Static Routing Exercise

Campus Network Design Workshop

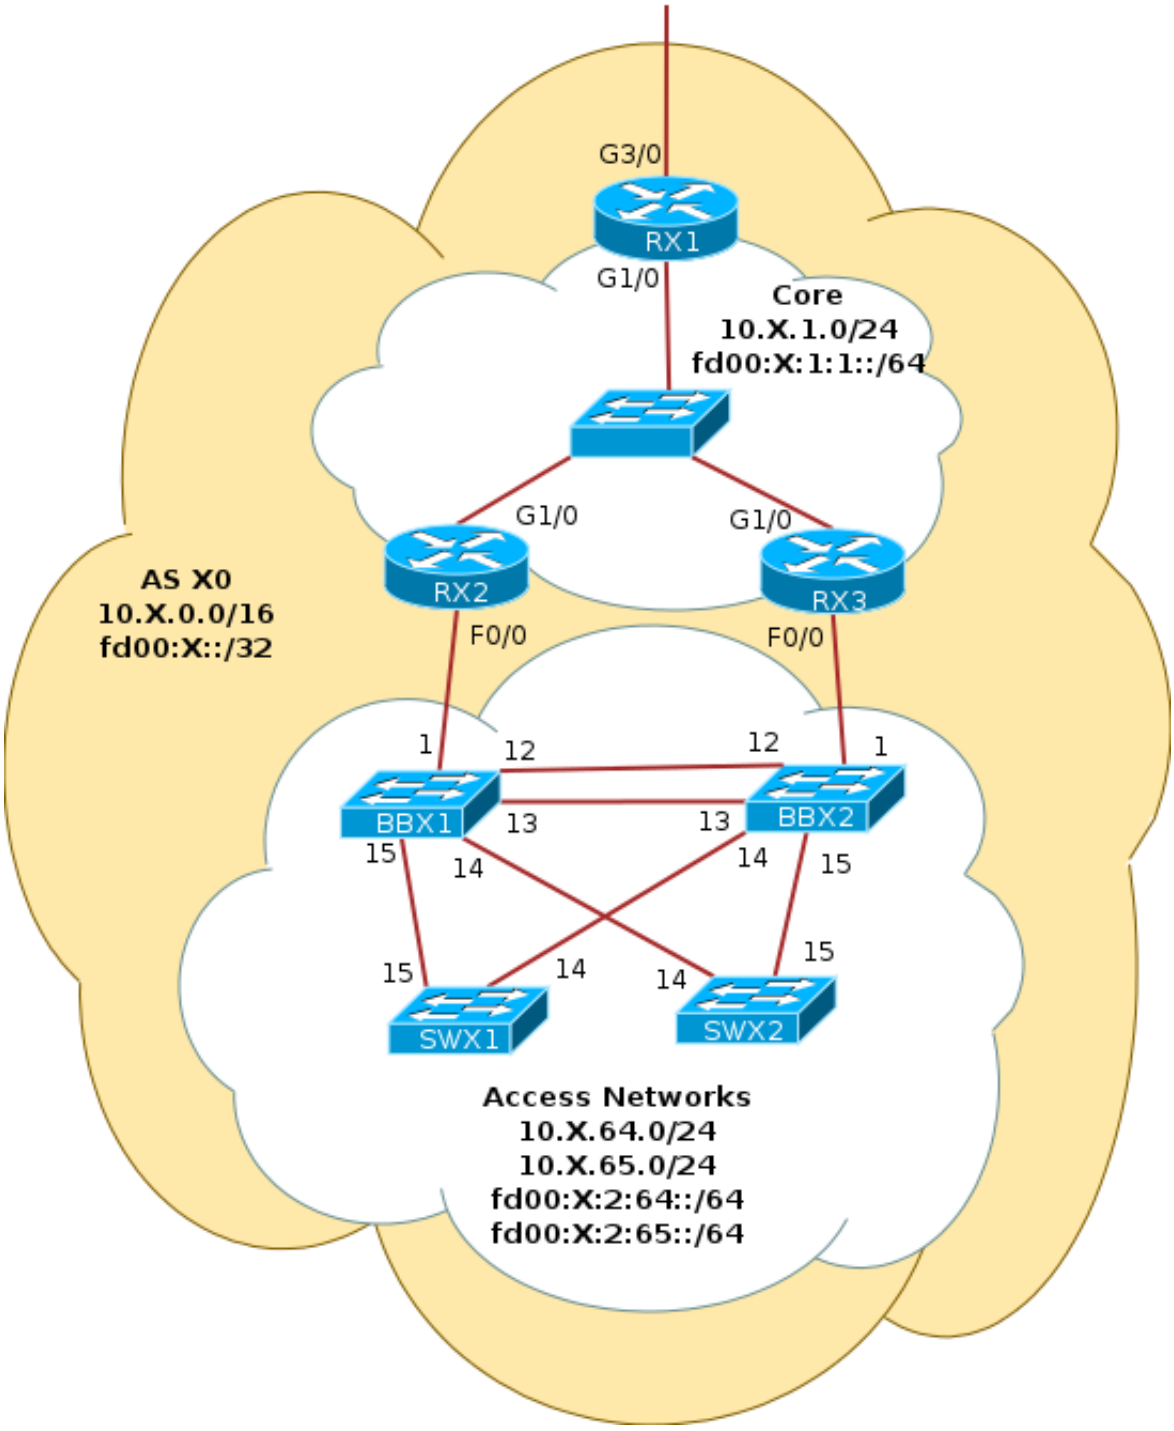

Physical Topology

Group interconnections:

```
    ISP Link network 10.254.255.0/24 and fd00:fe:0:0::/64
+---------------+---------------+---------------+---------------+----->
|               |               |               |               |
|.1             |.2             |.3             |.4             |.5
|g3/0           |g3/0           |g3/0           |g3/0           |g3/0
```

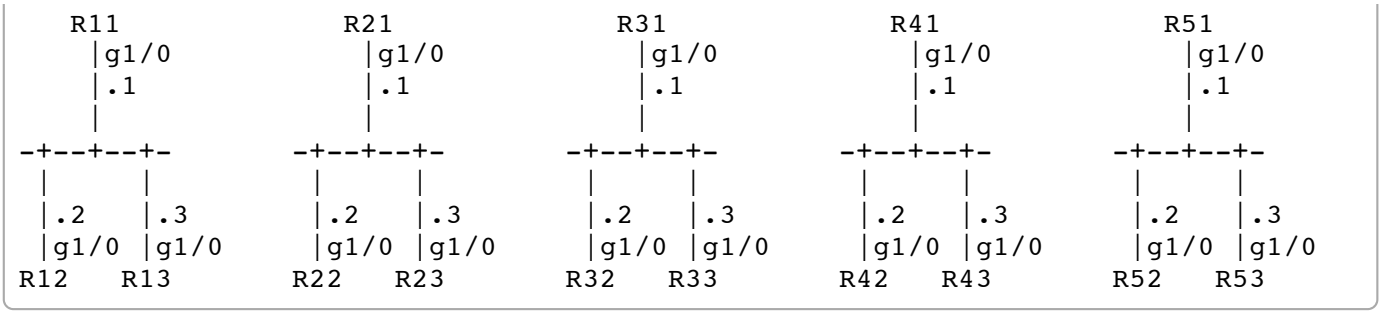

g1/0 is shorthand for the GigabitEthernet 1/0 interface on the router, etc.

Introduction

The purpose of this exercise is to learn how to configure static routes on a group of Cisco routers to allow full IP reachability between the various networks.

Router types used in the lab

Cisco 7206 VXR

Any Cisco IOS based router platforum is suitable for these exercises. We choose to use the Cisco 7206 VXR router.

The software used is from Cisco's IOS 15.1M or 15.2S release trains, but these exercises should be doable on any release from IOS 12.4 and 12.4T onwards. Obviously it is strongly recommended to use the newest release possible as it will have the most recent security fixes (not to mention new or more enhanced features when compared with the older releases).

Address Space Allocation

The allocations have been made assuming 5 groups in the workshop. If more groups are needed, simply extend the address plan following the same scheme.

Group	IPv4 Block	IPv6 Block
1	10.1.0.0/16	fd00:1::/32
2	10.2.0.0/16	fd00:2::/32
3	10.3.0.0/16	fd00:3::/32
4	10.4.0.0/16	fd00:4::/32
5	10.5.0.0/16	fd00:5::/32
ISP	10.254.0.0/16	fd00:fe::/32

Each group will then further partition their space like this:

IPv4	IPv6	Description
10.X.1.0/24	fd00:X:1:1::/64	Core Network
10.X.64.0/24	fd00:X:2:64::/64	Data Subnet (VLAN 64)
10.X.65.0/24	fd00:X:2:65::/64	VOIP Subnet (VLAN 65)

10.X.254.0/24	fd00:X:0:FE::/64	Router Loopback Subnet
10.X.255.0/24	fd00:X:0:FF::/64	Switch MGMT Subnet (VLAN 255)

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

The groups are connected together using a link subnet 10.254.255.0/24 and fd00:fe:0:0::/64, where each group's border router RX1 has address 10.254.255.X and fd00:fe:0:0::X on its external interface GigabitEthernet3/0.

Exercises

Basic Router Configuration

Configure each router based on the following example:

```
hostname Rxx      <-- e.g. 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
ipv6 cef
```

Explanations for some of the above commands

aaa new-model

The three AAA commands enable a more scalable method of authenticating user access to a router. AAA stands for authentication, authorization and accounting. The new-model mode gives greater flexibility, allowing provision of user accounts, as well as local and off-router authentication, authorization and accounting capabilities (amongst others).

service password-encryption

This specifies that all passwords stored on the router are obfuscated so they are not readable by anyone viewing the configuration. While better than plain text, the algorithm used is very simple, and easily reversible.

transport preferred none

This specifies that whatever is entered on the console and vty ports must be a command that the router understands. If 'none' was not specified, the router would try many of the supported transports (eg telnet, ssh,...) to resolve what the text sequence is.

no logging console

Given we are accessing the routers in the lab through their console ports, we don't really want the log messages cluttering up our work. Better to divert the console message into a log file, and view that log file as and when we need to. Besides, the console port on a router is a 9600 bits per second serial interface, and excessive console log messages can seriously impact the CPU performance.

Interface Configuration

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

- RX1: use Y=1
- RX2: use Y=2
- RX3: use Y=3

```
interface loopback 0
 ip address 10.X.254.Y 255.255.255.255
 ipv6 address fd00:X:0:fe::Y/128
!
interface GigabitEthernet1/0
 ip address 10.X.1.Y 255.255.255.0
 description Link to Core
 no ip redirects
 no ip directed-broadcast
 no ip proxy-arp
 ipv6 address fd00:X:1:1::Y/64
 no ipv6 redirects
 no shutdown
```

In addition, router RX1 will configure an interface to allow it to reach the other groups:

```
interface GigabitEthernet3/0
 ip address 10.254.255.X 255.255.255.0
 description Link to Groups
 no ip redirects
 no ip directed-broadcast
 no ip proxy-arp
 ipv6 address fd00:fe:0:0::X/64
 no ipv6 redirects
 ipv6 nd prefix default no-advertise
 ipv6 nd ra suppress all
 no shutdown
```

Routers RX2 and RX3 will configure access VLAN interfaces facing the backbone switches.

RX2:

```
interface Fast0/0
 no ip address
 no shutdown
!
interface Fast0/0.64
 encapsulation dot1Q 64
 ip address 10.X.64.2 255.255.255.0
 description Data Subnet (VLAN 64)
 no ip redirects
 no ip directed-broadcast
 no ip proxy-arp
```

```
 ipv6 address fd00:X:2:64::2/64
 no ipv6 redirects
 no shutdown
```

Do the same for VLANs 65 and 255.

RX3:

```
interface Fast0/0
 no ip address
 no shutdown
!
interface Fast0/0.64
 encapsulation dot1Q 64
 ip address 10.X.64.3 255.255.255.0
 description Data Subnet (VLAN 64)
 no ip redirects
 no ip directed-broadcast
 no ip proxy-arp
 ipv6 address fd00:X:2:64::3/64
 no ipv6 redirects
 no shutdown
```

Do the same for VLANs 65 and 255.

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).

IP directed broadcast is disabled by default in all Cisco IOS releases since 12.0, but most network operators still include it in their configuration templates.

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.

The proxy ARP feature is disabled by default in all Cisco IOS releases since 12.0, but most network operators still include it in their configuration templates.

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 illegitimate 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.

ICMP redirects are disabled by default in all Cisco IOS releases since 12.0, but most network operators still include it in their configuration templates.

ipv6 nd ra suppress [all]

Don't send Router Announcement messages on this interface. These are invitations to other devices on this network to use this device as their default gateway. You would turn this off where this router is not an appropriate choice of default gateway on that network.

The "all" flag is new in IOS 15.1(3)T3. It also disables solicited RA messages (where a client sends a broadcast asking for a router) as well as periodic messages.

There should never be any need for Router Advertisement on an infrastructure interface as the only devices there would be routers, servers, and other L3 devices which have their network information configured manually (rather than automatically).

ipv6 nd prefix default no-advertise

This prevents the router from sending any prefixes as part of router advertisements, so the client will not auto-configure itself with a global IPv6 address. This is helpful for IOS versions where you cannot suppress solicited RA messages.

Testing Connectivity

Ping your neighbor routers. For example from the first router in the group (RX1) you would do:

```
RX1# ping 10.X.1.2          <- RX2
RX1# ping 10.X.1.3          <- RX3
RX1# ping fd00:X:1:1::2     <- RX2
RX1# ping fd00:X:1:1::3     <- RX3
```

From RX2 you would ping .1 and .3, from RX3 you would ping .1 and .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 brief    : Summary of all IP interfaces
show ip interface <int>    : Shows detailed interface IP state and config
show ipv6 neighbors    : Shows IPv6 neighbors
show ipv6 interface brief  : Summary of all IPv6 interfaces
show ipv6 interface <int>  : Shows idetailed nterface IPv6 state and config
```

Now try pinging these other addresses in your network:

```
RX1# ping 10.X.254.2      <- RX2 loopback
RX1# ping 10.X.254.3      <- RX3 loopback
RX1# ping 10.X.64.2       <- RX2 data VLAN interface
RX1# ping 10.X.64.3       <- RX3 data VLAN interface
RX1# ping 10.X.65.2       <- RX2 VOIP VLAN interface
```

```
RX1# ping 10.X.65.3       <- RX3 VOIP VLAN interface
RX1# ping 10.X.255.2      <- RX2 management VLAN interface
RX1# ping 10.X.255.3      <- RX3 management VLAN interface
RX1# ping ipv6 fd00:X:0:FE::2   <-- ditto for IPv6
RX1# ping ipv6 fd00:X:0:FE::3
RX1# ping ipv6 fd00:X:2:64::2
RX1# ping ipv6 fd00:X:2:64::3
RX1# ping ipv6 fd00:X:2:65::2
RX1# ping ipv6 fd00:X:2:65::3
RX1# ping ipv6 fd00:X:0:FF::2
RX1# ping ipv6 fd00:X:0:FF::3
```

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

Now try pinging some addresses in the other groups (Z != X):

```
Rxx# ping 10.254.255.Z
Rxx# ping fd00:fe:0:0::Z
```

Does it work from your border router RX1? Does it work from the core routers RX2 and RX3? Can you explain what is happening and why?

Static routing

Look at the routing table (RIB):

```
show ip route
show ipv6 route
```

To view the forwarding table (FIB):

```
show ip cef
show ipv6 cef
```

Can you find route entries for the other groups, and for the ISP network, in the route table ? ... In the forwarding table ?

What do you need to do to be able to reach those groups ? What do those groups need to do to be able to reach your group ?

On your routers you will need to create static routes for:

- all of the other groups
- the ISP address space
- any networks in your group to which the router is not directly connected

What will those routes point to (next hop) on R11 ? What will those routes point to (next hop) on R12 and R13 ?

Remember the syntax for adding routes is:

```
ip route SUBNET MASK NEXT-HOP
ipv6 route SUBNET/PREFIXLEN NEXT-HOP
```

For example:

```
R11(config)# ip route 10.2.0.0 255.255.0.0 10.254.255.2
R11(config)# ipv6 route fd00:2::/32 fd00:fe:0:0::2
```

Based on the information above, create the required routes to be able to reach all the other groups, and the ISP address space.

Default route

Once you have all the class routing up, you can try adding a default route.

On the network where the RX1 routers are linked together, there is also an ISP router; it has IP address 10.254.255.254 and fd00:fe:0:0::fe/64

Where should the default route on RX2 and RX3 point?

When you have done this, can you now ping outside the class? (Note: IPv4 should work but there may not be external IPv6 connectivity available)

If ping to external hosts works, perhaps try a traceroute as well, and show your traceroute to the workshop instructors.