




How to configure OSPF-ECMP on ProCurve switches



Contents

1. Introduction	2
2. Network diagram	2
3. OSPF configuration.....	4
4. OSPF-ECMP.....	10
5. Firmware versions	12
6. Reference documents.....	12

1. Introduction

This application note explains configuration of OSPF-ECMP on a redundant 3-layer network architecture, in order to provide load-balanced routing capabilities.

2. Network diagram

The platform which will be used in this document to detail the different steps of the configuration consists of:

Two Distribution switches:

- 2 x ProCurve switches 8212zl, named "Distrib-1" and "Distrib-2" are set as Distribution switches. They aggregate multiple Edge switches redundantly.
- Distribution switches act as redundant default gateway for the user VLANs/ IP subnets define between Edge and Distribution switch. For details on this configuration refer to AN-I1 and AN-I2.

Two Core switches:

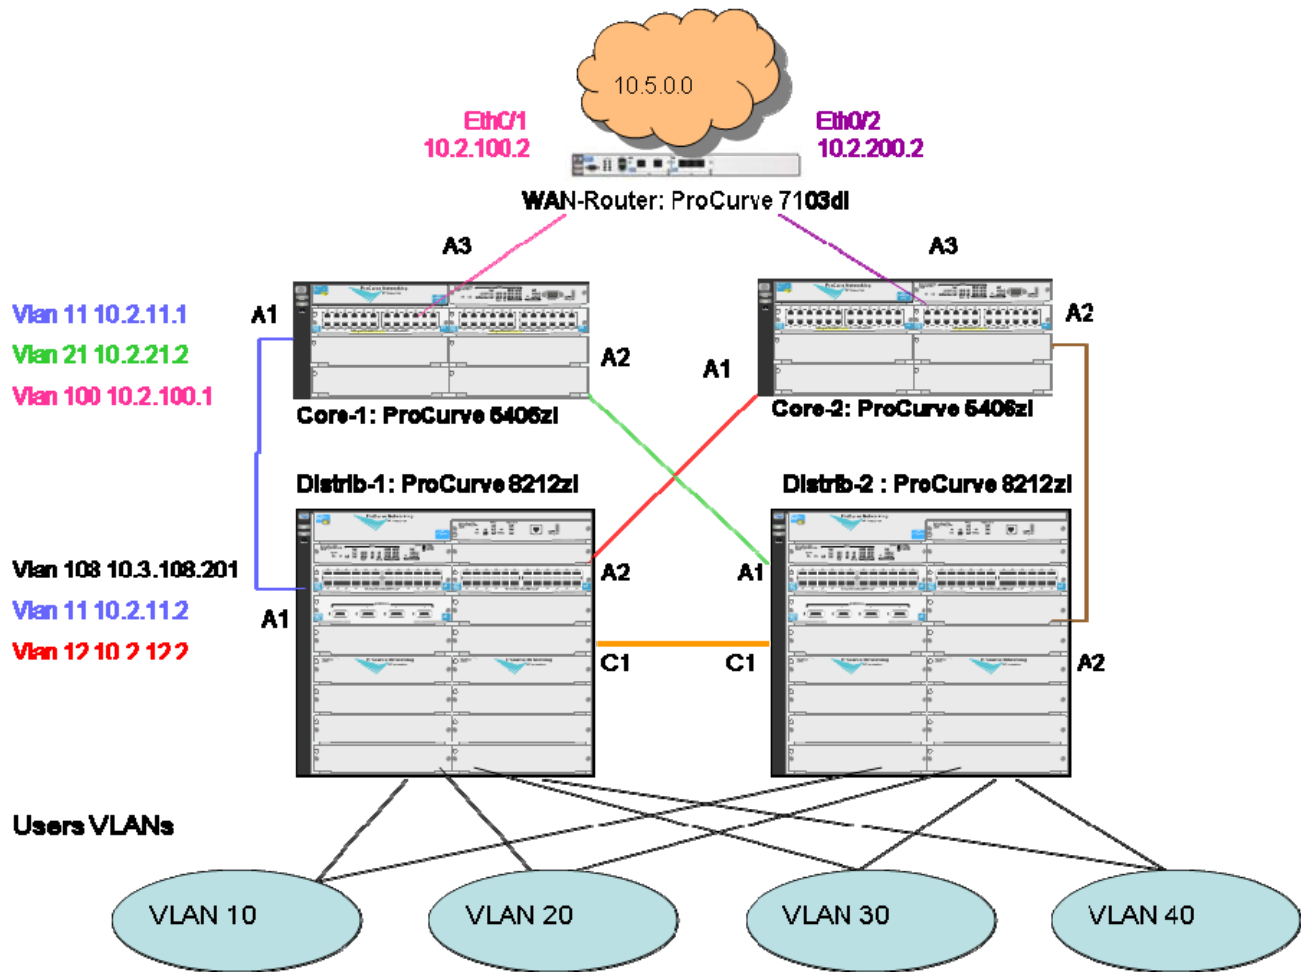
- The 2 Cores "Core-1" and "Core-2" are connected to each of the 2 Distribution switches via Gigabit uplinks. Each link is defined as a unique VLAN and IP Subnet ("Routed Link")

A Router:

- A ProCurve Secure Router 7102dl is redundantly connected to the 2 Core switches:

This hardware configuration is detailed in Figure-1 below.

Figure 1. Network setup



3. OSPF configuration

3.1 Interconnection VLANs/Routed Links

Distribution and Core are interconnected using “Routed Links”.

Why? Here we are pursuing several goals:

- 1- Limit extension of User Broadcast domains (VLANs) to the Edge and Distribution Switches.
- 2- Limit extension of the Spanning Tree Domain to the Edge and Distribution Switches.
- 3- Create redundant links between Core and Distribution.

We set each link between Core and Distribution into a unique VLAN and IP subnet. We also set “BPDU filtering” on those links to block Spanning tree extension.

The WAN-router is connected to 3 IP networks belonging to the IP range 10.5.0.0/16: 10.5.1.0/24, 10.5.2.0/24 and 10.5.3.0/24 that it advertises.

Routed interface configuration on Distrib-1

! VLAN 11 definition

```
Distrib-1(config)# vlan 11
Distrib-1(config-vlan-11)# name To-Core-1
Distrib-1(config-vlan-11)# ip address 10.1.11.2/24
Distrib-1(config-vlan-11)# untagged A1
Distrib-1(config-vlan-11)# exit
```

! VLAN 21 definition

```
Distrib-1(config)# vlan 21
Distrib-1(config-vlan-21)# name To-Core-2
Distrib-1(config-vlan-21)# ip address 10.1.21.2/24
Distrib-1(config-vlan-21)# untagged A2
Distrib-1(config-vlan-21)# exit
```

! BPDU-filtering is set on A1 and 12 to disable BPDUs on this ports and set the limit of spanning-tree domain

```
Distrib-1(config)# spanning-tree A1 bpdu-filter
Distrib-1(config)# spanning-tree A2 bpdu-filter
```

Routed interface configuration on Core-1

```
! IP Routing is enabled
Core-1(config)# ip routing

! VLAN 11 definition
Core-1(config)# vlan 11
Core-1(config-vlan-11)# name To-Dist-1
Core-1(config-vlan-11)# ip address 10.1.11.1/24
Core-1(config-vlan-11)# untagged A1
Core-1(config-vlan-11)# exit

! VLAN 12 definition
Core-1(config)#vlan 12
Core-1(config-vlan-12)# name To-Dist-2
Core-1(config-vlan-12)# ip address 10.1.12.1/24
Core-1(config-vlan-12)# untagged A2
Core-1(config-vlan-12)# exit

! VLAN 100 definition (between Core-1 and WAN-router)
Core-1(config)#vlan 100
Core-1(config-vlan-100)# name VLAN100
Core-1(config-vlan-100)# ip address 10.2.100.1/24
Core-1(config-vlan-100)# untagged A3
Core-1(config-vlan-100)# exit
```

3.2 OSPF topology configuration


OSPF is a link-state routing protocol which defines different types of areas of networks. For more information on OSPF configuration on ProCurve switches, please refer to the Multicast and Routing Guide for the ProCurve Series 3500yl, 6200yl, 5400zl and 8200zl Switches Chapter 5- IP Routing Features:

<http://cdn.procurve.com/training/Manuals/3500-5400-6200-8200-MRG-Jan08-5-IP-Routing.pdf>

In our topology, OSPF has been enabled on Distrib-1, Distrib-2, Core-1 and Core-2 switches and on the WAN-Router. Two areas have been defined:

- Backbone (Area 0) contains Interconnection VLAN/IP Subnet between Core and Distribution and network announced by the WAN router.
- Area 1 contains User VLAN/IP Subnet between Distribution and Edge

Figure2. OSPF Topology



To configure OSPF on Distribution routers, the first step is to define a router-id on each router and to configure them as members of both area0 and area1. Then, area 1 subnets (10.1.10.0/24, 10.1.20.0/24, 10.1.30.0/24 and 10.1.40.0/24) are summarized and advertised as 10.1.0.0/16 in area 0. Finally, OSPF interfaces are set for each VLAN: VLANs 11, 12, 21, 22 are attached to area 0 and VLANs 10, 20, 30 and 40 are attached to area 1.

OSPF Configuration on Distrib-1

```
! Router-id is configured
Distrib-1(config)# ip router-id 0.0.1.1

! OSPF is enabled and areas 0 and 1 are defined
Distrib-1(config)# router ospf
Distrib-1(ospf)#area 0
Distrib-1(ospf)#area 1

! All subnets in area 1 are summarized and advertised as 10.1.0.0/16 into Area 0
Distrib-1(ospf)#area 1 range 10.1.0.0/16

Distrib-1(ospf)#exit

! Set ip ospf interfaces
Distrib-1(config)# vlan 11 ip ospf area 0
Distrib-1(config)# vlan 12 ip ospf area 0
Distrib-1(config)# vlan 10 ip ospf area 1
Distrib-1(config)# vlan 20 ip ospf area 1
Distrib-1(config)# vlan 30 ip ospf area 1
Distrib-1(config)# vlan 40 ip ospf area 1
```

On Core switches, only area 0 and VLANs interfaces in this area have to be defined.

OSPF Configuration on Core-1

```
! Router-id is configured
Core-1(config)# ip router-id 0.0.0.1

! OSPF is enabled and area 0 is defined
Core-1(config)# router ospf
Core-1(ospf)#area 0

! Set ip ospf interfaces
Core-1(config)# vlan 11 ip ospf area 0
Core-1(config)# vlan 21 ip ospf area 0
```

On router 7103dl:

The command Router/OSPF/Default information originate always is used to generate a default route, that is always available. If the internet link fails, another route is automatically advertised.

OSPF Configuration on WAN-Router

```
! Router-id is configured
WAN-Router(config)# ip router-id 0.0.0.5

! OSPF is enabled and area 0 and area 2 are defined, with their networks; redistribution of connected subnets is configured
WAN-Router (config)# router ospf
WAN-Router(config-ospf)# network 10.2.0.0 0.0.255.255 area 0
WAN-Router(config-ospf)# network 10.5.1.0 0.0.0.255 area 2
WAN-Router(config-ospf)# redistribute connected subnets
WAN-Router(config-ospf)#area 0 default-cost 10
WAN-Router(config-ospf)#area 2 default-cost 10

! Default route generation
WAN-Router (config)# default-information-originate always
```

3.3 Verify OSPF configuration

The following commands enable to verify OSPF configuration:

“show ip ospf general” shows the generic OSPF configuration for the router: router ID, Administrative inter-area, intra-area and AS-external distances (default 110), default import metric (by default external type 2), Role (here Area Border and AS Border router), number of external LSAs exchanged.

```
Distrib-1# sh ip ospf general

OSPF General Status
OSPF protocol      : enabled
Router ID         : 0.0.1.1
RFC 1583 compatibility : compatible

Intra-area distance : 110
Inter-area distance : 110
AS-external distance : 110

Default import metric : 10
Default import metric type : external type 2

Area Border      : yes
AS Border       : yes
External LSA Count : 11

External LSA Checksum Sum : 313764
Originate New LSA Count : 0
Receive New LSA Count : 1
```

“show ip ospf area” lists the areas defined on the router and the number and roles of the routers in these areas:

```
Distrib-1# sh ip ospf area

OSPF Area Information

Area ID          Type   Cost  SPFR  ABR  ASBR  LSA  Checksum
-----
backbone         normal 1     158   2    2     0    0x034dd2ec
0.0.0.1          normal 1     158   1    1     0    0x0a860130
```


“ sh ip ospf neighbor” lists the neighbour routers:

```
Distrib-1# sh ip ospf neighbor

OSPF Neighbor Information
```

Router ID	Pri	IP Address	NbIfState	State	Rxmt QLen	Events	Helper Status
0.0.1.2	1	10.1.10.2	BDR	FULL	0	6	None
0.0.1.2	1	10.1.20.2	BDR	FULL	0	6	None
0.0.0.1	1	10.2.11.1	DR	FULL	0	6	None
0.0.0.2	1	10.2.12.1	DR	FULL	0	6	None

“ sh ip ospf interface” lists the interface of the router that are enabled for OSPF

```
Distrib-1# sh ip ospf interface

OSPF Interface Status
```

IP Address	Status	Area ID	State	Auth-type	Cost	Pri	Passive
10.1.10.1	enabled	0.0.0.1	DR	none	1	1	no
10.1.20.1	enabled	0.0.0.1	DR	none	1	1	no
10.1.30.2	enabled	0.0.0.1	DR	none	1	1	yes
10.1.40.2	enabled	0.0.0.1	DR	none	1	1	yes
10.2.11.2	enabled	backbone	BDR	none	1	1	no
10.2.12.2	enabled	backbone	BDR	none	1	1	no

“sh ip route ospf” lists the routes learnt through this protocol and their subtype: IntraArea, InterArea or External

```
Distrib-1# sh ip route ospf
```

IP Route Entries

Destination	Gateway	VLAN	Type	Sub-Type	Metric	Dist.
10.1.100.1/32	10.1.10.2	10	ospf	IntraArea	2	110
10.1.100.1/32	10.1.20.2	20	ospf	IntraArea	2	110
10.2.21.0/24	10.2.11.1	11	ospf	IntraArea	2	110
10.2.22.0/24	10.2.12.1	12	ospf	IntraArea	2	110
10.2.100.0/24	10.2.11.1	11	ospf	IntraArea	2	110
10.2.200.0/24	10.2.12.1	12	ospf	IntraArea	2	110
10.5.1.0/24	10.2.11.1	11	ospf	InterArea	3	110
10.5.1.0/24	10.2.12.1	12	ospf	InterArea	3	110
10.5.2.0/24	10.2.11.1	11	ospf	External2	20	110
10.5.2.0/24	10.2.12.1	12	ospf	External2	20	110
10.5.3.0/24	10.2.11.1	11	ospf	External2	20	110
10.5.3.0/24	10.2.12.1	12	ospf	External2	20	110

Note: Signification of “External2” subtype: 2 types of redistributions can be configured in OSPF:

External2: the cost is the same for both routers and never changes → Load balancing between 2 or several paths

External1: the cost is incremented, and each router uses the closest path: proximity load-balancing. A primary route with a lower cost is defined. By default, External2 is implemented.

4. OSPF-ECMP

OSPF-ECMP (Equal-cost multipath routing) allows OSPF to add routes with multiple next-hop addresses and equal costs to the same destination. Different routes to the same destination must be of the same type: intra-area (destination in the same area), inter-area (destination in other area) or external (destination in other autonomous system). The advantage of this function is to offer load-sharing of the traffic across multiple routes.

4.1 OSPF-ECMP configuration

To configure OSPF-ECMP, enable ip load-sharing and choose a number of equal-costs routes between 2 and 4:

Ip load-sharing 2

By default, ip load-sharing is enabled.

In our configuration, ECMP across 2 routes is enabled on Core Routers:

```
OSPF-ECMP Configuration on Core-1
```

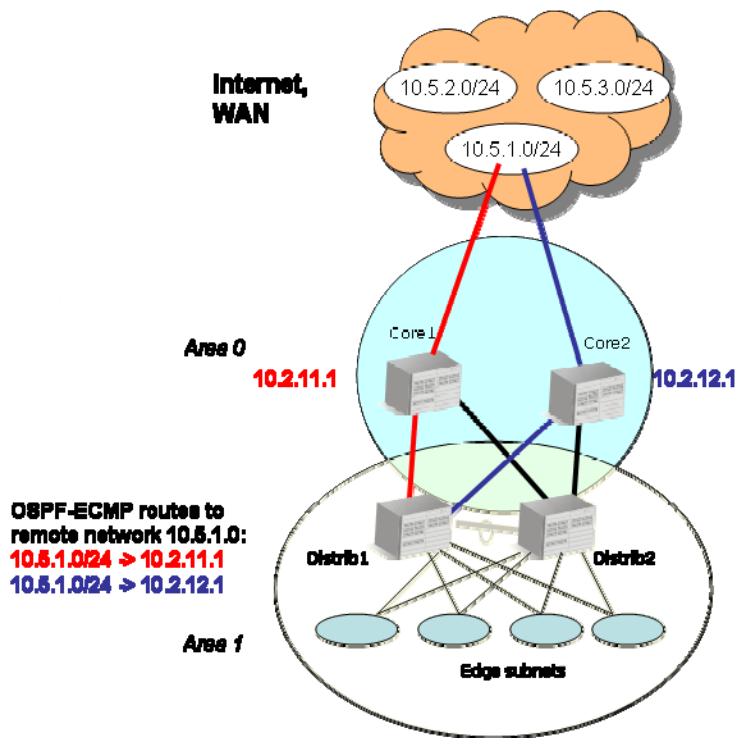
```
! ip load-sharing configuration
Core-1(config)# ip load-sharing
```

This enables the Distribution switches to have 2 routes for the same remote subnet:

For example, the switch Distrib-1 knows 2 equal-cost-routes to each of the 10.5.0.0 subnets:

Destination subnet	Next-hop	Type	Sub-Type	Metric	Administrative Distance
10.5.1.0/24	10.2.11.1	ospf	InterArea	3	110
10.5.1.0/24	10.2.12.1	ospf	InterArea	3	110
10.5.2.0/24	10.2.11.1	ospf	External2	20	110
10.5.2.0/24	10.2.12.1	ospf	External2	20	110
10.5.3.0/24	10.2.11.1	ospf	External2	20	110
10.5.3.0/24	10.2.12.1	ospf	External2	20	110

Figure 3. OSPF-ECMP equal-costs routes from Distrib-1 to 10.5.1.0/24



4.2 OSPF-ECMP operation verification

To verify OSPF-ECMP operation on the Distribution switches, use the command “show ip route ospf”:

OSPF-ECMP routes on Distrib-1						
Distrib-1# show ip route ospf						
IP Route Entries						
Destination	Gateway	VLAN	Type	Sub-Type	Metric	Dist.
0.0.0.0/0	10.2.11.2	11	ospf	External2	10	110
0.0.0.0/0	10.2.21.2	21	ospf	External2	10	110
10.1.100.1/32	10.1.10.2	10	ospf	IntraArea	2	110
10.1.100.1/32	10.1.20.2	20	ospf	IntraArea	2	110
10.2.21.0/24	10.2.11.1	11	ospf	IntraArea	2	110
10.2.22.0/24	10.2.12.1	12	ospf	IntraArea	2	110
10.2.100.0/24	10.2.11.1	11	ospf	IntraArea	2	110
10.2.200.0/24	10.2.12.1	12	ospf	IntraArea	2	110
10.5.1.0/24	10.2.11.1	11	ospf	InterArea	3	110
10.5.1.0/24	10.2.12.1	12	ospf	InterArea	3	110
10.5.2.0/24	10.2.11.1	11	ospf	External2	20	110
10.5.2.0/24	10.2.12.1	12	ospf	External2	20	110
10.5.3.0/24	10.2.11.1	11	ospf	External2	20	110
10.5.3.0/24	10.2.12.1	12	ospf	External2	20	110

5. Firmware versions

Switches firmware versions used in this application note are:

For ProCurve switches 3500yl, 5406zl and 8212zl: K.13.25.

6. Reference documents

This concludes the procedure for configuring OSPF on ProCurve switches.

For further information about how to configure OSPF on ProCurve switches 3500yl-5400zl-8212zl, please refer to the following links:

- ProCurve Multicast and Routing Guide for the ProCurve Series 3500yl, 6200yl, 5400zl and 8212zl series Chapter-5 IP Routing Features: <http://cdn.procurve.com/training/Manuals/3500-5400-6200-8200-MRG-Jan08-5-IP-Routing.pdf>

- Command Line Reference Guide: <http://cdn.procurve.com/training/Manuals/8200-6200-5400-3500-CLI-K13-Mar2008.pdf>

For further information, please visit www.procurve.eu



© 2008 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.