OSPF Exercise

Campus Network Design Workshop

Contents

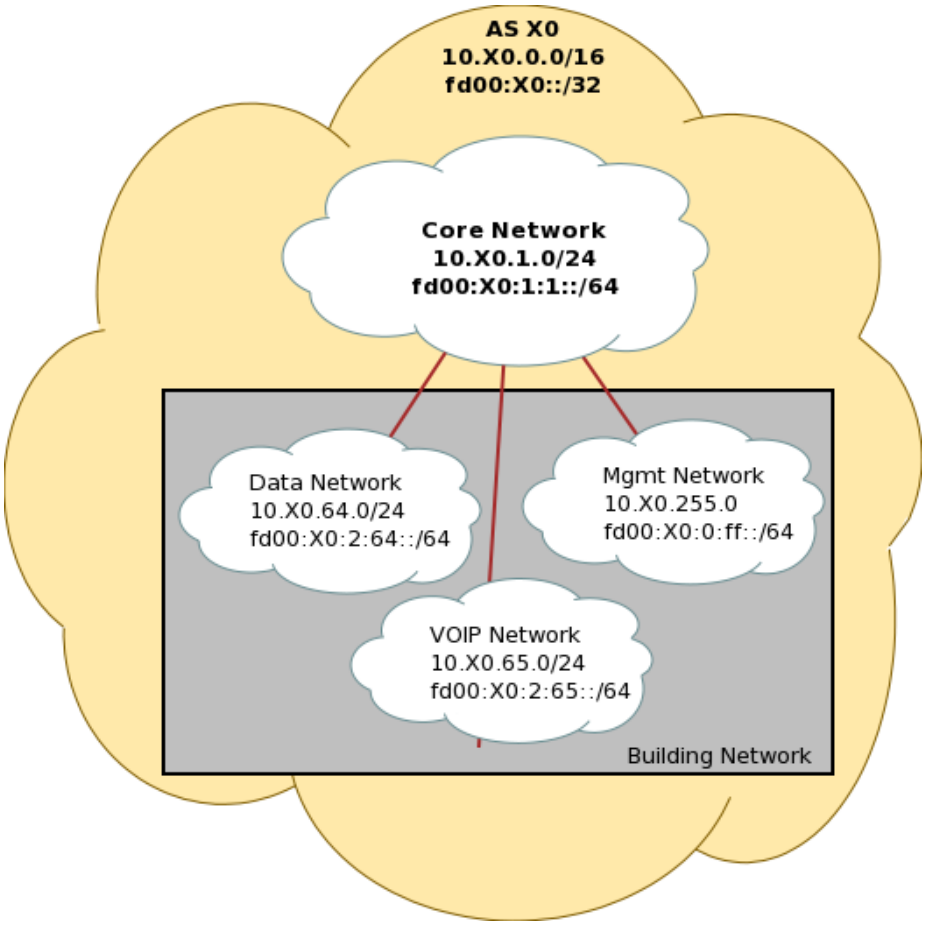

Figure 1: Logical Topology

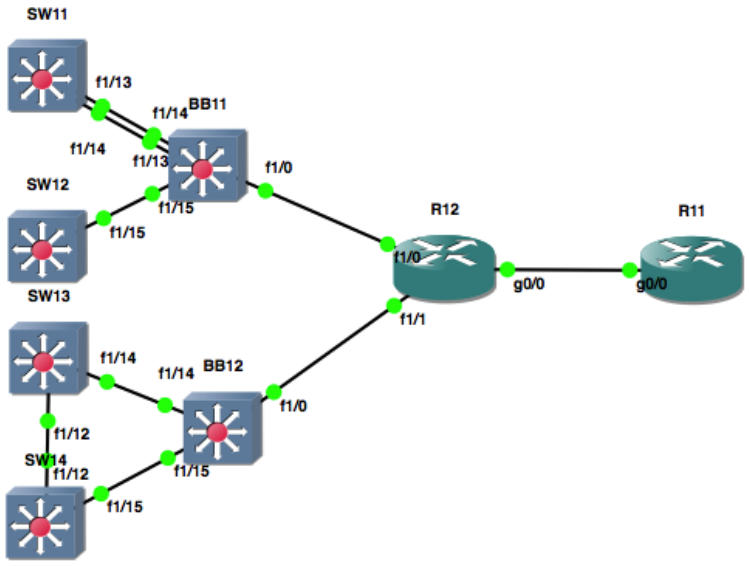

Figure 2: Physical Topology

1 Introduction

The purpose of this exercise is to learn how to configure OSPF on a group of Cisco routers so that they exchange network reachability information and maintain their own routing tables dynamically.

We've already configured two sets of switches with VLANS to represent Engineering and Computer Science (ECS) and the Library (LIB) and here we'll see how we connected those to the campus Core (or backbone) and then to the Internet.

All participants will work within a group as a team. Each group has two routers and six switches to work with. There is a certain dependency between the labs as the exercises progress. Make sure to maintain your configuration unless otherwise instructed. All exercises will use a common IP addressing scheme and network topology. As you go through the exercises all the examples are given from the point of view of R11, the border router in group 1.

Make sure to take the examples and adapt them to your own router, network topology and addressing scheme.

1.1 Router types used in the lab

Cisco 7206 VXR

1.2 Address Space Allocation

Group	IPv4 Block	IPv6 Block
1	10.110.0.0/16	fd00:110::/32
2	10.120.0.0/16	fd00:120::/32
3	10.130.0.0/16	fd00:130::/32
4	10.140.0.0/16	fd00:140::/32
5	10.150.0.0/16	fd00:150::/32
6	10.160.0.0/16	fd00:160::/32

Each group will then further partition their space like this:

IPv4	IPv6	Description
10.1X0.1.0/24	fd00:1X0:1:1::/64	Core Network
10.1X0.10.0/24	fd00:1X0:a::/64	Router Loopback Subnet
10.1X0.64.0/24	fd00:1X0:2:64::/64	ECS Data Subnet (VLAN 64)
10.1X0.65.0/24	fd00:1X0:2:65::/64	ECS VOIP Subnet (VLAN 65)
10.1X0.254.0/24	fd00:1X0:2:FE::/64	ECS MGMT Subnet (VLAN 254)
10.1X0.74.0/24	fd00:1X0:2:74::/64	LIB Data Subnet (VLAN 74)
10.1X0.75.0/24	fd00:1X0:2:75::/64	LIB VOIP Subnet (VLAN 75)
10.1X0.255.0/24	fd00:1X0:2:FF::/64	LIB MGMT Subnet (VLAN 255)

With X being your group number (1,2,3,4,5)

1.2.1 Using private address space

We are using private address space for the exercises - we recommend using public address space in campus networks wherever possible.

2 Exercises

2.1 Basic Router Configuration

We've already configured each router with following basic setup as we did for the switches. For example for R11:

```
hostname R11
!
aaa new-model
aaa authentication login default local
aaa authentication enable default enable
username nsrc secret nsrc
enable secret nsrc
service password-encryption
line vty 0 4
 transport preferred none
 line console 0
 transport preferred none
!
no logging console
logging buffered 8192 debugging
no ip domain-lookup
ipv6 unicast-routing
```

2.2 Interface Configuration

Configure each router's interface according to the diagram (where **X** represents your group):

2.2.1 RX1

Using address .1 for RX1:

```
interface loopback 0
 ip address 10.1X0.10.1 255.255.255.255
 ipv6 address fd00:1X0:a:1/128
!
interface GigabitEthernet1/0
 ip address 10.1X0.1.1 255.255.255.0
 description Link to Core
 ipv6 address fd00:1X0:1:1::1/64
 no ip redirects
```

```
 no ip directed-broadcast
 no ip proxy-arp
 no shutdown
```

2.2.2 RX2

Using address .2 for RX2:

Configure the loopback and Gigabit interfaces based on the setup for RX1.

On the access side, where you will connect the VLANs already configured on the switches:

RX2:

```
interface Fast1/0
 description Engineering and Computer Science (ECS)
 no ip address
 no shutdown
!
interface Fast1/0.64
 encapsulation dot1Q 64
 ip address 10.1X0.64.2 255.255.255.0
 description ECS DATA vlan 64
 ipv6 enable
 ipv6 address fd00:1X0:2:64::2/64
 no ip redirects
 no ip directed-broadcast
 no ip proxy-arp
 no shutdown
```

Do the same for ECS VLANs 65 and 254.

```
interface Fast1/1
 description Library (LIB)
 no ip address
 no shutdown
!
interface Fast1/1.74
 encapsulation dot1Q 74
 ip address 10.1X0.74.2 255.255.255.0
 description LIB DATA vlan 74
 ipv6 enable
 ipv6 address fd00:1X0:2:74::2/64
 no ip redirects
 no ip directed-broadcast
```

```
no ip proxy-arp
no shutdown
```

Do the same for Library VLANs 75 and 254.

2.3 Testing Connectivity - part 1

From R11, ping your neighbor routers:

```
R11# ping 10.110.1.2
R11# ping fd00:110:1:1::2
```

And verify the output of the following commands:

```
show arp           : Shows ARP cache
show interface <int>       : Shows interface state and configuration
show ip interface      : Shows interface IP state and config
show ipv6 neighbors    : Shows IPv6 neighbors
show ipv6 interface <int>  : Shows interface state and configuration
show ipv6 interface        : Shows interface state and configuration
```

Now try pinging these other addresses in your network:

```
R11# ping 10.110.0.2
R11# ping 10.110.64.2
R11# ping 10.110.65.2
R11# ping 10.110.254.2
R11# ping 10.110.74.2
R11# ping 10.110.75.2
R11# ping 10.110.255.2
R11# ping fd00:110:0:0::2
R11# ping fd00:110:2:64::2
R11# ping fd00:110:2:65::2
R11# ping fd00:110:2:74::2
R11# ping fd00:110:2:75::3
R11# ping fd00:110:2:FE::2
R11# ping fd00:110:2:FF::2
```

What is happening? Why can we not ping some of the addresses?

2.4 Dynamic Routing with OSPF

1. Configure a new OSPF routing process and configure OSPF on the interfaces **where adjacencies need to be established**, and also on **any interface that needs to have its subnets advertised by OSPF**.

In the case of R12, this includes the sub-interfaces for VLANs 64, 65, 74, 75, 254 and 255. Notice that we are configuring authentication for the OSPF adjacencies. This is important.

Notice that we will use the number "10" as the OSPF process number for routers R11, R12. This number is local to the router, so it doesn't need to match the process number of a neighboring router. However, it is strongly recommended that you use the same number throughout your network. Most people use their Autonomous System number (although OSPF has nothing to do with the BGP AS).

On R11:

```
router ospf 10
 log-adjacency-changes
 passive-interface default
 area 0 authentication message-digest
 no passive-interface GigabitEthernet0/0
 auto-cost reference-bandwidth 1000
 default-information originate
!
ipv6 router ospf 10
 log-adjacency-changes
 passive-interface default
 no passive-interface GigabitEthernet0/0
 area 0 authentication ipsec spi 256 md5 0123456789ABCDEF0123456789ABCDEF
 auto-cost reference-bandwidth 1000
!
interface Loopback0
 ip ospf 10 area 0
 ipv6 ospf 10 area 0
!
interface GigabitEthernet0/0
 ip ospf 10 area 0
 ip ospf authentication-key nsrc
 ipv6 ospf 10 area 0
```

and on R12:

```
router ospf 10
```

```
 log-adjacency-changes
 passive-interface default
 area 0 authentication message-digest
 no passive-interface GigabitEthernet0/0
 auto-cost reference-bandwidth 1000
!
ipv6 router ospf 10
 log-adjacency-changes
 passive-interface default
 no passive-interface GigabitEthernet0/0
 area 0 authentication ipsec spi 256 md5 0123456789ABCDEF0123456789ABCDEF
 auto-cost reference-bandwidth 1000
!
interface Loopback0
 ip ospf 10 area 0
 ipv6 ospf 10 area 0
!
interface GigabitEthernet0/0
 ip ospf 10 area 0
 ip ospf authentication-key nsrc
 ipv6 ospf 10 area 0

interface FastEthernet1/0.64
 ip ospf 10 area 0
 ip ospf authentication-key nsrc
 ipv6 ospf 10 area 0
```

and repeat for the VLAN interfaces for 65, 74, 75, 254 and 255

2.5 STOP – Checkpoint.

```
show ip ospf neighbor      : show adjacencies
show ip route              : show routes in routing table
show ip ospf               : shows general OSPF information
show ip ospf interface     : shows the status of OSPF in an interface
show ipv6 ospf neighbor
show ipv6 route
show ipv6 ospf
show ipv6 ospf interface
```

We set out to:

Configure a new OSPF routing process and configure OSPF on the interfaces **where adjacencies need to be established**, and also on **any interface that needs to have its subnets advertised by OSPF**.

Have we done that?

2.6 Verifying the routing table in R11

```
R11# show ip route ospf
```

How many routes do you have on R11 for each access network?

2.7 Testing connectivity - part 2

Now try pinging these addresses in your network:

```
R11# ping 10.110.0.2
R11# ping 10.110.64.2
R11# ping 10.110.65.2
R11# ping 10.110.254.2
R11# ping 10.110.74.2
R11# ping 10.110.75.2
R11# ping 10.110.255.2
R11# ping fd00:110:0:0::2
R11# ping fd00:110:2:64::2
R11# ping fd00:110:2:65::2
R11# ping fd00:110:2:74::2
R11# ping fd00:110:2:75::2
R11# ping fd00:110:2:FE::2
R11# ping fd00:110:2:FF::2
```

Do things work differently from the last time you tried this?

Can you ping the management interfaces of the switches you configured earlier?