

Layer 2 Network Design Lab

October 23, 2014

Part 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 5 groups of students, with 4 switches per group. The distribution of IP address space for the building (Layer 2) networks will be as follows:

- Group 1: 10.1.64.0/24
- Group 2: 10.2.64.0/24
- Group 3: 10.3.64.0/24
- Group 4: 10.4.64.0/24
- Group 5: 10.5.64.0/24

Switch types used in the lab

Cisco 3750

Lab access instructions

Refer to the file called [lab-access-dynamips.htm](#)

Hierarchical, redundant network

Our building network consists of two redundant backbone switches and 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 both backbone switches, so that if one of the backbone switches fails, the switch has an alternative connection.

Basic Switch Configuration

Follow these instructions to configure each switch:

1. Name the switch

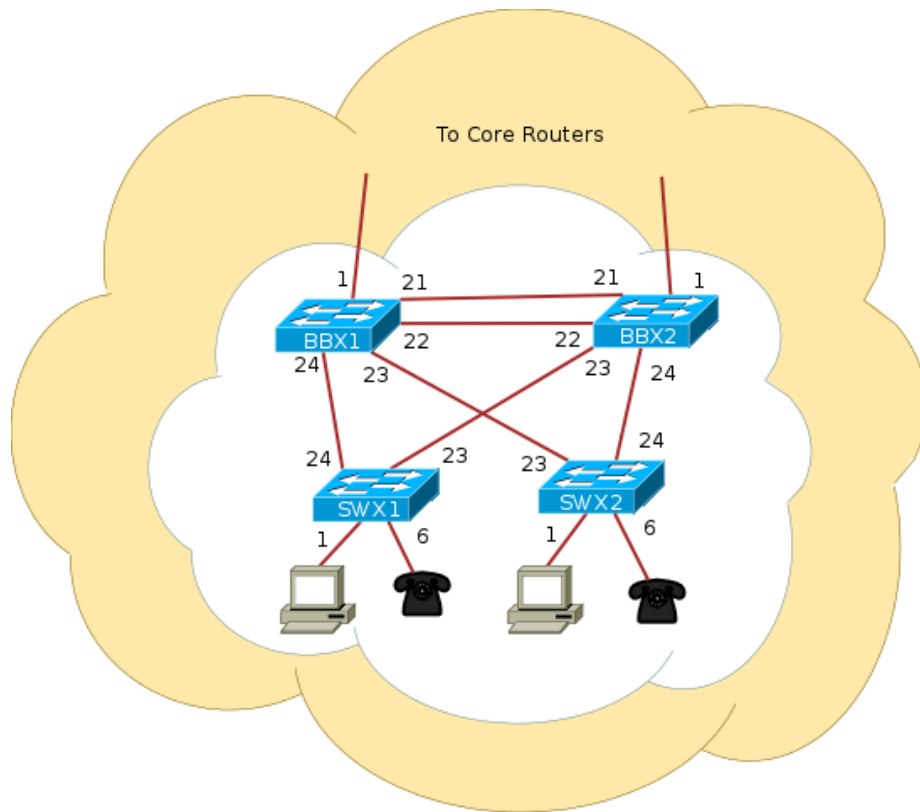


Figure 1: Lab topology

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

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. Exit configuration mode and save

```
end
write memory
```

IP Address Configuration

1. Assign each switch a different IP address as follows:

```
int vlan 1
ip address 10.X.64.Y 255.255.255.0
no shut
end
```

Replace the “X” with your group number, and replace “Y” like this:

1. BBX1: 10.X.64.4
1. BBX2: 10.X.64.5
1. SWX1: 10.X.64.6
1. SWX2: 10.X.64.7

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

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)

Spanning Tree Protocol

STP Status

Run the following commands and pay close attention to the output:

```
show spanning-tree bridge
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?

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 spanning-tree bridge
```

Why is it so important to set the priorities explicitly?

Disabling STP

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

ASK THE INSTRUCTOR BEFORE DISABLING STP!!!

THE INSTRUCTOR MAY DEMO A BROADCAST STORM INSTEAD

When you get the go-ahead from the instructor, execute the following on each switch:

```
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
```

If this still doesn't recover, it may be necessary to get the instructor to restart the simulator.

Simulate a backbone failure

1. Disconnect BBX1 from the rest of the network:

```
interface range fastEthernet 1/21 - 24
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 1/21 - 24
no shutdown
```

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

Part 2

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.

Configure the switches with DATA, VOIP and MGMT VLANs.

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 DATA
vlan 65
 name VOIP
vlan 255
 name MGMT
```

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

```
interface vlan 1
 no ip address
interface vlan 255
 ip address 10.X.255.Y 255.255.255.0
```

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 FastEthernet1/23
 switchport trunk encapsulation dot1q
 switchport mode trunk
```

Note 1: The Cisco default is to use dot1q encapsulation (rather than the Cisco proprietary ISL). But we include the dot1q command in the configuration in any case.

Note 2: Check Figure 1 to see which ports you need to modify. BBX1 and BBX2 are each connected to a router on Fast1/1. This port also needs to be a trunk.

Try pinging between switches again. It should work now.

5. Designate 5 edge ports each for 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
```

Verify which ports are members or trunks of each vlan:

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

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

Verify the Spanning Tree status:

```
show spanning-tree bridge
```

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

Note: This is called “Per-VLAN spanning tree”, or PVST. This means that the switches are creating 4 separate trees, each with its own parameters, status, calculations, etc. 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.

VLAN load-balancing with PVST

Your two aggregation switches are each connected to a core router (not shown in the pictures).

Suppose you wanted to load-balance the traffic from your various VLANs as they leave your aggregation switches towards your routers? How can you achieve this?

1. Configure BBX1 as the root switch for VLANs 64,65, and BBX2 as the root switch for VLAN 255. Also, make each switch a secondary root for the other VLAN(s):

On BBX1:

```
spanning-tree vlan 64 priority 12288
spanning-tree vlan 65 priority 12288
spanning-tree vlan 255 priority 16384
```

On BBX2:

```
spanning-tree vlan 64 priority 16384
spanning-tree vlan 65 priority 16384
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
```

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

```
show spanning-tree
```

STP Extended Features

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
```

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
```

Port Bundling

We now want more capacity and link redundancy between the aggregation switches.

1. Configure a Port Channel between BBX1 and BBX2:

```
interface port-channel 1
 description BBX1-BBX2 aggregate link
 !
interface range fast1/21 - 22
 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 1/21
shutdown
```

Is the channel still up?

4. Enable it again:

```
interface fast 1/21
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.

Reference

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 (BBX1) | |
| 16384 | Redundant Backbones (BBX2) | |
| 20480 | Secondary Backbone | This is for building complexes, where there are separate building (secondary) backbones that terminate at the complex backbone. |
| 24576 | Access Switches (SWX1, SWX2) | 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 1: Priority Table