Layer 2 Exercise

Campus Network Design Workshop

Contents

1 Part 1

1.1 Introduction

The purpose of these exercises is to build Layer 2 (switched) networks utilizing the concepts explained in today's design presentations. Students will see how star topology, aggregation, virtual LANs, Spanning Tree Protocol, etc. are put to work.

There will be 6 groups of students, with 6 switches per group. The distribution of IP address space for the building (Layer 2) networks will be as follows:

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 (1,2,3,4,5,6)

3

1.1.1 Using private address space

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

1.1.2 Switch types used in the lab

Cisco 3725 with 16 Port 10BaseT/100BaseTX EtherSwitch (NM-16ESW) module

Note: This Cisco model is actually a router, but the 16-port module provides basic Layer-2 capabilities, and we will use these as switches. Dynamips does not support the emulation of the Cisco Catalyst class of switches, unfortunately.
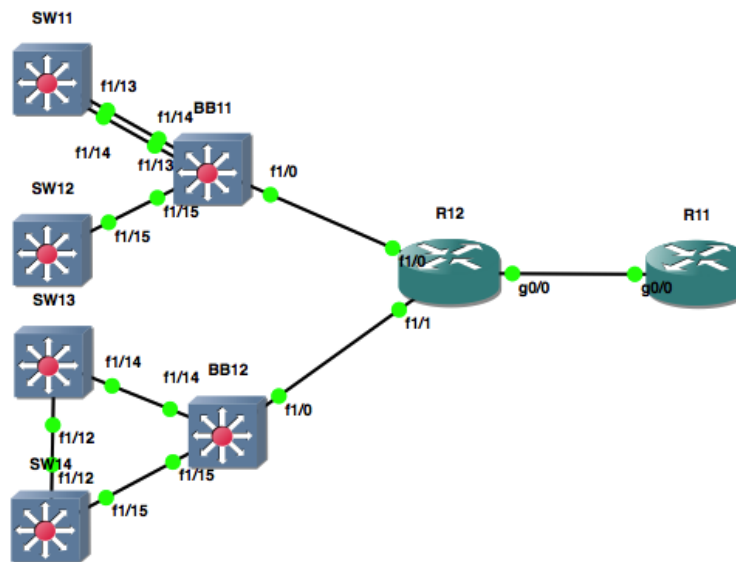
1.2 Hierarchical, redundant network

Figure 1: Lab topology

Our building network consists of two backbone switches each with two edge switches. The backbone switches connect to the core of our campus network and serve as aggregation points for all the edge switches. Edge switches serve the end users. Each edge switch has a connection to its backbone switch.

Switches BB11, SW11 and SW12 are used to represent one campus department, Engineering and Computer Science, and switches BB12, SW13 and SW14 represent the Library. We will also use the abbreviations ECS and LIB in the notes.

1.2.1 Lab access instructions

Refer to the file called dynamips-access-instructions.txt

1.3 Basic Switch Configuration

We've set up the switches in the lab with a very basic setup:

```
hostname BB11
```

```
!
aaa new-model
aaa authentication login default local
aaa authentication enable default enable
aaa session-id common
!
no logging console
logging buffered 8192 debugging
ipv6 unicast-routing
no ip domain-lookup

username nsrc secret nsrc
enable secret nsrc
service password-encryption

line con 0
  exec-timeout 0 0
  transport preferred none
  stopbits 1
line aux 0
  exec-timeout 0 0
  stopbits 1
line vty 0 4
  transport preferred none
  exec-timeout 0 0
```

You could have used these instructions to configure each switch:

1. Name the switch

   ```
   enable
   config terminal
   hostname <NAME>
   ```

2. Configure Authentication

   ```
   aaa new-model
   aaa authentication login default local
   aaa authentication enable default enable
   aaa session-id common
   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. Turn on IPv6

```
ipv6 unicast-routing
```

6. Exit configuration mode and save

```
end
write memory
```

1.3.1 IP Address Configuration

1. Assign each switch different IP addresses as follows:

```
int vlan 1
  ip address 10.1X0.33.Y 255.255.255.0
  ipv6 address fd00:1X0:2:33::Y/64
  no shut
end
```

Replace the "X" with your group number and replace "Y" like this:

```
 ECS switches:

 BBX1: 10.1X0.33.4 and fd00:1X0:2:33::4
 SWX1: 10.1X0.33.6 and fd00:1X0:2:33::5
 SWX2: 10.1X0.33.7 and fd00:1X0:2:33::6

 LIB switches:

 BBX2: 10.1X0.33.5 and fd00:1X0:2:33::7
 SWX3: 10.1X0.33.8 and fd00:1X0:2:33::8
 SWX4: 10.1X0.33.9 and fd00:1X0:2:33::9
```

Verify connectivity by pinging each switch. Do not continue until you can ping each switch from every other switch in the group.

HINT: If ping fails, but the configuration seems OK, try doing the following:

```
int vlan 1
 shutdown
 no shutdown
end
```

(this is not normal, but most likely a bug in the IOS code somewhere)

1.4   Spanning Tree Protocol

1.4.1   STP Status

Run the following commands on each switch and pay close attention to the output:

```
show spanning-tree brief
show spanning-tree blockedports
show spanning-tree
```

   a. What is the priority on each switch?

   b. Which switch is the root? Why?

   c. Which ports are blocked? Why?

1.4.2   STP Configuration

   1. Configure the STP priorities explicitly for each switch, according to the plan in Appendix A.

For example, on BB11:

```
BB11(config)#spanning-tree vlan 1 priority 12288
```

   2. Verify:

```
show spannning-tree brief
```

Why is it so important to set the priorities explicitly?

1.4.3 Disabling STP

We are now going to disable spanning tree to see what effect it has.

WARNING: Disabling spanning tree has a significant effect on the Dynamips server's CPU load. For this reason, we cannot have all groups disable spanning tree at the same time. We will take turns.

ASK THE INSTRUCTOR BEFORE DISABLING STP!!!

When you get the go-ahead from the instructor, execute the following on each of the LIB switches:

```
no spanning-tree vlan 1
```

Can the switches ping each other reliably now? Why?

Watch the port counters on the inter-switch links.

```
show interfaces stats
```

What happens with the counters of the connected interfaces? What is going on?

Very quickly enable STP again on all switches:

```
spanning-tree vlan 1
```

1.4.4 Simulate a backbone failure

1. Disconnect BBX1 from the rest of the network:

```
interface range fastEthernet 1/13 - 15
 shutdown
```

While it is cut off from the rest, verify spanning tree status on the other switches.

   a. Who is the root now?

   b. Verify port roles and status. Verify connectivity with ping.

2. Reconnect BBX1:

```
interface range fastEthernet 1/13 - 15
 no shutdown
```

What happens to the spanning tree when the switch comes back online?

2 Part 2

2.1 VLANs

We now want to segment the network to separate end-user traffic from VOIP and network management traffic. Each of these segments will be a separate subnet.

2.1.1 Configure the switches with separate DATA, VOIP and MGMT VLANs for

ECS and LIB.

VTP (VLAN Trunking Protocol) is a proprietary Cisco technology that allows for dynamic VLAN provisioning. We will not use it here.

1. Disable VTP by setting it to 'transparent mode':

```
vtp mode transparent
```

2. Add the VLANs to the VLAN database and give them names to better identify them:

```
vlan 64
 name "ECS DATA"
vlan 65
 name "ECS VOIP"
vlan 254
 name "ECS MGMT"

vlan 74
 name "LIB DATA"
vlan 75
 name "LIB VOIP"
vlan 255
 name "LIB MGMT"
```

3. Move the IP address to the MGMT vlan (notice the new subnet octets "254" and "255"):

On the ECS switches BB11, SW11 and SW12:

```
interface vlan 1
 no ip address
```

```
 no ipv6 address
interface vlan 254
 ip address 10.1X0.254.Y 255.255.255.0
 ipv6 address fd00:1X0:2:254::Y
ip default-gateway 10.1X0.254.2
```

On the LIB switches BB12, SW13 and SW14:

```
interface vlan 1
 no ip address
 no ipv6 address
interface vlan 255
 ip address 10.1X0.255.Y 255.255.255.0
 ipv6 address fd00:1X0:2:255::Y
ip default-gateway 10.1X0.255.2
```

Verify connectivity between switches. Can you ping? What's missing?

4. Configure trunk ports. Do the following for each port that needs to tag
   VLAN frames:

```
interface FastEthernet1/14
 switchport mode trunk
 switchport trunk encapsulation dot1q
```

Note: Check Figure 1 to see which ports you need to modify. BBX1 and BBX2
are each connected to the RX2 router on Fast1/0. This port also needs to be a
trunk.

Try pinging within the ECS and LIB switches again. It should work now.

5. Designate 5 edge ports for each ECS DATA and VOIP VLAN access:

On SWX1 and SWX2 only:

```
interface range Fast1/1 - 5
 switchport mode access
 switchport access vlan 64
!
interface range Fast1/6 - 10
 switchport mode access
 switchport access vlan 65
```

Designate 5 edge ports for each LIB DATA and VOIP VLAN access:

On SWX3 and SWX4 only:

```
interface range Fast1/1 - 5
 switchport mode access
 switchport access vlan 74
!
interface range Fast1/6 - 10
 switchport mode access
 switchport access vlan 75
```

Verify which ports are members or trunks of each vlan:

```
show vlan-switch id <VLAN ID>
```

Imagine that there are computers connected to the DATA vlan. Would they be able to ping the switch? Explain your response.

2.2   Verify the Spanning Tree status

```
show spanning-tree brief
```

Notice the root and bridge priorities on each VLAN (1,64,65,74,75,254,255). Are they the same?

This is called "Per-VLAN spanning tree", or PVST. This means that the switches are creating 7 separate trees, each with its own parameters, status, calculations, etc.

Use the commands in Appendix B to Configure PVST Priorities and check the root and bridge priorities again.

Note: Imagine if you had several hundred VLANs! This is certainly not ideal. There are better standards, like "Multiple Spanning Tree" (MST), that allow the administrator to create only the desired number of trees, and map groups of VLANs to each tree. Unfortunately, this Cisco device does not support MST.

2.3   STP Extended Features

2.3.1   PortFast

PortFast is a feature that allows end-user stations to be granted instant access to the L2 network. Instead of starting at the bottom of the Blocking-Listening-Learning-Forwarding hierarchy of states (30 seconds!), Portfast starts at the top.

The port starts in Forwarding state, and if a loop is detected, STP does all its calculations and blocks the necessary ports. This feature should only be applied to ports that connect end-user stations.

1. Configure end-user ports to be in PortFast mode:

```
interface range fast1/1 - 10
 spanning-tree portfast
```

2.3.2 BPDUGuard

With PortFast, end-user ports still participate in STP. That means that anything connected to those ports can send BPDUs and participate in (and affect the status of) the spanning tree calculations. For example, if the device connected to the edge port is configured with a lower bridge priority, it becomes the root switch and the tree topology becomes suboptimal.

Another useful Cisco feature that avoids this situation is BPDUGuard. At the reception of BPDUs, the BPDU guard operation disables the port that has PortFast configured.

1. Enable BPDUGuard on all ports with PortFast enabled:

```
spanning-tree portfast bpduguard
```

2.4 Port Bundling

We now want more capacity and link redundancy between a backbone and an edge switch.

1. Configure a Port Channel between BBX1 and SWX1:

On each switch:

```
interface port-channel 1
 description description BBX1-SWX1 aggregate link
!
interface range fast1/13 - 14
 channel-group 1 mode on
```

2. Verify the status:

```
show interface port-channel 1
```

What capacity do you have now on the new trunk? Hint: Look for the line that says BW ... Kbit/sec

3. Disable one of the ports in the bundle.

```
interface fast 1/13
 shutdown
```

Is the channel still up?

4. Enable it again:

```
interface fast 1/13
 no shutdown
```

Note: There is a standard protocol for port bundling. It's called "LACP" (Link Aggregation Control Protocol). This particular Cisco device does not support LACP, so these port channels are actually using a proprietary Cisco protocol called "EtherChannel". All modern switches support LACP, so we strongly recommend using it, instead of any proprietary versions.

3 Reference

3.1 Appendix A - Spanning Tree Configuration

Refer to this priority table below for the appropriate priorities on each switch.

Priority	Description	Notes
0	Core Node	The core switches/routers will not be participating in STP... reserved in case they ever are
4096	Redundant Core Node	Ditto
8192	Reserved	
12288	**Building Backbone**	
16384	Redundant Backbones	
20480	Secondary Backbone	This is for building complexes, where there are separate building (secondary) backbones that terminate at the complex backbone.
24576	**Access Switches**	This is the normal edge-device priority
28672	Access Switches	Used for access switches that are daisy-chained from another access switch. We're using this terminology instead of "aggregation switch" because it's hard to define when a switch stops being an access switch and becomes an aggregation switch.
32768	Default No	managed network devices should have this priority.

Table 3: Priority Table

3.2   Appendix B - Configuring PVST priorities

Configure BBX1 as the root switch for VLANs 64, 65 and 254 and BBX2 as the root switch for VLANs 74, 75 and 255.

On BBX1:

```
spanning-tree vlan 64 priority 12288
spanning-tree vlan 65 priority 12288
spanning-tree vlan 254 priority 12288
```

On BBX2:

```
spanning-tree vlan 74 priority 12288
spanning-tree vlan 75 priority 12288
spanning-tree vlan 255 priority 12288
```

On SWX1 and SWX2, the priorities are the same on every VLAN:

```
spanning-tree vlan 64 priority 24576
spanning-tree vlan 65 priority 24576
spanning-tree vlan 255 priority 24576
```

On SWX3 and SWX4, the priorities are the same on every VLAN:

```
spanning-tree vlan 74 priority 24576
spanning-tree vlan 75 priority 24576
spanning-tree vlan 255 priority 24576
```

2. Verify that the root switch is the correct one in all cases:

```
show spanning-tree brief
```