEN3, the SBx908 GEN3 adds on the cost as $100,000,000 / \text{link bandwidth}$. In this scenario, the links have 1 Gigabit bandwidth, so the cost is calculated at 0.1, which rounds up to 1. So the SBx908 GEN3 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:

```
service ospf
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 SBx908 GEN3 routing table:

```
sbx908gen3#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: [Table 1 on page 6](#).

Note: Configuring these **network** commands will also have the effect of enabling OSPF on the VLAN1, VLAN2, etc. interfaces of the x540L, which may not be desirable. "[Passive-interfaces](#)" on page 14, 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.

SBx908 GEN3 configuration

```
service ospf
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 x540L would receive the default route as an OSPF E2 route, with metric 1.

x540L IP route table:

```
x5401#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: [Table 1 on page 6](#).

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.

SBx908 GEN3 configuration

```
interface vlan50
  ip address 50.0.0.5/8
service ospf
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
```

x540L IP route table:

```
x5401#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 SBx908 GEN3 we want to prevent network 20.0.0.0/8, which is received from the x540L 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 SBx908 GEN3:

```
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.

```
service 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 SBx908 GEN3 does not now have the learned route to network 20.0.0.0/8 in its routing table:

```
sbx908gen3#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: [Table 1 on page 6](#).

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

This example demonstrates the effect of the **distribute-list access-list out** command for filtering OSPF routes. The x540L switch is configured to filter route 10.0.0.0/8.

```
access-list 1 deny 10.0.0.0
access-list 1 permit any
service ospf
router ospf
 network 192.168.1.0 255.255.255.0 area 0
 redistribute connected
 distribute-list 1 out connected
```

After adding the new configuration (above), the output from the **show ip ospf database** command (below) no longer displays the route 10.0.0.1/8, because it is now filtered out.

```
x5401#show ip ospf database

Router Link States (Area 0.0.0.0)

Link ID      ADV Router    Age  Seq#           CkSum  Link count
172.16.0.1   172.16.0.1    98   0x80000003    0x2930  1
192.168.1.1  192.168.1.1   98   0x80000004    0x0113  2
192.168.1.2  192.168.1.2   97   0x80000003    0xe0bc  1

Net Link States (Area 0.0.0.0)

Link ID ADV Router  Age  Seq#  CkSum
172.16.0.2 192.168.1.1 102 0x80000001 0xa660
192.168.1.2 192.168.1.2 97 0x80000001 0x8425

AS External Link States

Link ID ADV Router  Age  Seq#  CkSum  Route      Tag
20.0.0.0 192.168.1.2 141 0x80000001 0x87b6 E2 20.0.0.0/8 0
30.0.0.0 192.168.1.2 141 0x80000001 0x052f E2 30.0.0.0/8 0
```

And route 10.0.0.0/8 also no longer shows on the SBx908 GEN3 (below) and x950.

```
SBx908gen3#show ip route
Codes: C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       * - candidate default

O E2    20.0.0.0/8 [110/20] via 192.168.1.2, vlan24, 00:22:05
O E2    30.0.0.0/8 [110/20] via 192.168.1.2, vlan24, 00:22:05
C       172.16.0.0/16 is directly connected, vlan5
C       192.168.1.0/24 is directly connected, vlan24
```

SBx908gen3#show ip ospf database

```
Router Link States (Area 0.0.0.0)

Link ID          ADV Router      Age  Seq#           CkSum  Link count
172.16.0.1      172.16.0.1     1538 0x80000002    0x2b2f 1
192.168.1.1     192.168.1.1   1542 0x80000006    0xfc15 2
192.168.1.2     192.168.1.2   1332 0x8000000d    0xcc6 1

Net Link States (Area 0.0.0.0)

Link ID          ADV Router      Age  Seq#           CkSum
172.16.0.2      192.168.1.1   1542 0x80000001    0xa660
192.168.1.2     192.168.1.2   494  0x80000002    0x8226

AS External Link States

Link ID          ADV Router      Age  Seq#           CkSum  Route          Tag
20.0.0.0        192.168.1.2   1331 0x80000001    0x87b6  E2 20.0.0.0/8  0
30.0.0.0        192.168.1.2   1331 0x80000001    0x052f  E2 30.0.0.0/8  0
```

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 x950-28XTm will be removed before BGP is redistributed into OSPF on the SBx908 GEN3. The SBx908 GEN3 will not add this route to its LSA database, and so will not advertise it in OSPF to its neighbours.

SBx908 GEN3 configuration

```
access-list 2 deny 40.0.0.0 0.255.255.255
access-list 2 permit any
service ospf
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 SBx908 GEN3 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 x540L will be 200.0.0.0/24.

The x540L IP route table:

```
x5401#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: [Table 1 on page 6](#).

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

SBx908 GEN3 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.

```
service ospf
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
```

x540L The routes to 40.0.0.0/8 and 200.0.0.0/24, which have been redistributed from BGP into OSPF on the SBx908 GEN3 are now advertised to the x540L with a metric of 5:

```
x5401#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:

```
service ospf
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
```

x540L 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 SBx908 GEN3 are now advertised to the x540L with a metric of 25:

```
x5401#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 SBx908 GEN3 itself. They will be advertised to it's OSPF neighbours with a metric of 25.

SBx908 GEN3

```
sbx908gen3#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 [Table 1 on page 6](#).

Example 3: Modifying the metric on a specific BGP network

SBx908 GEN3 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.

```
service ospf
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
```

x540L In the route table on the x540L, we see that the route 40.0.0.0/8 has metric 35, and the route 200.0.0.0/24 has metric 30.

```
x5401#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: [Table 1 on page 6](#).

Full example configurations

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

Table 3: Full configuration for SBx908 GEN3

```
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
service ospf
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
```

Table 4: Full configuration for x540L-28XTm

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

Table 5: Full configuration for x950-28XTQm

```
vlan database
  vlan 5 name vlan5
interface port1.0.1
  switchport access vlan5
interface vlan1
  ip address 40.0.0.1/8
interface vlan5
  ip address 172.16.0.2/16
ip route 200.0.0.0/8 40.0.0.2
router bgp 62000
  neighbour 172.16.0.1 remote-as 65000
  network 40.0.0.0/8
  network 200.0.0.0/8
```