brainbuzz.com™
Cramsession

brainbuzz.com

Cramsession™ for Cisco CCNP 2.0 BCMSN

This study guide will help you to prepare for Cisco exam 640-504, CCNP 2.0, Building Cisco Multilayer Switch Networks. Exam topics include Building a Campus Network, Managing Campus Traffic Network and Basic Router and Switch Configuration, Spanning Tree Protocol and VLAN.

UPDATE Check for the newest version of this Cramsession
http://cramsession.brainbuzz.com/checkversion.asp?V=2451957&FN=cisco/CCNP2_BCMSN.PDF

RATE! Rate this Cramsession
http://cramsession.brainbuzz.com/cramreviews/reviewCram.asp?cert=Cisco+CCNP+2%2E0+%2D+BCMSN

FORUMS Feedback Forum for this Cramsession/Exam
http://boards.brainbuzz.com/boards/vbt.asp?b=650

More Cramsession Resources:

Search for Related Jobs
http://jobs.brainbuzz.com/JobSearch.asp?R=&CSRE

IT Resources & Tech Library
http://itresources.brainbuzz.com

SkillDrill - skills assessment
http://skilldrill.brainbuzz.com

CramChallenge - practice questions
http://www.cramsession.com/signup/default.asp#day

Certification & IT Newsletters
http://www.cramsession.com/signup/

Discounts, Freebies & Product Info
http://www.cramsession.com/signup/prodinfo.asp

Notice: While every precaution has been taken in the preparation of this material, neither the author nor BrainBuzz.com assumes any liability in the event of loss or damage directly or indirectly caused by any inaccuracies or incompleteness of the material contained in this document. The information in this document is provided and distributed "as-is", without any expressed or implied warranty. Your use of the information in this document is solely at your own risk, and Brainbuzz.com cannot be held liable for any damages incurred through the use of this material. The use of product names in this work is for information purposes only, and does not constitute an endorsement by, or affiliation with BrainBuzz.com. Product names used in this work may be registered trademarks of their manufacturers. This document is protected under US and international copyright laws and is intended for individual, personal use only. For more details, visit our legal page.

intense school

Powerful IT Training & Certification

Training so INTENSE you'll end up positively certifiable!

intense prep

Powerful IT Training & Certification

When success is critical and failure is not an option you need the kind of Intense training that is only available at Intense School. Information Technology training and certification through Intense School and Intense Prep are not for the weak of heart. They're for IT professionals who have the will to succeed and the conviction to reach their career goals. Our training isn't easy. It's **INTENSE!**

Intense School offers a wide range of INTENSE Information Technology Boot Camps specializing in:

- ***Windows 2000® MCSE® certification in 14 days***
- ***CCNA™ & CCDA™ certification in 6 days***
- ***CCNP™ certification in 16 days***

Intense School for intense reasons:

- 95%+ pass rate in all Intense Boot Camps in 2000 – the highest in the industry
- Proprietary certification focused Windows 2000® curriculum
- Intense instructors are tops in the field – including well-known IT author/guru Kurt Hudson
- Expanded schedule and new locations for 2001

Classes are filling fast! To check availability go to:
www.IntenseSchool.com
Or, call now:
1-800-330-1446 or 954-370-7583

Intense Prep provides INTENSE Information Technology practice tests to prepare IT professionals for certification on:

- ***Microsoft® including Windows 2000®***
- ***Novell®***
- ***Cisco®***
- ***A+/Network+***

Jump start you career!

- Study in the comfort of your own home
- Intense Prep CD's are put together by industry leading professionals
- Certification focused Windows 2000® practice tests
- All materials are competitively priced

Download a FREE demo today:
www.IntensePrep.com/downloaddemo/
Or, for more information, call now:
1-877-996-3100 or 954-577-3100

Ft. Lauderdale, FL | Columbus, OH | San Diego, CA

Microsoft Certified Technical Education Center

This material is not sponsored, endorsed or affiliated with Cisco® Systems, Inc. Cisco®, Cisco Systems®, CCNA™, CCDA™, CCNP™, are the trademarks or registered trademarks of Cisco® Systems, Inc. All other trademarks are trademarks of their respective owners.

The evolution of NT School to Intense School/ Intense Prep enables us to offer a host of new products and services in the field of accelerated IT training and certification.

nt school® HOW YOU GET THERE

Contents:

Early LANs and Broadcasts

The first Local Area Networks (LANs) were fairly small and required few network devices to provide basic connectivity. They were built around broadcast domains, with all devices sharing the same network access in a flat network, and all packets visible to every device on the network. Early on it was seen that broadcast congestion was a major impediment to growth because this would limit the number of devices that could share the network.

Just think "broadcasts bad". Broadcasts take up bandwidth, and worse yet, every device that can hear it must interrupt its other work to process it. Adding more devices to a flat network results in more broadcasts, less bandwidth available to other purposes, and more interruptions for end-devices to process.

When the network became saturated, a bridge, or what we would today call a two-port store-and-forward switch, would be installed. Bridges flood broadcasts, multicasts, and unknown unicasts to all segments; all the bridged segments together form a single broadcast domain. The primary advantage to bridges is that they do keep user traffic isolated and allow more hosts to be added to the network by reducing the size of the collision domain.

Layer-2 switches are micro-segmentation devices (a review of the OSI model is available at the end of this document). In other words, you can think of them as bridges with dozens of ports. I've heard it said jokingly that the only reason they named them switches was because if they called them a bridge, they'd never sell any. Because switches facilitated the move away from shared media for end-devices, they have the affect of increasing available bandwidth without adding additional complexity. The packets do not need to be modified when data is passed between devices on the same VLAN (Virtual Local Area Network). This allows data to travel at wire-speed through the switched network fabric.

One of the problems in a bridged or switched environment is containing broadcasts, which, since these devices work at layer-2, are forwarded to all ports. VLANs are a technology created to address this issue. Each port on a switch can be configured to be part of a specific VLAN, which can be thought of as a subnet. These represent a broadcast domain defined by a set of ports; and some form of layer-3 device, such as a router, is required to move traffic between them.

The now out-of-date 80/20 rule refers to the goal of keeping 80% of the traffic on a network segment bound for a local destination (peer-to-peer and workgroup servers), and that no more than 20% of the network traffic should be directed across the backbone. Under this older design principle, workgroup servers would be deployed as the primary target of local workstations, allowing most of the traffic to be contained within the local subnet. This was done to conserve valuable bandwidth when media was shared and bandwidth was extremely limited.

The philosophy of network design has been reversed in the last few years, and this is reflected in the new 20/80 rule, which has the bulk of traffic directed at shared resources on the core layer of the Cisco Hierarchical Model. These design principles have 80% of the traffic routed off the local domain, usually to the core, and 20% (I would say, probably significantly less than that) is kept within the same broadcast domain. This has driven the improvement in layer-3 device performance.

Routing

Because bridges limit collision domains, but not broadcast domains, routing was introduced on the LAN to provide control, and to actually segment the network into separate entities limiting the effect of broadcast traffic. This was an important step in the evolution of the LAN, but it must be remembered that router ports are expensive, both in pure dollars and in processing overhead.

Routing, which occurs at layer-3, is much more complex than bridges and switches because packets must be ripped apart and reassembled as they pass through the router. This activity is CPU intensive. Routers do allow a great deal of control over data through the use of access lists, static routes and dynamic routing protocols.

The *diameter* of a network is the number of router hops from any one device to another. Cisco recommends having a consistent diameter. Their way of achieving this is through the use of the Hierarchical Design Model (if you are unfamiliar with the HDM, it is defined at the end of this document).

Command-Line Interface

Those who have worked on Cisco routers in the past will be comfortable with the Cat1900/2820 and 2900XL series access switches. The command nomenclature is familiar and, other then a few new commands, the same rules apply.

The Cat5000/6000/6500 series of switches use a different style of CLI, which is based on the Unix csh or c-shell prompt. This is commonly called the Set-based CLI, since this is one of the three commands used on these devices. They are:

- Set – Implements configuration changes
- Show – Verifies and provides information on the configuration
- Clear – Removes configuration elements

Virtual LANs

A VLAN is an extended logical network that is configured independent of the physical network layout. Each port on a switch can be defined to join whatever VLAN suits the Network Architect’s plans. Since each VLAN is a separate broadcast domain, routing must be enabled between them if data is to be passed.

VLAN Stuff

- Switches are used in VLANs to act as entry points for end-station devices into the switched fabric, and to provide flexibility in configuring group users, ports, or logical addresses and to make filtering and forwarding decisions.
- Most VLANs use frame filtering (frame tagging) to examine particular information about each frame based on user-defined offsets, and uniquely assign a user-defined ID to each frame header.
- Each hub segment connected to a switch port can be assigned to only one VLAN.
- VLAN ports on a switch can be assigned statically using a VLAN management application or by working directly within the switch. A more convenient approach, Dynamic VLANs, are ports on a switch that can automatically determine their VLAN assignments.

VLAN Commands:

Assign Ports to a VLAN on a Set-Based Switch

Switch (enable) set vlan *vlan-number module/port*

Example: Port 6, on module 4 needs to be assigned to VLAN 3. Keep in mind that you can assign several ports at once by using wildcards, such as "4/1-12" for the first twelve ports on module 4.

```
Switch (enable) set vlan 3 4/6
```

Assign Ports to a VLAN on an IOS-Based Switch:

```
Switch(config-if)# switchport mode access
```
Defines the VLAN membership mode for the port

```
Switch(config-if)# switchport access vlan 6
```
Assign the port to VLAN 6

Verify Port VLAN Status on an IOS-Based Switch

show interface *interface-id* switchport

Verify Port VLAN Status on a Set-Based Switch

```
Switch (enable) show vlan
```

Switch (enable) show port

Trunking

A point-to-point link configured on a single Fast-Ethernet, Gigabit Ethernet, or Fast- or Gigabit EtherChannel bundle and another network device, such as a router or second switch. Trunks transport the packets of multiple VLANs over a single network link.

The available trunking encapsulation types for Ethernet are:

- Inter-Switch Link (ISL) - a Cisco-proprietary trunking encapsulation that adds a 26-byte header and 4-byte trailer to the frame.
- IEEE 802.1Q (dot1q)- an industry-standard trunking encapsulation that does not change the size of the frame. Because multiple vendors support dot1q, it is becoming more common in newer switched networks.
- Negotiate - The port negotiates with its neighbor port to mirror its encapsulation configuration, either ISL (preferred) or 802.1Q trunk. This configuration option is only available in switch software release 4.2 and later.

There are five trunking modes:

- On - Forces the port to become a trunk port, even if the neighboring port does not agree to the change.
- Off – Forces the port to become non-trunking, even if the neighboring port does not agree to the change.
- Desirable - Causes the port to actively seek to convert the link to a trunk. The port becomes trunked if the neighboring port is set to either "on", "desirable", or "auto" modes.
- Auto - Makes the port available to serve as a trunk link. The port becomes a trunk port if the neighboring port is set to either "on" or "desirable" modes. This is the default mode for both Fast- and Gigabit Ethernet ports.
- Nonegotiate - Puts the port into permanent trunking mode but the neighboring port must be manually configured as a trunk port in order to establish a trunk.

Trunking Facts:

- For trunking to be auto-negotiated on Fast Ethernet and Gigabit Ethernet ports, the ports must be in the same VTP domain.

- Not all switches support all encapsulation methods; for instance the Cat2948G and Cat4000 series switches support only 802.1Q encapsulation. In order to determine whether a switch supports trunking, and what trunking encapsulations are supported, look to the hardware documentation or use the "show port capabilities" command.
- For trunking to be enabled on EtherChannel bundles the speed and duplex settings must be configured the same on all links. If part of an EtherChannel bundle fails, traffic will still be passed, but at a slower rate.

Trunking Commands

The command to enable trunking on a SET based switch is:

```
Switch (enable) set trunk module/port on|desirable|auto|off|nonegotiate
    [vlan-numbers] [isl|dot1q]
```

The command to disable trunking on a SET based switch is:

```
Switch (enable) clear trunk module/port vlan-numbers
```

The command to verify trunking status on a SET based switch is:

```
Switch (enable) show trunk [module/port]
```

The command to enable trunking on a 2900XL is:

```
Switch(config-if)# switchport mode trunk
Switch(config-if)# switchport trunk encapsulation isl|dot1q
```

The command to verify trunking on a 2900XL is:

```
Switch# show interface
```

The command to enable trunking on a Cat1900/2820 is:

```
Switch(config-if)# trunk [on|off|desirable|auto|nonegotiate]
```

The command to verify trunking on a Cat1900/2820 is:

```
Switch# show trunk
```

Spanning Tree Protocol (STP)

When multiple bridges or switches are installed, the possibility of loops forming and causing broadcast storms is a significant concern. A layer-2 loop occurs when a frame is transmitted from an end-device and detected by two different bridges or switches. These switches populate their address tables with the MAC address from the source address of the frame. Once their table is updated they forward the frame to the second segment, and then pick up the same MAC address from the other switch, feeding each other back the same information repeatedly. In other words, multi-active paths between stations create loops in the network, causing hosts to receive redundant messages and forcing switches to learn duplicate host MAC addresses from multiple ports.

The Spanning Tree Protocol (STP) was developed to prevent loops in the network and to route around failed elements. It is a link management protocol that provides path redundancy and prevents undesirable loops in the network. The Spanning Tree Algorithm (STA) calculates the best loop-free path throughout a switched network - switches send and receive spanning-tree frames at regular intervals, using them to construct a loop-free path, and forcing redundant data paths into a standby (blocked) state. All this is done in a way that is transparent operationally to the network hosts. The election of root bridges is performed through an exchange of data messages called Bridge Protocol Data Units (BPDUs). STP can be manually disabled on a per-VLAN or a global basis.

The following are characteristics of the STP broadcast domain:

- Where redundant links exist, any but the one with the least distance from the root switch are blocked.
- STP convergence can take upwards of 50 seconds.
- Broadcast traffic within the layer-2 domain (VLAN) interrupts every host.
- Broadcast storms within the layer-2 domain affect the whole domain.
- Isolating problems can be time consuming.
- Network security options within the layer-2 domain (VLAN) is limited.

STP can be configured two ways:

- Per-VLAN Spanning Tree (PVST) – A Cisco proprietary method of connecting through 802.1Q VLAN trunks, the switches maintain one instance of the spanning tree for each VLAN allowed on the trunk, versus non-Cisco 802.1Q switches which maintain one instance for ALL VLANs. This is the default STP used on ISL trunks. Since each VLAN has its own instance of STP, there is more granular control of the path selection process, and fewer sub-optimal paths may be invoked. Because the size of the STP topology is reduced, this can have the effect of reducing convergence time and increasing scalability and stability.
- Common Spanning Tree (CST) - When connecting a Cisco switch to a non-Cisco device through an 802.1Q trunk, the Cisco switch combines the spanning-tree instance of the 802.1Q native VLAN of the trunk with the spanning-tree instance of the non-Cisco 802.1Q switch. The primary advantages of CST are that only one set of BPDU's are used; it is only necessary to track changes for a single instance of STP, and non-Cisco

switches can be added to the mesh. However, with only one STP algorithm running, sub-optimal paths are more likely to be selected than under other methods. With CST, less bandwidth will be used to negotiate a root bridge, although with only one root bridge for the entire network, it may take longer for STP to recalculate when a change occurs.

Bridge Protocol Data Units (BPDUs) are multicast frames sent out periodically to announce the presences, resources and recent changes of a switch's configuration. They:

- Propagate bridge IDs in order for the selection of the root switch to take place.
- Are used to determine loop locations within a network.
- Provide notification of network topology changes.
- Remove loops by placing redundant switch ports in a backup state.

The Bridge ID defines which device will be the root bridge. It is made up of two parts; the 2-byte priority, a default value that can be changed by the Network Architect; and the 6-byte MAC address of the switch or bridge.

There are two factors involved in the root port selection:

- Path Cost, which is the sum of all links crossed to get to the root bridge
- Port ID

As BPDU leaves a port, it applies the root port cost. Path Cost is the total sum of all of the port costs, and is what STP uses to determine which ports should forward and which ports should block. If the path cost is the same for several ports, STP will use the lowest port ID.

STP Timers

- Hello timer - How often the switch broadcasts Hello messages to other switches.
- Forward delay timer - Amount of time a port will remain in the listening and learning states before going into the forwarding state.
- Maximum age timer – How long protocol information received on a port is stored by the switch.

STP Port States

Ports on an STP domain will progress through the following states:

- Blocking – Listens for BPDUs from other bridges, but does not forward them or any traffic.
- Listening – An interim state while moving from blocking to learning. Listens for frames and detects available paths to the root bridge, but will not collect host MAC addresses for its address table.
- Learning – Examines the data frames for source MAC addresses to populate its address table, but no user data is passed.
- Forwarding – Once the learning state is complete, the port will begin its normal function of gathering MAC addresses and passing user data.
- Disabled – Either there has been an equipment failure, a security issue, or the port has been disabled by the Network Administrator.

Notes about STP Port States:

- A port in blocking state does not participate in frame forwarding - switch always goes into blocking state immediately following switch initialization.
- When a port changes from the listening state to the learning state it is preparing to participate in frame forwarding.
- Port in the Forwarding state actually forwards frames (User data, BPDUs, etc.).

Root Switch Selection

Rather than allowing STP to define the root bridge, a good Network Architect will select the switch to be the root that minimizes unnecessarily convoluted data migrations. They will change the Spanning-tree bridge priority from the default value (32768) to a significantly lower value, ensuring that the switch becomes the root for the specified VLANs.

Spanning Tree Commands:

Enable spanning tree on per-VLAN or global basis:

```
Switch (enable) set spantree enable [vlans]
Switch (enable) set spantree enable all
```

Verify that spanning tree is enabled:

```
Switch (enable) show spantree [vlan]
```

Configure a switch as the root switch:

Switch (enable) set spantree root vlans [dia network_diameter] [hello hello_time]

Change the global port priority for a port:

Switch (enable) set spantree portpri *mod_num/port_num* priority

Change the port-VLAN priority for a VLAN on a switch port

Switch (enable) set spantree portvlanpri *mod_num/port_num* priority [vlans]

Change the global port cost for a switch port

Switch (enable) set spantree portcost *mod_num/port_num* cost

Change the port-VLAN cost for a VLAN on a switch port

Switch (enable) set spantree portvlancost *mod_num/port_num* cost cost [vlans]

Set the bridge priority for a VLAN

Switch (enable) set spantree priority *bridge_priority* [vlan]

Set the Hello time for a VLAN

Switch (enable) set spantree hello *interval* [vlan]

Set the forward delay time for a VLAN

Switch (enable) set spantree fwddelay *delay* [vlan]

Set the maximum aging time for a VLAN

Switch (enable) set spantree maxage *agingtime* [vlan]

PortFast

By default, all ports on a switch are assumed to have the potential to have bridges or switches attached to them. Since each of these ports must be included in the STP calculations, they must go through the four different states whenever the STP algorithm runs (when a change occurs to the network).

Enabling PortFast on the user access ports is basically a commitment between the Network Architect and the switch agreeing that the specific port does not have a switch or bridge connected, and therefore this port can be placed directly into the Forwarding state; this allows the port to avoid being unavailable for 50 seconds while it cycles through the different bridge states, simplifies the STP recalculation and reduces the time to convergence.

The command to enable PortFast on a SET based switch is:

```
Switch (enable) set spantree portfast module-number/port-number enable
Switch (enable) set spantree portfast 4/1-12 enable
```

The command to disable PortFast on a SET based switch is:

```
Switch (enable) set spantree portfast module-number/port-number disable
Switch (enable) set spantree portfast 4/1-12 disable
```

The command to verify PortFast status on a SET based switch is:

```
Switch (enable) show spantree
```

The command to enable PortFast on a 2900XL is:

```
Switch(config-if)# spanning-tree portfast
```

The command to disable PortFast on a 2900XL is:

```
Switch(config-if)# no spanning-tree portfast
```

The command to verify PortFast on a 2900XL is:

```
Switch# show spanning-tree
```

The command to enable PortFast on a Cat1900/2820 is:

```
Switch(config-if)# spantree start-forwarding
```

The command to disable PortFast on a Cat1900/2820 is:

```
Switch(config-if)# no spantree start-forwarding
```

The command to verify PortFast on a Cat1900/2820 is:

```
Switch# show spantree
```

UplinkFast

Convergence time on STP is 50 seconds. Part of this is the need to determine alternative paths when a link between switches is broken. This is unacceptable on networks where realtime or bandwidth-intensive applications are deployed (basically any network).

If the UplinkFast feature is enabled (it is not the default) AND there is a least one alternative path whose port is in a blocking state AND the failure occurs on the root port of the actual switch, not an indirect link; then UplinkFast will allow switchover to the alternative link without recalculating STP, usually within 2 to 4 seconds. This allows STP to skip the listening and learning states before unblocking the alternative port.

The command to enable UplinkFast on a SET based switch is:

```
Switch (enable) set spantree uplinkfast enable
```

The command to disable UplinkFast on a SET based switch is:

```
Switch (enable) set spantree uplinkfast disable
```

The command to verify UplinkFast status on a SET based switch is:

```
Switch (enable) show spantree
```

The command to enable UplinkFast on a 2900XL is:

```
Switch(config)# spanning-tree uplinkfast
```

The command to disable UplinkFast on a 2900XL is:

```
Switch(config)# no spanning-tree uplinkfast
```

The command to verify UplinkFast on a 2900XL is:

```
Switch# show spanning-tree
```

The command to enable UplinkFast on a Cat1900/2820 is:

```
Switch(config)# uplink-fast
```

The command to disable UplinkFast on a Cat1900/2820 is:

```
Switch(config)# no uplink-fast
```

The command to verify UplinkFast on a Cat1900/2820 is:

```
Switch# show uplink-fast
```

BackboneFast

BackboneFast is used at the Distribution and Core layers, where multiple switches connect together, and is only useful where multiple paths to the root bridge are available.

This is a Cisco proprietary feature that speeds recovery when there is a failure with an active link in the STP. Usually when an indirect link fails, the switch must wait until the maximum aging time (max-age) has expired before looking for an alternative link. This delays convergence in the event of a failure by 20 seconds (the max-age value). When BackboneFast is enabled on all switches, and an inferior BPDU arrives at the root port - indicating an indirect link failure - the switch rolls over to a blocked port that has been previously calculated.

BackboneFast Stuff:

- The Primary difference between UplinkFast and BackboneFast is that BackboneFast can detect indirect link failures and is used at the Distribution and Core layers; while UplinkFast is aware of only directly connected links, and is used primarily on Access layer switches. If UplinkFast is turned on for the root switch, it will automatically disable it.
- There is no BackboneFast command for IOS based switches. Since this is an enhancement for Core and Distribution layer devices only, and these are all Set-based switches.

The command to enable BackboneFast on a SET based switch is:

```
Switch (enable) set spantree backbonefast enable
```

The command to disable BackboneFast on a SET based switch is:

```
Switch (enable) set spantree backbonefast disable
```

The command to verify BackboneFast status on a SET based switch is:

```
Switch (enable) show spantree backbonefast
Switch (enable) show spantree summary
```

VLAN Trunk Protocol (VTP)

In a switched environment a subnet corresponds to a VLAN, and a VLAN may map to a single Layer 2 switch, or it may span several switches, especially at the access layer. Also, it is likely that one or more VLANs may be present on any particular switch.

VLAN Trunk Protocol (VTP) is a layer-2 messaging protocol that centralizes the management of VLAN additions, deletions and changes on a network-wide basis. This simplifies the management of large switched networks with many VLANs.

A VTP domain is specified by the Network Engineer and consists of one or more interconnected switches that share the same VLAN configuration. A switch can only be configured as a member of a single VTP domain. Changes to the global VLAN configuration for the domain can be implemented using either the CLI or an SNMP session.

Switches defined as part of VTP domains can be configured to operate in any of three VTP modes:

- *Server* – Advertise VLAN configuration to other switches in the same VTP domain and synchronize with other switches in the domain. Can create, modify, and delete VLANs as well as modify VLAN configuration parameters such as VTP version and VTP pruning for the entire domain. This is the default mode for a switch.
- *Client* - Advertise VLAN configuration to other switches in the same VTP domain and synchronize their VLAN configuration with other switches based on advertisements received over trunk links; however, they are unable to create, change, or delete VLAN configurations.
- *Transparent* - Does not advertise its VLAN configuration and does not synchronize its VLAN configuration with other switches. In VTP version 2, transparent switches do forward VTP advertisements.

Advertisement types include: requests from clients, summary advertisements and subset advertisements. An advertisement contains the VLAN IDs, the Emulated LAN names for ATM LANE, the 802.10 SAID values for FDDI, the VTP domain name, the VTP configuration revision number, the MTU size and the Frame format.

Early design specifications touted the ability of VTP to create global VLAN groups that the Network Engineer could use to have VLANs that would span vast networks;

however, in recent years it has become obvious that this has generated unnecessary and expensive wide area traffic for not much gain. Most design specifications now suggest creating VTP domains for each facility, and limiting the VTP advertisements sent over limited and expensive wide area links.

VTP advertisements carry configuration revision numbers that are incremented every time a VLAN is modified. This is used to identify the most recent changes to the network topology. When a switch finds an advertisement with a higher configuration revision number, it will save the new VTP database over the old one. A VLAN that does not exist in the new database is automatically deleted from the switch, and any ports that were in the VLAN will be orphaned.

It is a common problem for a newly ordained Network Engineer (also called a "newby" or "loser-boy") to add a switch that has been used on a separate or test network to a production network, not being aware of the revision number. Since test networks change much more often than production networks, the new switch likely has a higher configuration revision number than the production VTP domain. The result is that the entire domain's VTP database gets overwritten and any ports assigned to the lost VLANs lose their VLAN membership and become unavailable to users. If you receive a call that all the switched ports on a network have suddenly locked up and no traffic is being passed, one of the first places to look is the red-faced newest member of the network team who thought he was helping when he put a new switch in the network. Good documentation and control over physical access to network devices are probably your best defense against this problem. Also, to prevent this problem, the command "clear config all" should be used on any switch before it is added to a production network.

VTP pruning is a technique to limit VTP broadcast from branches of the network that do not contain member ports of a specific VLAN. By default, VTP pruning is disabled. VTP pruning must be enabled on a VTP server and promulgates pruning eligibility through the entire management domain. By default, VLAN 1 is always pruning-ineligible, and VLANs 2 through 1000 are pruning-eligible.

Configure a VTP Domain

Enter the VTP configuration mode

```
Switch (enable) vlan database
```

Set the VTP domain name to "Primary"

```
Switch (vlan) vtp domain Primary
```

Set the VTP domain password to "scubbie"

```
Switch(vlan) vtp password scubbie
```

VTP version 2 is enabled (to return to ver.1 - “**no vtp v2-mode”**)

Switch(vlan) **vtp v2-mode**

Set the switch to VTP server mode. The client or transparent arguments could be used instead.

Switch(vlan) **vtp server**

Verify VTP Operation

Display the VTP switch configuration and statistics

Switch **show vtp status**

Display the VTP counters for messages sent and received

Switch **show vtp counters**

Adding VLANs to a VTP domain

Enter the VTP configuration mode

vlan database

vlan vlan-id **name** vlan-name

Example: Add VLAN 6 to the domain and name it “accounting”. If a name is not specified, it defaults to the VLAN number designation, as in this case would be “VLAN0006”:

vlan 6 name accounting

Verify VLAN

Display the VLAN configuration

show vlan name *vlan-name*

show vlan name accounting

Displays a list of configured VLANs

show vlan brief

Deleting VLANs from a VTP domain

Enter the VTP configuration mode

vlan database

no vlan *vlan-id*

Example: Remove vlan 6 from the VTP domain and orphan any ports assigned to that VLAN:

no vlan 6

VLAN commands - Brief

- Vlan database - enter into VLAN configuration mode
- Vtp domain domain-name - configure a VTP administrative-domain's name
- Vtp password password-value - set the password for the VTP domain
- Vtp server - configure the switch as a server
- Vtp client - put the switch in VTP client mode
- Vtp transparent - put the switch in VTP transparent mode
- Show vtp status - show VTP configuration
- No vtp v2-mode - disable VTP version 2

VTP Stuff

- Global Information in a VTP Advertisement includes VTP Domain Name, VTP Configuration Revision Number, Update Identity, Update Timestamp, MD5 Digest.
- VLAN Information in a VTP Advertisement includes VLAN ID, VLAN Name, VLAN Type, VLAN State.
- VTP Version 2 has features not supported in VTP version 1, including Token Ring LAN Switching and VLANs, unrecognized Type Length Value, Version Dependent Transparent Mode and Consistency Checks. Please note that all the switches in the VTP domain must run the same VTP version.
- In general, don't enable VTP version 2 in the VTP domain unless all the switches are running version 2 as well. However, if the network is Token Ring, you must enable VTP version 2.
- VTP Pruning increases bandwidth by controlling traffic flow to the vital trunk links and to block flooded traffic to VLANs in the pruning eligible list. Enabling VTP pruning on a VTP server will enable it on the entire management domain.

Configuration Guidelines

- Max 250 active VLANs supported by a switch. Watch out though, as some switch models only support 6 VLANs.
- When creating a VLAN, the switch must be in VTP server or transparent mode.
- Default VLAN Configurations - Ethernet Parameters have an ID Range 1-1005. No limit on VLAN Name, and the MTU Size is 1500.

LAN Emulation

Am emulated LAN is a group of ATM-attached devices treated as an independent broadcast domain. Think of it as a single Ethernet segment or independent Token Ring. ELANs are made up of two components: the LAN Emulation Client (LEC) and LAN Emulation (LANE) services. The LEC can be located in the same device(s) as the LANE Services. LANE services are made up of a LAN Emulation Configuration Server (LECS), the LAN Emulation Server (LES), and the Broadcast and Unknown Server (BUS), and all of them can be located in the same device or distributed among one, two, or three devices.

To join an emulated LAN, LEC needs to contact the LECS in order to obtain its ATM address via reconfigured address for the LECS, ILMI or the well-known address of the configuration service. When two hosts are in the same emulated LAN, switches are enough for data transmissions. When two systems reside in different emulated LANs, a layer-3 router or switch must be used to interconnect them, regardless of the physical connection.

LAN Switching

All nodes on an Ethernet network can transmit at the same time, so the more nodes you have, the greater the possibility of collisions happening, which can slow the network down.

LAN segmentation means to break up collision domains by decreasing the number of workstations per segment using bridges or switches. Switches are sometimes called micro-segmentation devices, because there may be as little as one host per collision domain.

Switching is a layer-2 data manipulation that forwards through the network by destination MAC addresses.

These are the common Cisco switching techniques:

- Store-and-forward – receives the complete frame before forwarding. Copies the entire frame into the buffer and then checks for CRC errors. Higher latency then other techniques. This technique is used on Cat5000s.
- Cut-through – checks the destination address as soon as the header is received, and immediately forwards it out, lowering the latency level.
- Fast switching - The default switching type. It can be configured manually through use of the "ip route-cache" command. The first packet is copied into packet memory, while the destination network or host information is stored in the fast-switching cache.
- Process Switching - This technique doesn't use route caching, so it runs slow; however, slow usually means SAFE. To enable, use the command "no protocol route-cache".
- Optimum Switching – From its name you can understand what it is – high performance! This is the default on 7500's.

Multi-layer switching

Multi-layer Switching is the ability to use a combination of layer-2 switching technology, with layer-3 routing and layer-4 application specificity.

Layer-2 Switching

Layer-2 switching is hardware-based, using Application-Specific Integrated Circuits (ASICs) to bridge a network. The performance difference between a Layer-2 switch and a shared hub is significant. A layer-2 switch can be thought of as a bridge on steroids. It has all the same characteristics and limitations as bridging.

Problems with layer-2 switched networks:

- They provide scaling and performance issues on large bridged networks.
- The broadcast radiation increases with the number of hosts; broadcasts are seen by all end stations.
- STP can have slow convergence on large networks.

Layer-3 Switching

Traditional routers use CPUs that are general purpose devices, while a layer-3 switch uses an ASIC, a piece of high-speed hardware designed to perform a more limited set of tasks, in this case to achieve efficient routing (in some cases and under certain circumstances, wire-speed). For most purposes you can consider a layer-3 switch a device that integrates layer-2 and layer-3 (and sometimes layer-4) functionality in a single piece of equipment.

Depending on the network design, including what protocols, interfaces, and features are required, layer-3 switches can be used in place of routers and allow almost wire-speed routing. Standard routing protocol can be used for route determination, including OSPF, EIGRP, RIP, and IS-IS.

A router is used to determine conversations between end-devices, and then switching techniques continue the conversations. It has the following advantages: Hardware-based packet forwarding, high-performance packet switching, scalability, low latency, lower per-port cost, flow accounting, security and control over Quality of Service (QoS).

Layer-4 Switching

Layer-4 switching refers to hardware-based routing, using ASICs, which takes application specificity into consideration.

TCP or UDP flows include port number in the packet heading, which serves to identify the application under consideration.

Cisco routers have the ability to control traffic based on Layer-4 information using extended access lists and NetFlow switching.

Switching Hardware

To support multi-layer switching, you will need to have the following:

- Multilayer Switching-Switching Engine (MLS-SE) - Catalyst 2926G series switch, or Catalyst 5000 series switch with the NFFC (NetFlow Feature Card) or NFFC II. The NFFC is a daughter-card upgrade to the Supervisor Engine that is an ASIC-based layer-3 switching engine.
- Multilayer Switching-Route Processor (MLS-RP) - A Route Switch Module (RSM) or an externally connected Cisco router with software that supports MLS. The RSM is an IOS-based router on a blade that uses the same Reduced Instruction Set Computing (RISC) processor as the RSP2 engine in 7500 series routers. When MLS is enabled, the RSM or externally attached router continues to handle all non-IP protocols while offloading the switching of IP packets to the MLS-SE.
- Multilayer Switching Protocol (MLSP) – A protocol running between the MLS-SE and MLS-RP.

You’ll hear the configuration with an external router referred to as a “router-on-a-stick” or a “one-armed-router”.

Implementation Issues

- When using an external router, the ideal set up is one directly attached external router per switch to ensure proper caching.
- You can use Cisco high-end routers for MLS when they are externally attached to the switch, make the attachment with multiple Ethernet connections on an one per subnet basis or by using Fast or Gigabit Ethernet with Inter-Switch Link encapsulation.
- Router interfaces with input access lists or reflexive access lists cannot participate in MLS. However, you can translate input access lists to output access lists to provide the same effect.
- When an output access list is applied, the MLS cache entries for that interface are purged. However, entries associated with other interfaces are not affected at all.
- Flow mask mode is destination-ip when there is no access list on any MLS-RP interface. When there is a standard access list, the mode is source-destination-ip. When there is an extended access list, the mode is ip-flow.

Setting up Multi-layer Switching

These are the commands necessary to configure an internal or external Multi-layer Switch Route Processor:

Enabling MLSP

```
Switch(enable) mls rp ip
```

Entering into the router interface

```
Switch(enable) interface
```

Assign VLAN ID to the route processor interface

Switch(enable-if)# mls rp *vlan-id*

Place the external route processor in the interface of the VTP domain switch

Switch(enable-if)# mls rp *vtp-domain*

Enable the RSM interface

```
Switch(enable-if)# mls rp management-interface
```

MLS Flows

- Based only on layer-3 addresses.

- NFFC (or NFFC II) maintains layer-3 switching table (MLS cache) for the layer-3-switched flows.
- Whenever the Layer-3 switching entry for a flow ages out, the flow statistics will be exported to a flow collector application.
- Maximum MLS cache size is 128K.
- Cache larger than 32K increases the likelihood that a flow will get forwarded to the router.
- When a layer-3 packet is switched from source to destination, the switch performs a packet rewrite based on information learned from the router and stored in the MLS cache.
- If Host A and B are on different virtual LANs, when Host A sends a packet to the MLS-RP to be routed to Host B, the MLS-SE recognizes that the packet was sent to the MAC address of the MLS-RP, and will check the MLS cache to find the matching entry.
- MLS-SE uses flow mask modes to determine how MLS entries are created; the flow mask mode is based on the access lists configured on the MLS router interfaces.

Multicasting

- Unicast – A frame that will only be processed by the destination host (one machine to one machine)
- Broadcast – A frame that every host on the broadcast domain must process (one machine to all machines)
- Multicast – A frame that will only be processed by multicast members on the broadcast domain (one machine to a select list of machines)

Using Multicasts, an application can send a single stream of packets to a defined group of computers, instead of sending it one by one to each recipient, or flooding the network with broadcasts. Class-D addresses are reserved for multicast traffic and are allocated dynamically.

To manage multicast by allowing directed switching of multicast traffic, and also to dynamically configure switch ports so that IP multicast traffic is forwarded only to the appropriate ports, Cisco switches use:

- Internet Group Management Protocol (IGMP) - Standard protocol to manage the multicast transmissions passed to routed ports. One of the problems with this protocol is if a VLAN on a switch is set to receive, all the workstations on that VLAN will get the multicast stream.

- Cisco Group Management Protocol (CGMP) - Cisco proprietary protocols to control the flow of multicast streams to individual VLAN port members. Solves the problem sited above. Requires IGMP to be running on the router.

CGMP and IGMP software components run on both the Cisco routers and switches. Remember that CGMP is Cisco proprietary. When the CGMP/IGMP-capable router receives an IGMP control packet, it creates a CGMP or IGMP packet that contains the request type, the multicast group address, and the MAC address of the host. These request types can either be "join" or "leave" messages. The router sends the packet to a well-known address to which all switches listen, so that the supervisor engine module interprets the packet and modifies the forwarding table automatically. If a spanning-tree VLAN topology changes, the CGMP/IGMP-learned multicast groups on the VLAN are purged and the CGMP/IGMP-capable router generates new multicast group information. If a CGMP/IGMP-learned port link is disabled, the corresponding port is removed from any multicast group.

CGMP/IGMP-capable routers send periodic multicast group queries, so if a host wants to remain in a multicast group, it must respond to the query. If, after a number of queries, the router receives no reports from any host in a multicast group, the router sends a CGMP/IGMP command to the switch to remove the group from the forwarding tables. CGMP fast-leave-processing allows the switch to detect IGMP version-2 leave messages sent to the all-routers multicast address by hosts on any of the supervisor engine module ports.

Multicasting Commands

Display information on dynamically learned and manually configured multicast router ports

```
show multicast router mod_num/port_num vlan_id
```

Display total number of multicast address groups in each VLAN

```
show multicast group count vlan_id
```

CGMP Commands

Enable CGMP on the switch

```
Switch (enable) set cgmp enable
```

Verify that CGMP is enabled

```
Switch (enable) show cgmp statistics vlan_num
```

Enable CGMP fast-leave processing on the switch

```
          Switch (enable) set cgmp leave enable
```

Verify that CGMP fast-leave processing is enabled

```
          Switch (enable) show cgmp leave
```

Display information on those multicast router ports learned dynamically using CGMP

```
          Switch (enable) show multicast router cgmp mod_num/port_num vlan_id
```

Display information about multicast groups learned dynamically through CGMP

```
          Switch (enable) show multicast group cgmp mac_addr vlan_id
```

Display total number of multicast address groups in each VLAN that were learned dynamically through CGMP

```
          Switch (enable) show multicast group count cgmp vlan_id
```

Display CGMP statistics

```
          Switch (enable) show cgmp statistics vlan_id
```

Disable CGMP fast-leave processing on the switch

```
          Switch (enable) set cgmp leave disable
```

Disable CGMP on switch

```
          Switch (enable) set cgmp disable
```

IGMP Commands

Enable IGMP snooping on the switch

```
          set igmp enable
```

Verify that IGMP snooping is enabled

```
          show igmp statistics vlan_num
```

Enable IGMP fast-leave processing on the switch

```
          set igmp fastleave enable
```

Verify that IGMP fast-leave processing is enabled

```
          show igmp leave
```

Display information only on those multicast router ports learned dynamically using IGMP snooping

```
          show multicast router igmp mod_num/port_num vlan_id
```

Disable IGMP snooping on the switch

```
          set igmp disable
```

Protocol Independent Multicast (PIM)

PIM is used to forward multicast packets through a network. It must be enabled for a Cisco interface to perform IP multicast routing. Enabling PIM on an Interface also enables IGMP operation on that interface.

Interface can be configured to be in dense mode, sparse mode, or sparse-dense mode - the modes determine how the router populates its multicast routing table and how the router forwards multicast packets it receives from its directly connected LANs. For PIM to work, it must be in one mode, although there is no default mode setting as multicast routing is disabled on an interface by default.

- Dense-mode interfaces are always added to the table. Dense mode is used when multicast group members are densely distributed throughout the network and there is plenty of bandwidth available. Dense mode PIM floods the multimedia packet to all routers and prunes routers that do not support members of that particular multicast group.
- Sparse-mode interfaces are added to the table only when periodic "join" messages are received from downstream routers, or when there is a directly connected member on the interface. Sparse mode is used when members are more spread out and there is limited bandwidth available. Sparse mode PIM relies on rendezvous points. For this purpose the PIM neighbor with the highest IP address is elected to be the Designated Router (DR). If no PIM queries are received from this DR after a certain period of time, the election mechanism will run again.
- Sparse-dense mode interfaces are treated as dense mode if the group is in dense mode, or in sparse mode if the group is in sparse mode.

A significant difference between Dense and Sparse modes is that a dense mode router assumes all other routers are willing to forward multicast packets for a group, while a sparse mode router requires an explicit request for the traffic.

PIM Commands

Enable dense-mode PIM on the interface

```
ip pim dense-mode
```

Enable sparse-mode PIM on the interface

```
ip pim sparse-mode
```

Quality of Service (QoS)

Quality of Service refers to the capability to provide higher levels of access network resources based on the type of traffic. It is defined as being over various network technologies, including Frame Relay, Asynchronous Transfer Mode (ATM), Ethernet and 802.1 networks, SONET, and IP-routed networks. QoS was created to provide:

- Dedicated bandwidth
- Control of jitter and latency
- Enhanced control of potential data loss

Cisco IOS QoS software allows control over complex networks to provide the ability to predictably service a variety of applications and traffic types. It provides:

- Enhanced control over, and more efficient use of, network resources
- Better network analysis management and accounting tools
- The ability to consistently service the most important traffic, while still providing access for less time-sensitive applications
- Enables ISPs to offer tailored grades of service to their customers

There are three fundamental pieces for QoS implementation:

- Within a single network element - queuing, scheduling, and traffic shaping tools
- Signaling techniques for coordinating QoS from end to end between network elements
- QoS policy, management, and accounting functions to control and administer end-to-end traffic across a network

There are three basic levels of end-to-end:

- Best-effort service - basic connectivity with no guarantees
- Differentiated service - soft QoS - some traffic is treated better than the rest via statistical preference
- Guaranteed service - hard QoS - absolute reservation of network resources for specific traffic

There are two traffic-shaping tools:

- Generic traffic shaping (GTS) - GTS reduces outbound traffic flow by constraining specified traffic to a particular bit rate while queuing bursts of the specified traffic
- Frame Relay traffic shaping (FRTS) - FRTS provides parameters useful for managing network traffic congestion: committed information rate (CIR), FECN and BECN, and DE bit

CDP (Cisco Discovery Protocol)

A proprietary Data Link layer protocol used between Cisco devices to pass information about local conditions. CDP uses a data-link, multicast address with no protocol ID or network layer field, and cannot be filtered.

The only way to prevent their being passed is to configure "no cdp enable" on those interfaces on which you do not want to run CDP. You can configure a MAC-layer filter to deny a multicast address as an alternative method to block these packets.

Asynchronous Transfer Mode (ATM)

Developed as a compromise between voice and data needs, ATM is commonly found either on large telecom networks or built into networks that have a strong need for QoS (Quality of Service) needs.

ATM uses *Cells* that are uniform in size - 53 bytes; 5 bytes for a header, and 48 bytes for payload. This allows a great deal of control over traffic and allows for QoS, but is wasteful in that the header is a greater percentage of the traffic than in other methods.

ATM is connection-oriented with traffic traveling from end-to-end over either:

- *SVC (Switched Virtual Circuits)* – Dynamically created on-demand circuits.
- *PVC (Permanent Virtual Circuits)* – Permanently allocated circuits that are always established and active.

There are two types of interfaces:

- NNI (Network-to-Network) – connections within the network cloud between two ATM devices.
- UNI (User-to-Network) – connects a workstation to an ATM switch.

There are four major layers in the ATM reference model (equivalent to the OSI Model):

- Higher layers – ATM signaling, addressing and routing.
- AAL (ATM Adoption Layer) – Converts from higher level to ATM cells.
- ATM – Defines ATM cell relaying and multiplexing.
- Physical – Defines the physical network media and framing.

DDR Dial-on-Demand Routing

DDR has two important applications:

- When there is a WAN link that needs to be available, but rarely sees traffic, the Network Architect might provision a pay-per-use wide area connection - such as BRI - and use DDR on the routers to only activate the link when there is “interesting traffic”, and rip it down when the conversation is over.
- When there is a critical WAN link and there must be a redundant connection. If there were a T1 between two sites, and it was imperative that the link see very little downtime, DDR might be enabled on a BRI ISDN port. If the T1 were to fail, the BRI would establish connectivity over at least one of its data channels (B-channel), and could be configured to enable the second channel if traffic needs were to reach a defined threshold.

DDR spoofs routing tables to provide the image of full-time connectivity using Dialer interfaces and filters out interesting packets for establishing, maintaining, and releasing switched connections. Interesting traffic is defined by an access list.

Encapsulation Methods for DDR:

- *PPP* – recommended, as it supports multiple protocols and is used for synchronous, asynchronous, or ISDN connections. It is also non-proprietary.
- *HDLC* - supported on synchronous serial lines and ISDN connections only, and supports multiple protocols, with NO authentication.

- *SLIP* - works on asynchronous interfaces and is IP only, with NO authentication.
- *X.25* - works on synchronous serial lines and a single ISDN B channel.

IEEE Protocols

IEEE - The Institute of Electrical and Electronics Engineers, a professional organization that, among other things, developments communications and network standards.

IEEE 802.1 - IEEE specification that describes an algorithm that prevents bridging loops by creating a spanning tree. The algorithm was invented by Digital Equipment Corporation. The Digital algorithm and the IEEE 802.1 algorithm are not exactly the same, nor are they compatible.

IEEE 802.2 - IEEE LAN protocol that specifies an implementation of the LLC sublayer of the data link layer. IEEE 802.2 handles errors, framing, flow control, and the network layer (Layer 3) service interface. Used in IEEE 802.3 and IEEE 802.5 LANs.

IEEE 802.3 - IEEE LAN protocol that specifies an implementation of the physical layer and the MAC sublayer of the data link layer. IEEE 802.3 uses CSMA/CD access at a variety of speeds over a variety of physical media. Extensions to the IEEE 802.3 standard specify implementations for Fast Ethernet. This is the specification that describes Ethernet.

IEEE 802.4 - IEEE LAN protocol that specifies an implementation of the physical layer and the MAC sublayer of the data link layer. IEEE 802.4 uses token-passing access over a bus topology and is based on the token bus LAN architecture. This is the specification that describes Token Ring Bus.

IEEE 802.5- IEEE LAN protocol that specifies an implementation of the physical layer and MAC sublayer of the data link layer. IEEE 802.5 uses token passing access at 4 or 16 Mbps over STP cabling and is similar to IBM Token Ring.

IEEE 802.6 - IEEE MAN specification based on DQDB technology. IEEE 802.6 supports data rates of 1.5 to 155 Mbps. This is the specification that describes Municipal Area Networks (MAN).

More 802.x standards

802.7	Broadband
802.8	Fiber-optic LANs
802.9	Integrated Voice & Data
802.10	LAN/MAN Security
802.11	Wireless
802.12	VGAnyLAN (HP's answer to FastEthernet)

Tag Switching

Tag switches support multicast by utilizing data link layer multicast capabilities: all tag switches that are part of a given multicast tree on a common sub-network must agree on a common tag so that forwarding of multicast packets to all downstream switches on that sub-network is possible.

Tag switching can mark packets as belonging to a particular class after they have been classified the first time, which is an important aspect of QOS.

The tag-switching forwarding paradigm is based on label swapping, which is the same as in ATM forwarding, tag-switching technology can be applied to ATM switches.

Tag information can be carried in a packet in many ways, such as a small "shim" tag header inserted between the Layer 2 and the network-layer headers, as part of the Layer 2 header (e.g. ATM), as part of the network-layer header (e.g. Ipv6). This is why tag switching can be implemented over any media type.

When a packet with a tag is received, the switch uses the tag as an index in its Tag Information Base (TIB). If the switch finds a matching entry, then for each component in the entry the switch replaces the tag in the packet with the outgoing tag. The switch also replaces the link-level information in the packet with the outgoing link-level. This is called Label Swapping.

There are two principal components:

- *Forwarding component* - uses tag information carried by packets and the tag-forwarding information maintained by a tag switch to perform packet forwarding.
- *Control component* - maintaining correct tag-forwarding information among a group of interconnected tag switches.

Remote Monitoring (RMON)

RMON has 4 groups:

- Statistics Group for port utilization and error statistics

- History Group for periodic statistics
- Alarm Group for sampling interval and threshold
- Event Group for logging events to network management station

Cisco Device Management

There are two ways of managing routers and switches:

- In-band management – Telnet or SNMP network connection through modem or line module.
- Out-of-band management – The console port direct connection to the Supervisor module.

General Troubleshooting Tips

- SPAN is the Enhanced Switched Port Analyzer that monitors traffic for analysis by other tools.
- CWSI CiscoWorks Switched Internet Solutions is a management suite that consists of CiscoView, VlanDirector, and TrafficDirector.
- A cable tester device is used to look for cable breaks.
- Time Domain Reflectometer measures cable length and impedance; loose or incorrect device connection can also be detected.
- Always isolate network segment problems by checking the devices on the same segment to see if they can communicate. In an IP environment, use the "ping" and "traceroute" commands.
- Switch LEDs indicate problem based on color: red = failure, orange = less severe problem. If the Output Fail LED = Red, check the power supply.
- To troubleshoot other problems, try using the show commands to find out what is going on: Sh config, Sh int, Sh module, Sh spantree, Sh trunk, Sh vlan, Sh port, Sh mac, Show test and Show log, etc.
- A Protocol Analyzer can capture and display protocol information, while Network monitors can continuously monitor network traffic.
- A STP failure generally results in a bridging loop.
- For point-to-point links, a duplex mismatch occurs when one side of the link is hardcoded full duplex, while the other side is auto-negotiation, and eventually the link ends up in half-duplex.

- When a link is experiencing many physical errors, a number of consecutive BPDUs could be lost, leading a blocking port to transition to forwarding.
- STP is software based: if the CPU is over-utilized, it is possible that it can lack the resources necessary to send out BPDUs. Also, software bugs are possible.
- When the age field of a BPDU goes beyond max age, it is discarded - this occurs if the diameter of the STP network is too large, making the root switch too far from some distant switches.
- To limit the risk implied by the use of the STP, it is recommended that you reduce (as much as possible) the number of blocked ports - prune VLAN not needed off your trunks and use the PortFast command on those user ports that will never have switches or bridges installed.
- Keep traffic off the administrative VLAN and avoid having a single VLAN spanning the entire network.
- Avoid hand-tuning STP parameters - Catalyst software provides macros that perform fine-tuning of most significant STP parameters.

IOS commands for troubleshooting:

- "debug spantree events" displays STP events to help determine problems. Be careful that this doesn't overwhelm the CPU of the switch.
- "logging buffered" captures debug information in the device's buffers.
- "show interface" verifies interface utilization, packet corruption, speed and duplex status of the specified port.
- "show processes cpu" checks CPU utilization.

Catalyst OS Commands for troubleshooting:

- "set logging level spantree 7 default" increases the default level of STP related event to debugging.
- "set logging buffer 500" sets a maximum number of messages in the switch's buffers.
- "show port <module#/port#>" give you details of the port configuration.
- "show system" give indication on the backplane utilization.
- "show spantree statistics <module#/port#> <vlan#> gives accurate information on suspected ports.

Hot Standby Routing Protocol (HSRP)

Provides a means of having two default gateways to protect against an equipment failure locking out a group of users from the wider internetwork.

The default priority for each router is 100, but can be change to give one priority as the most likely default gateway (if say, one unit were faster than another).

TAC and CCO

Cisco's Technical Assistance Center (TAC) provides 7x24x365 technical support on all their products. The Center follows the sun, with offices around the globe. It is staffed by Customer Support Engineers (CSEs).

Cases can be opened by phone, e-mail or through Cisco Connection Online (Cisco's exceptionally well designed website at www.cisco.com). When a case is logged, a call number is generated and assigned to a CSE who will work with client to answer questions, provide advice on system use, help with system configuration, or correct a system malfunction.

There are four priority levels:

- Priority 1 - existing network is "down", which is critical
- Priority 2 - network is severely degraded, which has a significant impact
- Priority 3 - operational performance of the network is impaired, although business operations remain functional
- Priority 4 - little or no impact to the business operation at this moment

The OSI Model

The OSI is a common tool for conceptualizing how network traffic is handled. In this document we will be interested primarily in the lower four levels. Just a reminder, that you can use the old mnemonic "**A**ll **P**eople **S**eem **T**o **N**eed **D**ata **P**rocessing" as a way to help remember the sequence.

7. Application – User interface tools (such as Telnet, SMTP, FTP, etc.)
6. Presentation – Encoding/Decoding (such as ASCII, MPEG, GIF, JPEG, etc.)
5. Session – Creating, managing and terminating Presentation layer
4. Transport – Error checking and recovery, flow control and multiplexing (TCP, SPX, etc.)
3. Network – Routing (IP, IPX, etc.)
2. Data Link (LLC/MAC)

LLC – Manages communications
MAC – Manages addressing and access to the physical layer

1. Physical – Establish and maintain physical connectivity

Cisco Hierarchical Internetworking Model

- *Access* – The point at which users join the network. VLANs, WAN connections, RAS services are all at this layer. Cat1900 or 3500 series switches with 10BaseTx and 100BaseTx ports might be appropriate for the Access layer, where high port density and low per-port costs are major concerns.
- *Distribution* – Control layer; Aggregation of traffic, access lists, compression, encryption and other services that provide the glue between Access and Core layers. Cat6000 series with 16-port Gigabit Ethernet modules, or Cat5500 series switches with internal RSMs might be appropriate at the Distribution layer. Performance becomes more of an issue at this layer.
- *Core* – Concentrates all traffic traversing the network. The focus is on speed. Fast switching, Gigabit Ethernet, and ATM are commonly deployed at this layer. Cat8500 or 6500 series switches - high speed and expensive - are probably most appropriate at the core, where passing traffic is the primary concern.

Special thanks to

Dennis Laganiere for contributing to this Cramsession and Michael Yu for the original document. Please visit their sites at http://www.michaelyu.freeservers.com/

http://www.routedpacket.com/