Advanced Routing Workshop

Basic Routing Lab

Contents

1 Introduction

The purpose of this exercise is to:

- Configure the basics of a Cisco router
- Enable OSPF to exchange internal routing information
- Configure static routing towards a service provider

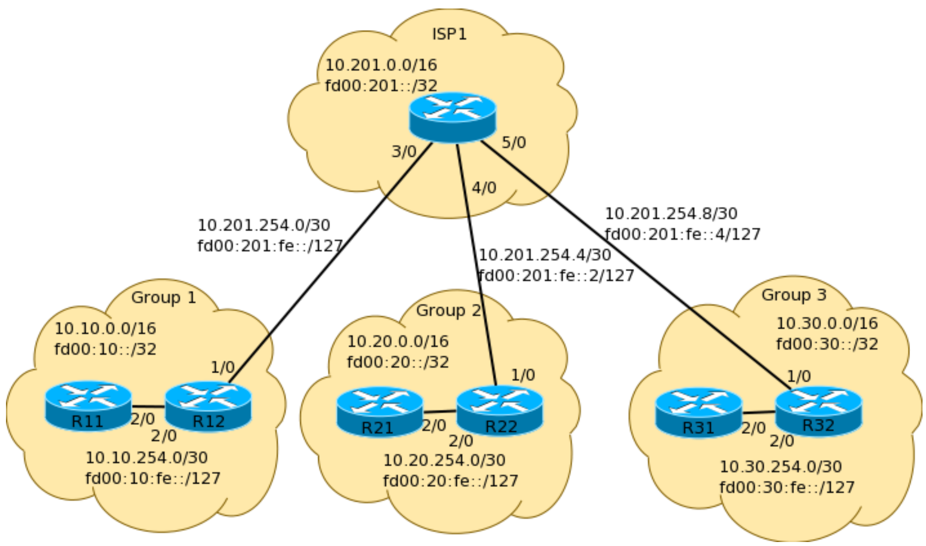

Figure 1: Physical Topology - Module 1

The network configuration is designed to be modular to allow the lab to grow as needed depending on the number of partipants. Each module will contain 1 ISP and 3 customer networks (universities, etc). Modules will be interconnected (see Fig. 3)

2 Logistics

Each participant will be assigned to a network. Depending on the number of participants, either a single person or a group will be responsible for the configuration of a router. You may be asked to rotate and work on a different router so that you have the opportunity to understand the network from another point of view.

As you go through the exercises, you will see examples of configurations for one or more routers. **Make sure to take those examples and adapt them to**

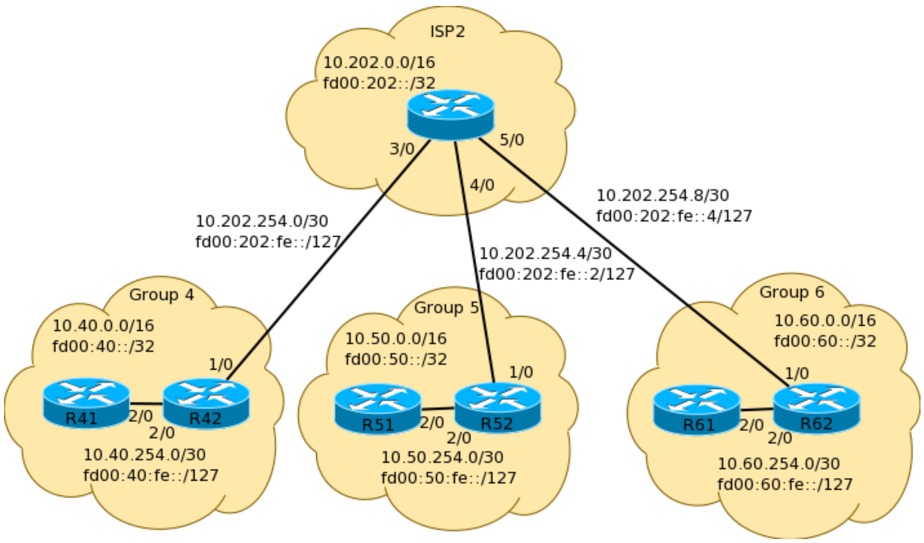

Figure 2: Physical Topology - Module 2

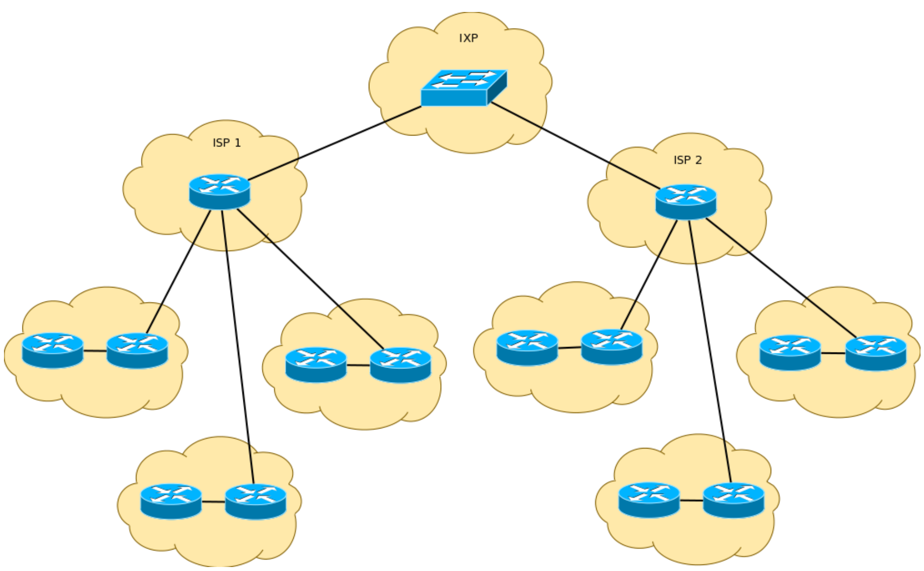

Figure 3: Topology with 2 modules

your own router, network topology and addressing scheme. Use the diagrams to guide you.

Refer to the *Lab Access Instructions* document for information about logging into the routers that have been assigned to you.

3 Address Space Allocation

3.1 End networks (universities, etc)

Group	IPv4	IPv6	ASN
1	10.10.0.0/16	fd00:10::/32	10
2	10.20.0.0/16	fd00:20::/32	20
3	10.30.0.0/16	fd00:30::/32	30

The list will continue in the same pattern if there are more groups.

Each group will then further partition their space as follows:

IPv4	IPv6	Description
10.X0.0.0/17	fd00:X0::/40	End user space
10.X0.254.0/24	fd00:X0:fe::/64	Point-to-point links
10.X0.255.0/24	fd00:X0:ff::/64	Router loopbacks

Where X is your group number (1,2,3…)

Prefixes for point to point links will be of length /30 for IPv4 and /127 for IPv6 (we will adopt the recommendations of RFC6164 for IPv6 inter-router links):

IPv4	IPv6	Description
10.X0.254.0/30	fd00:X0:fe::/127	P2P #1
10.X0.254.4/30	fd00:X0:fe::2/127	P2P #2
10.X0.254.8/30	fd00:X0:fe::4/127	P2P #3

… and so on.

Router loopback addresses will be of size /32 for IPv4 and /128 for IPv6:

IPv4	IPv6	Description
10.X0.255.1/32	fd00:X0:ff::1/128	RX1 Loopback
10.X0.255.2/32	fd00:X0:ff::2/128	RX2 Loopback

3.2 Commercial Internet Service Providers (ISPs)

ISP	IPv4	IPv6	ASN
1	10.201.0.0/16	fd00:200::/32	201
2	10.202.0.0/16	fd00:201::/32	202

... and so on.

3.3 Internet Exchange Points (IXPs)

IXP	IPv4	IPv6
1	10.251.1.0/24	fd00:251:1::/64

4 Basic Router Configuration

1. Name the router

```
enable
config terminal
hostname R11
```

2. Configure Authentication

```
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
```

3. Configure logging

```
no logging console
logging buffered 8192 debugging
```

4. Disable DNS resolution

```
no ip domain-lookup
```

5. Make sure the router understands CIDR. This is the default setting in recent IOS versions, but just in case.

```
ip subnet-zero
ip classless
```

6. Disable source routing

```
no ip source-route
```

7. Activate IPv6 routing

```
ipv6 unicast-routing
```

8. Exit configuration mode and save

```
end
write memory
```

9. Configure your interfaces according to the diagram

Notice that for the links to the ISP we will use the ISP's addresses, while for internal links we use internal addresses.

On R11:

```
interface GigabitEthernet2/0
 description P2P Link to R12
 ip address 10.10.254.1 255.255.255.252
 no ip directed-broadcast
 no ip redirects
 no ip proxy-arp
 ipv6 address fd00:10:fe::/127
 ipv6 nd ra suppress
 no shutdown
!
```

On R12:

```
interface GigabitEthernet1/0
 description P2P Link to ISP1
 ip address 10.201.254.2 255.255.255.252
 no ip directed-broadcast
 no ip redirects
 no ip proxy-arp
 ipv6 address fd00:201:fe::1/127
 ipv6 nd ra suppress
 no shutdown
!
interface GigabitEthernet2/0
 description P2P Link to R11
 ip address 10.10.254.2 255.255.255.252
 no ip directed-broadcast
 no ip redirects
 no ip proxy-arp
 ipv6 address fd00:10:fe::1/127
 ipv6 nd ra suppress
 no shutdown
```

Explanations for some of the above commands:

no ip directed-broadcast

An IP directed broadcast is an IP packet whose destination address is a valid broadcast address for some IP subnet, but which originates from a node that is not itself part of that destination subnet.

Because directed broadcasts, and particularly Internet Control Message Protocol (ICMP) directed broadcasts, have been abused by malicious persons, we recommend disabling the ip directed-broadcast command on any intereface where directed broadcasts are not needed (probably all).

no ip proxy-arp

Proxy ARP is the technique in which one host, usually a router, answers ARP requests intended for another machine. By "faking" its identity, the router accepts responsibility for routing packets to the "real" destination. Proxy ARP can help machines on a subnet reach remote subnets without the need to configure routing or a default gateway.

Disadvantages of proxy arp:

- It increases the impact of ARP spoofing, in which a machine claims to be another in order to intercept packets.
- It hides network misconfigurations in hosts

- Hosts will have larger ARP tables

no ip redirects

ICMP redirects can be sent to a host when the router knows that another router in the same subnet has a better path to a destination. If a hacker installs a router in the network that causes the legitimate router to learn these ilegitimate paths, the hacker's router will end up diverting legitimate traffic thanks to ICMP redirects. Thus, we recommend that you disable this feature in all your interfaces.

ipv6 nd ra supress

IPv6 router advertisements are sent periodically by routers to inform hosts that the router is present, and to allow hosts to autoconfigure themselves using stateless autoconfiguration mechanisms. This is not necessary on point-to-point interfaces.

10. Do some PING tests

```
R12# ping 10.10.254.1          <- R11
R12# ping fd00:10:fe::0        <- R11
R12# ping 10.201.254.1         <- ISP1
R12# ping fd00:201:fe::0       <- ISP1
```

and then verify the output of the following commands:

```
show arp                    : Show ARP cache
show interface <int>        : Show interface state and config
show ip interface           : Show interface IP state and config
show ipv6 neighbors         : Show IPv6 neighbors
show ipv6 interface <int>   : Show interface state and config
show cdp neighbors          : Show neighbors seen via CDP
```

11. Create Loopback interface

On R11:

```
interface loopback 0
 ip address 10.10.255.1 255.255.255.255
 ipv6 address fd00:10:ff::1/128
```

do the same for R12 (obviously, using different addresses).

12. Verify and save the configuration.

```
show running-config
write memory
```

5 Routing

5.1 OSPF

1. Try pinging the loopback addresses of your neighbor

```
R11# ping 10.10.255.2          <- R12 loopback
R11# ping fd00:10:ff:2         <- R12 loopback
```

Q. What is happening?

2. Configure a new OSPF routing process.

Notice that we will use the number "10" as the OSPF process number for routers R11 and 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 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 ASN).

On R11 and R12:

```
router ospf 10
 log-adjacency-changes
 passive-interface default
!
ipv6 router ospf 10
 log-adjacency-changes
 passive-interface default
 area 0 authentication ipsec spi 256 md5 0123456789ABCDEF0123456789ABCDEF
```

3. Now configure OSPF on the interfaces where adjacencies need to be established:

On R11 and R12:

```
interface GigabitEthernet2/0
 ip ospf 10 area 0
 ip ospf authentication message-digest
 ip ospf authentication-key N$RC
 ip ospf network point-to-point
 ipv6 ospf 10 area 0
 ipv6 ospf network point-to-point
```

Notice two things:

a) We are configuring authentication to have control over who becomes an adjacent router and protect against ilegitimate routing information. We configure authentication per interface in IPv4 and per area in IPv6, just because it is simpler that way.

b) We use the "network point-to-point" statement because we are using point-to-point links over a broadcast network. There is no reason for OSPF to elect a Designated Router (DR) and Backup Designated Router (BDR).

Then, since we have configured OSPF to make all interfaces passive by default (recommended!), we need to explicitly activate the *OSPF Hello* function on the interfaces where routing information needs to be exchanged:

R11 and R12:

```
router ospf 10
 no passive-interface GigabitEthernet2/0
!
ipv6 router ospf 10
 no passive-interface GigabitEthernet2/0
```

Now configure OSPF on any interface that needs to have its subnets advertised by OSPF, if OSPF has not been enabled already:

On R11 and R12:

```
interface Loopback0
 ip ospf 10 area 0
 ipv6 ospf 10 area 0
```

5. STOP. Checkpoint

Now try the following show commands:

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

show ipv6 ospf neighbor
show ipv6 route
show ipv6 ospf
show ipv6 ospf interface
```

Repeat the last ping tests.

Q. Can you ping the loopback address of the neighboring router now?

5.2 Static default routes

1. Configure static default routes to reach the outside world.

On R11:

```
ip route 0.0.0.0 0.0.0.0 10.10.254.2
ipv6 route ::/0 fd00:10:fe::1
```

On R12:

```
ip route 0.0.0.0 0.0.0.0 10.201.254.1
ipv6 route ::/0 fd00:201:fe::
```

Do some ping and traceroute tests.

```
R11# ping 10.20.255.1
R11# ping 10.30.255.1
R11# traceroute 10.20.255.1
R11# traceroute 10.30.255.1
```

Q. Can you reach the routers in other networks?

Don't forget to save your configurations.

6 Appendix A - ISP1 Sample Configuration

```
hostname ISP1
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
ip subnet-zero
ip classless
no ip source-route
ipv6 unicast-routing
!
interface Loopback0
 ip address 10.201.255.1 255.255.255.255
 ipv6 address fd00:201:ff::1/128
!
interface GigabitEthernet1/0
 description Link to IXP
 ip address 10.251.1.1 255.255.255.0
 no ip directed-broadcast
 no ip redirects
 no ip proxy-arp
 ipv6 address fd00:251:1::1/64
 ipv6 nd ra supress
 no shutdown
!
interface GigabitEthernet3/0
 description P2P Link to R12
 ip address 10.201.254.1 255.255.255.252
 no ip directed-broadcast
 no ip redirects
 no ip proxy-arp
 ipv6 address fd00:201:fe::/127
 ipv6 nd ra supress
 no shutdown
!
```

```
ip route 10.10.0.0 255.255.0.0 10.201.254.2
ipv6 route fd00:10::/32 fd00:201:fe::1
!
ip route 0.0.0.0 0.0.0.0 10.251.1.2
ipv6 route ::/0 fd00:251:1::2
```