



Operational Best Practices for the Cisco Catalyst 6500 Series

1. INTRODUCTION

Companies today place high demands on network infrastructure. The ability to maintain performance, availability, security, and manageability of those network devices is of paramount importance to the successful operation of their data center. With the wide-scale deployment of the Cisco® Catalyst® 6500 Series in networks worldwide, Cisco Systems® has been able to better understand the demands placed on the switching infrastructure. More importantly, Cisco Systems has been able to direct the development of the Catalyst 6500 family of switches to meet the challenges corporate networks place on these devices.

This document will attempt to present a best practices guide for operations management for features that customers can implement on the Catalyst 6500 to meet these challenges. It will focus on the areas of performance, scalability, security, availability, and manageability. Examples of the command-line interface (CLI) will be used in strategic places to show how commands can be used to deploy the proper best practice guidelines. This document will predominantly use command examples from the Cisco IOS® Software CLI, although references to the Catalyst OS commands will be made where appropriate.

2. BEST PRACTICES: MANAGEABILITY

Some of the switch management best practices that follow tend to be a low-focus area for many network administrators. While some of these recommendations are not critical, they do make it easier for the operations group to manage the network. The following set of management features is suggested for implementation.

2.1 Configure Hostname

The default hostname of the Catalyst 6500 is “Router.” Consideration should be given to choosing a name that reflects some meaning to operations people managing the switch. A common practice among many customers is to incorporate a reference to the location (for example, building/floor/wiring closet) or function (for example, core/distribution/server farm) that this switch serves. Having a well-thought-out naming convention will prove helpful for administrators when using network management software and when troubleshooting the network.

Changing the hostname can be achieved using the following command:

```
Router(config)# hostname ?  
WORD This system's network name  
Router(config)# hostname 6K-LV2-CL3  
6K-LV2-CL3#
```

In this example, after the hostname is entered, the CLI prompt changes to reflect the entered name. The hostname is derived from the device type (Catalyst “6K”), floor location (Level 2), and wiring closet location (Closet 3). Sometimes the name can also incorporate the supervisor type or other relevant information (for example, S720B-LV2-CL3).

2.2 Configure Login Banner

The configuration of the login banner provides a way to display a message to anyone who accesses the switch. Login banners can be useful to reemphasize the potential law infringement represented by the unauthorized access to the device in question. It can also be used to provide information about the location of the device, the contact details of the administrator, or a message of the month. The login banner is set as follows:

```
6K-LV2-CL3(config)# banner login ^d Access to this device or the attached networks is
prohibited without express permission from the network administrator.
```

```
Violators will be prosecuted to the fullest extent of both civil
and criminal law.^d
```

When users access the switch configured with the preceding banner, they will be greeted with this login message.

2.3 Sync Up the System Clock to Syslog Messages

This is a simple but often overlooked task. If system logs are kept, it is definitely worth setting the system clock so that log records are time-stamped accordingly. The system clock is set from the enable prompt as follows:

```
6K-LV2-CL3# clock set 15:10:00 2 May 2006
```

```
6K-LV2-CL3# show clock
```

```
15:10:04.691 UTC Tue May 2 2006
```

Similarly, in the Catalyst OS the corresponding command to set the clock is performed as follows:

```
Sup720BXL-DC> (enable) set time Thursday 05/02/2006 09:07:30
```

```
Thu May 02 2002, 09:07:30
```

```
Sup720BXL-DC> (enable) show time
```

```
Thu May 02 2006, 09:07:38
```

2.4 Use Network Time Protocol to Synchronize Time

Debugging or troubleshooting network problems is not easy at the best of times, but the level of difficulty increases in diagnosing problems that occur across multiple nodes. Implementing Network Time Protocol (NTP) provides a means to ease the burden of troubleshooting in this scenario. NTP will synchronize the clocks of all nodes such that time-stamped SYSLOG records across all devices are time synchronized. Trying to determine when a particular event occurs and the effect it has across multiple nodes is made easier when all device clocks are in sync. NTP is thus deemed to be a critical requirement for customer networks that want to facilitate improvements in the monitoring of networks, troubleshooting, and debugging.

An example of an NTP configuration is shown below. This example shows the switch will peer with an NTP server whose IP address is 172.16.1.3 and use a predefined key of cisco123 to authenticate against this NTP device.

```
6K-LV2-CL3(config)# ntp peer 172.16.1.3 key 1011 source loopback 1
```

```
6K-LV2-CL3(config)# ntp authenticate
```

```
6K-LV2-CL3(config)# ntp authentication-key 1011 md5 cisco123
```

A common use of the NTP synchronized time is in debug and log time-stamps, which can be enabled as follows:

```
6K-LV2-CL3(config)# service timestamps debug datetime msec localtime show-timezone
6K-LV2-CL3(config)# service timestamps log datetime msec localtime show-timezone
```

The corresponding Catalyst OS version of the example above would look like this:

```
Sup720BXL-DC> (enable) set ntp client enable
Sup720BXL-DC> (enable) set ntp authentication enable
Sup720BXL-DC> (enable) set ntp key 1011 trusted md5 cisco123
Sup720BXL-DC> (enable) set ntp server 172.16.1.3 key 1011
```

2.5 Set Simple Network Management Protocol Parameters

Simple Network Management Protocol (SNMP) provides the transport vehicle for allowing a network management platform to both extract information from the switch and also make changes to elements of the configuration. It is recommended that after SNMP is enabled, the descriptive components of SNMP should also be completed. This will allow SNMP management applications to provide clearer descriptions of devices and better identification for the originator of SNMP alerts.

Three SNMP configuration elements should be added if SNMP is enabled. Those SNMP entities are “contact,” “location,” and “chassis-id.” These are shown in the following CLI example.

```
6K-LV2-CL3(config)# snmp-server contact ?
LINE  identification of the contact person for this managed node

6K-LV2-CL3(config)# snmp-server location ?
LINE  The physical location of this node

6K-LV2-CL3(config)# snmp-server location chassis-id ?
LINE  <cr>
```

Both contact and location are self-explanatory. Chassis-ID is used to provide a unique description of that device.

With the Catalyst OS the same information can be entered by using the set system command.

2.6 Backup Cisco IOS Software and Configuration

Normal operation of the switch requires an image to boot the system. Bootable images are typically stored on the supervisor bootflash located on the supervisor module. The switch can load an OS image from local flash memory (for example, compact flash), or from a device reachable through the network such as a TFTP/FTP server. It is a good practice to have a backup copy of the current OS image located either on the supervisor flash or on a separate compact flash. Both Cisco Catalyst 6500 Series Supervisor Engine 720 and Cisco Catalyst 6500 Supervisor Engine 32 support a compact flash slot on the front panel. Should the image on bootflash be deleted accidentally or become inaccessible, locating a backup image on the flash card will save time in getting the switch back up and running.

There is a 16-bit configuration register that is used to instruct the switch at bootup time where to locate the boot p image. The default configuration register value on the Catalyst 6500 is 0x2102 (the register is stored as a hexadecimal number). The default value instructs the switch to boot using an image specified by the **boot system** command that is found in the configuration file. The Catalyst 6500 supports multiple **boot** commands in the configuration file. The switch will work in a top-down fashion, starting with the first image in the list. If that image is not accessible, it will try booting the image on the second and subsequent lines.

The following gives an example of such **boot** commands:

```
6K-LV2-CL3(config)# boot system flash sup-bootflash:s72033-jk9sv-mz.122-18.SXD.bin
6K-LV2-CL3(config)# boot system flash disk0:s72033-jk9sv-mz.122-18.SXD.bin
```

The commands above only differ in the location of the OS image. The first line points to “sup-bootflash,” which refers to the onboard flash resident on the supervisor module. The Catalyst 6500 has two flash memories, one being the route processor flash and the other being the switch processor flash. The switch processor flash (“sup-bootflash”) must be used to boot a Cisco IOS Software image in the so-called “native” configuration. The route processor flash instead is referred to simply as “bootflash.”

The second line in the above example points to the same image, which is located on compact flash. The compact flash drive is referred to as disk0. The Supervisor Engine 720 has a second compact flash slot, and this is referred to as disk1. In both cases, sup-bootflash and disk0 use a “:” as an ending delimiter in the **boot** command. This must be present for the system to successfully locate the device.

Another tip to consider when typing in the **boot** command is to copy and paste the image name after it is loaded into the **boot system flash** command. Failure to type the exact name of the image as it is known on the storage device will result in the system failing to find the image and potentially have it fall back to ROMMON mode. After an image is loaded onto the switch, one should issue a **show sup-bootflash** command to see and verify the image name. Then one could simply highlight the image name, copy it, and then paste it into the **boot system flash** command. An example of the **show** command is shown below:

```
6K-LV2-CL3# show sup-bootflash:
-#-  ED  ----type----  --crc---  -seek--  nlen  -length-  -----date/time-----  name
1    ..  image          C10CAA82  2E7D270    30  76599844  May 02 2006 19:56:17 +00:00 s72033-
advipservicesk9_wan-mz.122-18.SXF.bin
```

In addition to having a copy of the OS image on compact flash, it can prove useful to have a copy of the configuration located there as well. If no configuration management tool is used, then copying the configuration to the bootflash or compact flash drive prior to making a configuration change could help you out of an awkward situation, should a mistake be made.

Similarly, in Catalyst OS the set boot system flash command can be used to configure the boot image(s) thus:

```
Sup720BXL-DC> (enable) set boot system flash
Usage: set boot system flash device:[filename] [prepend] [mod]
```

With the Catalyst OS, the switch processor bootflash is simply called “bootflash:” or “bootdisk:” if on a Supervisor Engine 32. If multiple boot statements are found in the configuration, the system will move top down through the list until it finds an accessible and working image. Adding an additional boot statement to an existing set will see this new boot statement added at the bottom of the list. If this boot statement were required to be located at the top of the list, the “prepend” option can be used to force this boot statement to the top of the list.

2.7 Reduce Time to Track Down Problems

For a network administrator to successfully troubleshoot a network problem, the administrator might need to take a sniffer to the troublesome network node to assist in resolving the problem. Remote Switched Port Analyzer (SPAN) is one option that can be used to reduce the burden of physically going to the wiring closet to connect the sniffer. SPAN is a feature of the Catalyst 6500 that can redirect a copy of the data from a source port to a destination port for analysis by an attached probe or sniffer or by an IDS appliance or blade. While SPAN requires the source and destination ports to be local on the same switch, Remote SPAN (RSPAN) allows the source and destination ports to be located on different switches.

Remote SPAN uses a reserved SPAN VLAN to transmit a copy of the data from the source port across the network to the destination port. All network devices in the path between the source port and the destination port must have this VLAN defined and active. The idea behind this suggestion is to partially set up the RSPAN VLAN on relevant network switch nodes and complete the destination RSPAN configuration on a local switch where a sniffer can be located. Should a problem arise, all that is needed is to complete the configuration by defining the RSPAN source port, and traffic will start flowing to the sniffer. This eliminates having to visit wiring closets and also provides a more timely way to start resolving any network issues that might arise.

The main configuration elements that are required to be set up are shown below.

1. First a VLAN needs to be created on each network node in the path of the source and of the destination port as follows:

```
6K-LV2-CL3(config)# vlan 555
6K-LV2-CL3(config-vlan)# remote-span
```

This example sets up VLAN 555 as a remote span VLAN.

2. Set up the RSPAN destination port on the local administrator's switch (this is where the sniffer would be permanently located) as follows:

```
6K-LV2-CL3(config)# monitor session 1 source remote vlan 555
6K-LV2-CL3(config)# monitor session 1 destination interface g3/1
```

This configuration statement sets up port 1 on module 3 to receive SPAN traffic from RSPAN VLAN 555.

3. If, for example, traffic for a set host on a given switch started to show signs of trouble, the RSPAN session could be completed as follows:

```
6K-LV2-CL3(config)# monitor session 1 source interface g9/16
6K-LV2-CL3(config)# monitor session 1 destination remote vlan 555
```

2.8 NetFlow Data Export for Traffic Analysis

NetFlow is a feature of the Catalyst 6500 that allows statistics to be collected for flows that pass through the switch. Traffic for each flow is monitored and a flow record is maintained for data associated with that flow. NetFlow records are maintained for data that is processed by both the policy feature card (PFC) and the multilayer switch feature card (MSFC). All supervisors can collect NetFlow records for routed traffic. With the introduction of the PFC3B and PFC3BXL daughter cards, the Supervisor Engine 720 can be configured to collect statistics for both bridged and routed traffic. (When using Catalyst OS, PFC and PFC2 support the same capability.)

NetFlow record analysis has often been used in conjunction with other features to facilitate troubleshooting, problem diagnosis, and detection of security breaches. For instance, some universities use NetFlow records to track down students who initiate denial-of-service attacks on campus network devices. Cisco IT routinely uses NetFlow to detect network issues such as worm outbreaks. Therefore, NetFlow can be considered both a management/traffic analysis feature and a security feature.

The Catalyst 6500 currently exports NetFlow records in the v5, v7, or v8 format (whereas the v9 format is planned for a future software release, as of this writing). A collector must be used to receive the NetFlow records exported from the Catalyst 6500. The Network Analysis module (WS-SVC-NAM-1 and WS-SVC-NAM-2) is an example of a NetFlow collector. The NetFlow collector is used to store and collate the NetFlow records and pass them to a NetFlow analyzer software component, which will present the traffic statistics in a tabular or graphical format (for example, inside a GUI window).

It is highly recommended that the collection of NetFlow records be enabled and a suitable collector be implemented to collate exported records for later analysis.

NetFlow Data Export (NDE) can be configured on the switch as follows:

```
6K-LV2-CL3(config)# mls flow ip full /*sets the desired NetFlow hw mask
6K-LV2-CL3(config)# mls nde sender version 5 /*enables NDE on the PFC w/ v5 format
6K-LV2-CL3(config)# ip flow-export version 5 /*sets v5 format on the MSFC as well
6K-LV2-CL3(config-if)# ip flow ingress /*enables NDE on the MSFC
6K-LV2-CL3(config)# ip flow-export source <source interface>
6K-LV2-CL3(config)# ip flow-export destination <ip-address> <UDP port number>
```

The commands above enable NetFlow Data Export for flows processed in hardware (on the PFC) and for flows handled in software (on the MSFC) in addition to defining the destination IP address(es) of the device(s) to which the NetFlow records will be exported. (To configure redundant NDE data streams, which improves the probability of receiving complete NetFlow data, starting from Cisco IOS Software Release 12.2(18)SXE for Supervisor Engine 720 and Release 12.2(18)SXD for Supervisor Engine 2, one can enter the **ip flow-export destination** command twice and configure a different destination IP address in each of the two commands.)

The NetFlow collector/analyzer component plays a crucial role for the network administrator to be able to examine and interpret the bulk of raw data extracted by the NetFlow-enabled devices. For this task, Cisco Systems offers both a Flow Collector engine and a Data Analyzer engine. Third-party software vendors offer an ample choice of very specialized applications based on NetFlow: for example, Arbor Networks sells the security-oriented and very popular PeakFlow application, NetQoS offers a powerful QoS-oriented solution, Crannog Software sells advanced traffic profiling tools, and so forth. In addition, many freeware or shareware collectors are available on the Web.

If security is the user's major concern, Cisco Systems offers a very powerful security monitoring, analysis, and reporting appliance-based solution called Cisco Security MARS (which is available in various models and which was formerly sold by the Protego acquisition). Cisco Security MARS can collect and correlate data derived from NetFlow statistics as well as from the IDS/IPS and FWSM blades and from the CSA software component (for example, it can use syslog data generated by traffic hitting a deny plus log statement on the firewall). This combination of strengths offers superior flexibility in terms of traffic profiling and statistical anomaly detection for security applications.

It is recommended to evaluate the deployment of one of the aforementioned tools to best use the benefits of the NetFlow feature on the switches.

2.9 Power Redundancy

When a Catalyst 6500 is primed with two power supplies, the default mode of operation is to run the units in combined mode. Running the switch in combined mode provides 167% of the total capacity of the two power supplies. One consideration for running in combined mode is that as more line cards and inline devices are added to the switch, a time could arise when the power load required is greater than what a single power supply can provide. A power failure to one power supply could cause the shutdown of inline-powered devices and/or of the line cards themselves.

For this *reason it is recommended that power supplies be run in redundant mode*. This mode of operation will allow the redundant power supply to take over and provide power for the active configuration should the primary power supply fail. The power mode can be viewed using the following command:

```
6K-RACK2-REDUN# show power
system power redundancy mode = combined
system power total =      1921.92 Watts (45.76 Amps @ 42V)
system power used =      1173.06 Watts (27.93 Amps @ 42V)
system power available =  748.86 Watts (17.83 Amps @ 42V)
<snip>
```

Changing the power mode to redundant is as simple as entering the following command:

```
6K-RACK2-REDUN(config)# power redundancy-mode ?
  combined  combine power supply outputs (no redundancy)
  redundant either power supply can operate system (redundancy)
6K-RACK2-REDUN(config)# power redundancy-mode redundant
```

2.10 Power Shutdown Sequence

The previous section highlighted the recommendation for running in redundant mode. Should there be a requirement to run in combined mode, an extra precaution needs to be taken. In the event of a failing power supply that leaves the system without enough power to drive the total configuration, inline devices and line cards will be shut down by the power management software to fall within operating power capabilities.

In order for the user to recognize the event and its consequences, the shutdown sequence logic should be known and understood: in particular, the Catalyst 6500 control software will first start to power down inline devices from the highest numbered port to the lowest numbered port, then from the bottom slot up. If after powering down all inline devices there is still insufficient power, the system will start to power down line cards from the bottom slot up with the exception of the supervisor module and of the services modules (that is, firewall module, VPN module, IDS module, and so on). Supervisor and services modules are the last modules left operating. For this reason, important hosts should be plugged into ports on the top-most modules (that is, slots 1, 2, 3, and so on) and from the lower numbered ports up. It is worth noting that the order of shutdown is fixed by the system and cannot be changed.

3. BEST PRACTICES: SECURITY

The need for increased security is evident with the ever increasing threat of worms, viruses, hackers, and denial-of-service attacks. There are a number of ways in which a Catalyst 6500 can be hardened against denial-of-service attacks to minimize the effect on attached devices. The following is a list of best practices to better protect the Catalyst 6500.

3.1 Disable Cisco Discovery Protocol

Cisco Discovery Protocol is an effective management, configuration, and troubleshooting tool that is used in a number of ways by the Catalyst 6500. Inline-powered devices, for example, use Cisco Discovery Protocol for device discovery and configuration. Cisco Discovery Protocol is also an effective management tool that gives an insight into the device that is connected on the other end of a link. An example of what information Cisco Discovery Protocol provides is shown in the following output:

```
6K-LV2-CL3# show cdp neighbor detail
-----
Device ID: c6503
Entry address(es):
  IP address: 192.168.10.1
Platform: cisco WS-C6503, Capabilities: Router Switch IGMP
Interface: GigabitEthernet1/1, Port ID (outgoing port): FastEthernet3/1
Holdtime : 144 sec

Version :
Cisco Internetwork Operating System Software
IOS (tm) s72033_rp Software (s72033_rp-JK9SV-M), Version 12.2(18)SXD, RELEASE SOFTWARE (fc2)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2004 by cisco Systems, Inc.
Compiled Thu 29-Jul-04 01:36 by cmong

advertisement version: 2
VTP Management Domain: 'ENG'
Native VLAN: 211
Duplex: full
```

The show output above provides a comprehensive set of information about devices directly connected to this switch.

In Catalyst OS a very similar command is available:

```
Sup720BXL-DC> show cdp neighbors detail
Port (Our Port):4/4
Device-ID:69046406
Device Addresses:
```



```
IP Address:172.20.25.161
Holdtime:150 sec
Capabilities:TRANSPARENT_BRIDGE SWITCH
Version:
  WS-C6009 Software, Version NmpSW: 5.4(1)CSX
  Copyright (c) 1995-1999 by Cisco Systems
Port-ID (Port on Device):4/8
Platform:WS-C6009
VTP Management Domain:unknown
Native VLAN:1
Duplex:half
```

There may be situations where the Catalyst 6500 is used to connect to partner companies through an extranet, or to nodes in the network that have no need for this level of knowledge. In such cases, this information could potentially prove useful to someone trying to access the device to initiate an attack. Therefore, Cisco Discovery Protocol can be disabled on a per interface basis to selectively prevent this information from being obtained by untrusted adjacent devices.

It is recommended that Cisco Discovery Protocol only be enabled on switch ports whose neighbors require this level of information. Cisco Discovery Protocol should be disabled instead on any Catalyst 6500 switch port that connects to an untrusted device. Other more controlled management tools can be used as an alternative to provide insight into the capabilities of untrusted network devices.

3.2 Protect Against TCP SYN Flooding Attacks and Other Advanced Attacks

A TCP SYN flood attack is a form of denial-of-service (DoS) attack. It works by targeting a host with a continuous stream of new TCP sessions. The attack works by the attacker initiating numerous TCP sessions but not completing the initiation process. The device under attack will have to process and hold state for all the incoming TCP connections, which can result in reducing that device's ability to communicate with other devices in the network.

The Catalyst 6500 supports features that can prevent TCP SYN flooding attacks in services modules like the firewall services module (FWSM) and the content switching module (CSM). These services modules offer an effective way to mitigate this form of attack. The FWSM, for instance, supports an implementation of the TCP intercept feature that uses the "SYN cookies" algorithm and the concept of embryonic connection limit to stop TCP SYN attacks. When the embryonic limit is exceeded, every inbound SYN is intercepted by the FWSM, which responds on behalf of the target device with a SYN/ACK. The FWSM maintains pertinent state information, drops the packet, and waits for the client's acknowledgement. If the ACK is received, then a copy of the client's SYN is sent to the server and the normal three-way TCP handshake is completed between the FWSM and the server and the connection resumes as per normal. If, however, the client does not acknowledge the ACK, then the FWSM will drop that session.

An analogous mechanism is available on the CSM, which allows the user to configure a threshold (in number of pending sessions) beyond which the SYN cookies algorithm is engaged. With both the FWSM and the CSM, this threshold-based algorithm allows the SYN flood protection to kick in only when needed without affecting the regular three-way TCP handshake (and related TCP optimization features/options) as long as the number of half-open sessions stays below a predefined limit.

As a matter of practice, it is *recommended that these types of SYN flood protection features be used when possible.*

Furthermore, when considering a more general scenario than the above one, a first line of protection against generic network intrusion attempts can be deployed by making use of the inherent capability of the Catalyst 6500 to run access control lists (ACLs) in hardware. The Catalyst 6500 supports regular Cisco IOS Software IP ACLs (also known as router ACLs, or RACLs), VLAN-based ACLs (also known as VACLs), and port-based ACLs (PACLs; supported on PFC3 only). Any of these types of ACLs can be chosen based on the specific application to enforce a first level of traffic filtering to block spoofing attempts or unauthorized network accesses.

Moreover, the Catalyst 6500 supports another category of hardware-based ACLs, called QoS ACLs, that can be used in conjunction with QoS policers to rate limit certain categories of traffic (like ping messages) and avoid their potential use in DoS attacks.

All of the aforementioned ACL features are run in hardware in the forwarding engine and can also be combined together for a more effective filtering action. On top of that, they do not require dedicated hardware to be supported, because they are part of the standard capabilities of a Catalyst 6500 system.

There is a caveat: although hardware-based ACLs are a great way of enforcing security policies at wire speed, they have some notable limitations: they are stateless and their logging capabilities are still CPU bound.

With the introduction of PFC3B/BXL one of their limits has been eliminated, namely the lack of accurate ACE hit counters. This capability is critical for network issue diagnosis and troubleshooting and is therefore highly recommended for all deployments making use of hardware-based ACLs.

The second and more sophisticated line of defense that can be deployed is a firewall device such as the firewall services module (FWSM). In particular, an FWSM possesses several capabilities that are extremely useful but that are not part of the tool set of the switch forwarding engine's ACL acceleration logic: statefulness and deeper flow inspection.

A hardware-based ACL on a Supervisor Engine 720 performs packet by packet checks and as such is incapable of detecting any sort of advanced attacks relying on sequence number variations, invalid flag combinations, time-to-live (TTL) evasion, or checksum errors. On the other hand a firewall's bread and butter is represented exactly by these advanced detection tasks, and the FWSM is capable of executing them at reasonably high speeds (the FWSM does not support TCP normalization yet, as of this writing, but it is planned for Release 3.1).

Besides, an FWSM module is capable of providing the user with a detailed session log for a hit against an ACE rule with a peak performance of 35K syslogs per second, a performance that is significantly higher than what the CPU-bound logging mechanism of the supervisor engine can achieve.

As a general recommendation, it is suggested to explore the application of the FWSM or other firewalling device (typically in a redundant configuration) in the network to use the advanced features that this type of technology can offer.

In addition, for large server farms a very advanced and scalable anti-(d)DoS solution is represented by the new Traffic Anomaly Detector and Anomaly Guard services modules. These modules are capable of detecting and blocking malicious traffic flows in real time without affecting the other legitimate traffic. Thanks to their sophisticated security algorithms, they represent the innovative DoS protection technology for demanding applications, and as such they are highly recommended.

3.3 Secure Access to the Switch CLI

The switch provides console and telnet access as the two primary ways to access the command line interface (CLI). The switch is generally located in a room with access controlled by some form of security (for example, badge reader, keypad, and so on). In this regard, console access tends to be secured with only authorized personnel having access to the room. The other means to access the switch CLI is through telnet access. Most organizations simply secure telnet access with a local password. The problem herein lies with the fact that most telnet passwords do not conform to any form of stringent security guidelines and are more than often based on word strings that can be cracked with a brute-force dictionary-based attack.

Telnet security can be improved by using either (preferably both) of the following methods. access control lists (ACLs) can be used to identify access from known hosts in the network. Telnet connections from hosts outside of the ACL permit range will simply be denied. This in itself is only part of the recommended solution. Further reinforcing telnet security can be achieved by using an external AAA server. This method will prompt requests from valid devices for a username/password combination. Together, these features provide a more secure solution for reducing unauthorized access to the switch CLI.

An example of one possible configuration with a strictly controlled access to the system is reported below:

```
6K-LV2-CL3(config)# service tcp-keepalives-in
6K-LV2-CL3(config)# service tcp-keepalives-out
6K-LV2-CL3(config)# line vty 0 4
6K-LV2-CL3(config-line)# access-class 101 in
6K-LV2-CL3(config-line)# exec-timeout 15 0
6K-LV2-CL3(config-line)# login local
6K-LV2-CL3(config-line)# transport input ssh
6K-LV2-CL3(config-line)# no transport output
6K-LV2-CL3(config)# access-list 101 remark VTY Access ACL
6K-LV2-CL3(config)# access-list 101 permit tcp host <admin-host-ip> host 0.0.0.0 eq 22
log-input
6K-LV2-CL3(config)# access-list 101 deny ip any any log-input
```

3.4 Secure Transmission of Console Data

Tools that can launch man in the middle (MiM) attacks are a common means used by hackers to get access to active telnet sessions. Not only can these tools allow a third party to view what is being displayed in that session, but they also provide a way to capture the username and password for that session. A simple and effective way to stop this type of threat is to use features like Dynamic ARP Inspection (discussed later in this paper). However, other means are still available to capture the data in transit.

For this reason, it is recommended that Secure Shell (SSH) become the standard way of accessing a console remotely. SSH provides the means to encrypt the data so that any traffic captured in transit is going to be protected from view. The Catalyst 6500 supports both version 1 and version 2 of SSH. Version 2 is generally considered to be the more solid implementation of SSH and therefore should be used for implementation.

SSHv2 can be configured on the switch/router as follows:

```
6K-LV2-CL3(config)# crypto key generate rsa /*recommended min modulus size: 1024
6K-LV2-CL3(config)# ip ssh version 2 /*enables version 2
```

3.5 Protect the Switch Control Plane

Like other network devices, the Catalyst 6500 processes packets in two places, in the hardware and in the switch CPUs. The so-called data plane corresponds to the aggregate of the user traffic and is typically processed in hardware. The control plane term instead refers to the aggregate of the protocol traffic, which is typically processed by the switch CPUs. The performance of the control plane is limited by the speed of the CPUs and is orders of magnitude less than what the data plane capacity can be. For this reason, it is important to protect the operational integrity of the control plane. Compromising the control plane can affect the switch performance even if the capacity of the data plane is underutilized.

Initiators of denial-of-service attacks know very well that a network device can be compromised if its control plane is overwhelmed. Cisco Systems has recognized the need to protect the control plane with the introduction of two new features available on Supervisor Engine 720 and on Supervisor Engine 32: special hardware-based rate limiters and Control Plane Policing (CoPP). These new capabilities offer the means to protect the switch control plane from overloads caused either accidentally or maliciously.

3.5.1. Special Purpose Hardware Rate Limiters

Built into the PFC3 hardware of Supervisor Engine 720 and Supervisor Engine 32 are 10 hardware registers that can be utilized to implement a rate limiting action for specific traffic types. A rate limiter is enabled defining a packet per second threshold for a given type of CPU-bound traffic type. Traffic matching that traffic type in excess of the defined rate is simply dropped.

Layer 2 and Layer 3 rate limiters can be applied at the same time on the switch. The rate limiters that can be enabled by default are shown in the following CLI output highlighted in BOLD:

```
6K-LV2-CL3# show mls rate-limit
```

```
Sharing Codes: S - static, D - dynamic
```

```
Codes dynamic sharing: H - owner (head) of the group, g - guest of the group
```

Rate Limiter Type	Status	Packets/s	Burst	Sharing
MCAST NON RPF	Off	-	-	-
MCAST DFLT ADJ	On	100000	100	Not sharing
MCAST DIRECT CON	Off	-	-	-
ACL BRIDGED IN	Off	-	-	-
ACL BRIDGED OUT	On	1000	255	Group:1 S
IP FEATURES	Off	-	-	-
ACL VACL LOG	On	2000	1	Not sharing
CEF RECEIVE	Off	-	-	-
CEF GLEAN	Off	-	-	-
MCAST PARTIAL SC	On	100000	100	Not sharing
IP RPF FAILURE	On	100	10	Group:0 S
TTL FAILURE	Off	-	-	-
ICMP UNREAC. NO-ROUTE	On	100	10	Group:0 S
ICMP UNREAC. ACL-DROP	On	100	10	Group:0 S
ICMP REDIRECT	Off	-	-	-
MTU FAILURE	Off	-	-	-
MCAST IP OPTION	Off	-	-	-
UCAST IP OPTION	Off	-	-	-
LAYER_2 PDU	Off	-	-	-

```

LAYER_2 PT Off - - -
  IP ERRORS On 100 10 Group:0 S
CAPTURE PKT Off - - -
  MCAST IGMP Off - - -
MCAST IPv6 DIRECT CON Off - - -
MCAST IPv6 ROUTE CNTL Off - - -
MCAST IPv6 *G M BRIDG Off - - -
  MCAST IPv6 SG BRIDGE Off - - -
  MCAST IPv6 DFLT DROP Off - - -
MCAST IPv6 SECOND. DR Off - - -
  MCAST IPv6 *G BRIDGE Off - - -
    MCAST IPv6 MLD Off - - -
IP ADMIS. ON L2 PORT Off - - -
  UCAST IP TINY FRAG Off - - -
  MCAST IP TINY FRAG Off - - -

```

It is recommended that the network administration team investigate this rate limiting functionality. The switch defaults to enabling a number of rate limiters and predefines a set rate for that traffic. Based on the local traffic mix, the need to modify or enable additional rate limiters may be required. The following command shows the default settings on the Catalyst 6500:

```
6K-LV2-CL3# show mls rate-limit usage
```

```

                Rate Limiter Type      Packets/s  Burst
                -----
Layer3 Rate Limiters:
  RL# 0: Free          -                -        -
  RL# 1: Free          -                -        -
  RL# 2: Free          -                -        -
  RL# 3: Used
                    MCAST DFLT ADJ      100000    100
  RL# 4: Free          -                -        -
  RL# 5: Free          -                -        -
  RL# 6: Used
                    IP RPF FAILURE       100       10
                    ICMP UNREAC. NO-ROUTE 100       10
                    ICMP UNREAC. ACL-DROP 100       10

```

	IP ERRORS	100	10
RL# 7: Used			
	ACL VACL LOG	2000	1
RL# 8: Rsvd for capture	-	-	-

Layer2 Rate Limiters:

RL# 9: Reserved			
RL#10: Reserved			
RL#11: Free	-	-	-
RL#12: Free	-	-	-

3.5.2. Control Plane Policing (CoPP)

Control plane policing provides a broader level of protection than that offered by hardware-based rate limiters. In many respects, control plane policing should be viewed as a complement to the rate limiting function. The recommendation is to enable both CoPP and the appropriate hardware rate limiters.

Control plane policing introduces a new virtual interface on the switch: the control plane interface. This new interface type adds to existing virtual interfaces such as the loopback interface, the bridge virtual interface, the dialer interface, the tunnel interface, and the VLAN interface. When the control plane interface is enabled, it allows a rate limiter (policer) to be applied. This rate limiting action applies to the traffic selected by a user-configured ACL and destined to the control plane. This more flexible coverage can provide better protection against a global level control plane attack. Used in conjunction with the hardware rate limiters, CoPP provides additional protection against specific CPU-bound traffic storms (for example, with both a CoPP ACL using the fragment keyword and the IP option special rate limiters it is possible to create a protection for the route processor CPU from two uncommon types of traffic that can potentially be used for a DoS attack).

3.6 Mitigate Address Spoofing

Address spoofing is a means to gain access into a network. The aim of a spoofing attack is to fool the network into thinking the data is from a valid local source. Address spoofing is one method used to initiate a denial-of-service attack. One form of attack, known as an atomic attack, consists of a single packet attack. If a spoofing attack successfully launches an atomic attack and the packet gains entry into the network, then this attack can be successful in compromising the intended target.

Both Supervisor Engine 32 and Supervisor Engine 720 support the multipath unicast reverse path forwarding (uRPF) check in hardware. uRPF works by inspecting the forwarding table to determine if the packet has arrived from a network that is known to exist out of a given interface. If the arriving packet has a source address of A and arrived on interface X, yet the forwarding table shows network A to be associated to interface Y, then the system knows this is a spoofed address and can drop the packet. Supervisor Engine 720 and Supervisor Engine 32 have the added benefit of providing uRPF multipath lookups in hardware. When a routing protocol supports equal or unequal cost load balancing, a network entry can be installed in the forwarding table associated with multiple interfaces. Multipath uRPF can support a lookup to a forwarding entry associated with one or two interfaces without degrading system performance.

uRPF provides two forms of checks, strict mode and loose mode. Strict mode will only permit the packet to pass if the source address of the arriving packet can be found in the portion of the forwarding table associated to the interface from where the packet arrived. Loose mode check relaxes that check somewhat, and will only verify that the source address exists in the forwarding table irrespective of the interface it is associated with.

To better secure the network against spoofing attacks, it is recommended that uRPF strict mode be enabled.

Strict uRPF can be enabled on a per interface basis as follows:

```
6K-LV2-CL3(config-if)# ip verify unicast source reachable-via rx
```

Loose uRPF can also be enabled on a per interface basis as follows:

```
6K-LV2-CL3(config-if)# ip verify unicast source reachable-via any
```

Note that loose and strict uRPF differ in the last keyword of the configuration line. Other uRPF configuration options are available and these can be referenced in the Catalyst 6500 documentation on www.cisco.com.

3.7 Implement Port Security

Port security provides administrators with a tool to secure fixed host ports from being used by other devices. When implemented in the server farm, port security can be used to lock down ports to specific hosts using the MAC address as the key to keeping the port active. In this way, a switch port link will go into error disable status if someone, for example, attempts to use a server's Ethernet link to connect another device into the network.

It is recommended that switch ports used for all static devices such as servers, printers, scanners, access points, and so on be "locked down" by using port security. An example of how this feature is applied is shown below:

```
6K-LV2-CL3(config-if)# switchport port-security mac-address 0000.0C3F.2A6D
```

This command example locks down the given port by only allowing the device with the MAC address 0000.0C3F.2A6D to access that switch port.

Port security also supports another more flexible option by which it can limit the number of devices that can be seen on a given switch port. This becomes especially useful in protecting a switch from a MAC flooding attack. A MAC flooding attack works by sending random MAC addresses into the switch in an attempt to fill the L2 forwarding table (sometimes also referred to as "L2 CAM table"). The effect of this is that, if there is no space left to store the new MAC addresses of the incoming traffic, address learning will stop until some space is freed up in the table and the switch will have to flood all the traffic destined to the yet-to-be-learned MAC addresses. An unscrupulous user could then exploit this extra flooding and use a sniffer to collect all the data flooded in his/her VLAN.

At a minimum, port security should be implemented to limit the number of MAC addresses that can be learned on untrusted ports. This will limit the potential for a MAC flooding attack to be launched. The command used to enable this option of port security is shown below:

```
6K-LV2-CL3(config-if)# switchport port-security maximum 10 /*Modify value as needed!
```

The example limits the given port from learning more than 10 MAC addresses.

Please note that special care is necessary when configuring port security for dual-homed servers that are using NIC teaming for redundancy and load balancing: in these cases if one NIC loses connectivity the redundant NIC becomes active and/or inherits the same MAC address as the primary one, so on the same switch port more than one valid MAC address may be seen (including the virtual MAC address if configured).

Another security measure to be taken to minimize the scope of a MAC flooding attack is to disable the Dynamic Trunking Protocol (DTP) on all the ports connected to untrusted devices or to devices that do not require trunking. DTP is a Cisco Systems proprietary protocol that facilitates the bring-up process of trunk links between switches. For that reason its per-port configuration is by default set to auto (Catalyst OS) or to dynamic desirable (Cisco IOS Software). However, when the automatic bring up process of trunks is not required, the off option

(switchport mode access in Cisco IOS Software) offers a higher level of security because it prevents a rogue device from successfully completing a DTP message exchange, which could turn an access port into a trunk and allow the malicious device to gain access to all the VLANs active on that link.

3.8 Mitigate Man-in-the-Middle Attacks

A man-in-the-middle attack provides a way for a user to intercept traffic on an active session between two nodes in the network. This attack often makes use of gratuitous ARP packets, which allow a host to send messages telling other devices in the local subnet about a bogus IP to MAC address association (thereby “poisoning” the other devices’ ARP tables with this invalid information).

The attack works as follows. There are two well-behaved devices that intend to talk to each other, for example, host X and host Y, and there is a third malicious host Z that intends to launch a man-in-the-middle attack. So Z will first send a gratuitous ARP message to host X saying that it “owns” the MAC address of host Y. Then it will also send a gratuitous ARP message to host Y saying it “owns” the MAC address of host X. After doing that, host Z will have successfully poisoned the ARP cache of both host X and host Y. Host X will now send all the traffic destined to the IP address of host Y to host Z, and host Y, likewise, will send all traffic destined to the IP address of host X to host Z. Host Z would typically run some form of local forwarding to knit the two sessions together so that both X and Y will be unaware of the presence of Z as extra forwarding hop between them. In this manner, host X and host Y are none the wiser to the man in the middle, and host Z is able to inspect all the data exchanged between these two devices.

Cisco Systems introduced a new feature called Dynamic ARP Inspection (DAI) in its Catalyst OS Release 8.3 software, which can be used to mitigate this type of attack. DAI is used to create a binding table for each host that uses DHCP to request its IP address. DAI will use DHCP Snooping to snoop on the dynamically assigned IP address and bind that to the MAC address of the device on the corresponding switch port. If the user initiates an attack using bogus ARP information, the switch will inspect the ARP messages against the entry for that port in its binding table. If it finds the advertised MAC/IP association does not match the binding entry, it will drop the ARP packet. DAI is now available starting from Release 12.2(18)SXE of Cisco IOS Software.

For those devices that do not use DHCP (such as servers and other static hosts like printers and so on), a static ARP ACL can be used to filter out the invalid MAC/IP associations.

It is recommended that customers consider implementing this feature to harden their network against these forms of attacks.

3.9 Isolate VLAN 1 from User Traffic

Historically the Catalyst 6500 has used VLAN 1 as a special management VLAN. This VLAN can also be used as native (that is, untagged) VLAN on 802.1Q trunks and serves as the default VLAN for access ports. It is used by various protocols like STP, Cisco Discovery Protocol, VTP, and PAGP. One issue is that because of its nature VLAN 1 can end up spanning the entire network. This can lead to the network becoming somewhat unstable if STP loops are formed. Of more concern is that, as an untagged VLAN, VLAN 1 has the potential of being misused to initiate a VLAN hopping attack. With this attack, the user can use VLAN 1 to hop into another VLAN bypassing the Layer 3 device that normally polices the movement of traffic between VLANs.

For this reason, the recommendation is to avoid using VLAN 1 for any user traffic and as default VLAN of any port. Moreover it is recommended to use another numbered VLAN to carry the network management traffic (telnet, SSH, SNMP, and so on).

A generic recommendation is to remove VLAN 1 from all ports and trunks that do not need it. Reducing the spread of VLAN 1 will minimize the opportunity of it being compromised.

3.10 Protect the STP Root Switch

The implementation of the Spanning Tree Protocol (STP) necessitates the definition of the so-called spanning tree root. The STP root plays an important role in a spanning tree domain. It is the peak of a tree-based topology that defines a loop-free Layer 2 network over which data is forwarded. The STP root is derived as the switch having the lowest bridge priority among all switches in the same STP domain.

Should there be a tie (that is, two or more devices have the same bridge priority), then the device with the lowest-order MAC address assumes the STP root role.

Should another unauthorized device take over the STP root role, the STP topology could change, creating an inefficient path for data to traverse. Even worse is the consideration that someone could insert a device to assume the STP root role to divert traffic over a path that could be used to snoop on data.

It is recommended that, when the correct STP root device is configured, the STP Root Guard feature be enabled to protect the integrity of the assigned STP root device. Should the STP root receive a superior BPDU (indicating another device wants to take over the STP role), then the active STP root will place that port into a root inconsistent state and disable the port. This protects the STP root role from being accidentally or intentionally taken away.

It is recommended that the STP Root Guard feature be enabled on all appropriate interfaces on the designated root switch using the following command:

```
6K-LV2-CL3(config-if)# spanning-tree guard root
```

Please note that in the vast majority of the cases this type of feature is more helpful in preventing possible misconfigurations (especially with switches placed in the network without being properly set up, for example, blade server switches using a default configuration) rather than blocking STP attacks, which are known to be very rare.

3.11 Improve Link Availability

When a host activates a switch port, by default the switch will go through normal spanning tree listening and learning states to help ensure a loop is not going to be created with the activation of this port. This can manifest into a problem if, when the host is booting, it reaches the point of issuing a DHCP request before the switch port is active. This can leave the host without an IP address and isolate it from the network.

Spanning Tree Portfast is designed to avoid this happening. When enabled, it bypasses the STP listening and learning states allowing the port to move to the forwarding state when it is activated. This feature assumes that an end-station is connected to the Portfast port. If another STP-capable switch instead inadvertently uses this port, it could potentially create a loop in the network. Loops would have adverse effects on the stability of the network. STP BPDU Guard is a further enhancement to the Portfast feature in that it shuts the port down

if a BPDU (Bridge Protocol Data Unit) is detected. A BPDU is the vehicle a switch uses to exchange spanning tree information with other switches.

STP Portfast is generally advocated for all connected hosts. It is also recommended that use of Portfast should be complemented with STP BPDU Guard, when applicable, to enjoy the benefits of both. These can be enabled as follows:

```
6K-LV2-CL3(config-if)# spanning-tree portfast
```

```
6K-LV2-CL3(config-if)# spanning-tree bpduguard
```

Please note that, similarly to RootGuard, BPDU Guard is very helpful in preventing not-so-uncommon misconfigurations, and as an added benefit it can also protect from rare STP attacks.

3.12 Secure Servers with Private VLANs

Private VLANs extend the capabilities of normal VLANs by providing a framework for defining special subdomains with a constrained forwarding capability. Private VLANs introduce the concept of two “secondary” VLAN types known as a community VLAN and an isolated VLAN. These two entities exist always only in association with what is referred to as a primary VLAN. A promiscuous port is the element in a private VLAN domain that acts for example as a Layer 3 port to provide a path through which hosts in a community or in an isolated VLAN can talk to the rest of the world (ACLs permitting).

Many community VLANs along with a single isolated VLAN can be associated to the same primary VLAN. In this environment, the L2 forwarding rules of engagement dictate that hosts within a community VLAN can communicate with one another but cannot communicate with hosts in another community VLAN. Isolated VLANs are more restrictive and do not allow hosts in the same isolated VLAN to communicate at Layer 2 with each other (unless of course L3 communication can be established among them using the promiscuous port).

Private VLANs offer customers an added security feature. Should a host, server, or group of devices have a requirement to limit communications between one another, private VLANs can be used to support this capability. The simplicity of private VLANs provides a way to easily secure communication between devices in the same IP subnet. In addition, by default, private VLANs enable the so-called sticky ARP capability on the associated L3 VLAN interfaces. This measure makes potential poisoning attempts of the default gateway’s ARP table much harder, while the intrinsically higher Layer 2 segregation that private VLANs provide offers a better protection against host-to-host ARP poisoning attempts.

Please notice though that the sticky ARP functionality makes IP/MAC bindings in the ARP table rather static, which can create problems with dual-attached servers using the burned-in MAC addresses for redundancy purposes. In those cases it is recommended to configure the redundant network interface cards (NICs) to use a virtual MAC address instead.

4. BEST PRACTICES: AVAILABILITY

An increased reliance on network services means an increased need for higher availability. Most networks implement some of the more obvious physical high-availability features such as redundant hardware, redundant cabling plant, separate power supplies, and so on. Software availability is also increased with the implementation of features like Nonstop Forwarding (NSF) and Stateful Switchover (SSO).

To further improve availability of the network infrastructure, a number of not so obvious features can be used to further bolster availability of online services. Improving availability is not only about maintaining the operational running of a system when a fault occurs; it is also about getting the system back online in a timely manner when a component fails. Best practice availability recommendations are detailed below.

4.1 Improve System Health Checks

Generic On-Line Diagnostics (GOLD) is a recent enhancement available on the Catalyst 6500 with the introduction of Release 12.2(18)SXD of Cisco IOS Software. GOLD provides a way to inspect the health of numerous components that make up the switch. The diagnostic aspect of this facility provides checks at system startup time as well as checks during the operational running of the switch.

When configured for boot-up diagnostics, GOLD will check the health of components in the chassis. As an example, if GOLD detects a malfunctioning line card, it will shut it down before the system becomes active.

Diagnostics in this manner can be enabled as follows:

```
6K-LV2-CL3(config)# diagnostic bootup level ?
  complete  Complete level
  minimal   Minimal level
```

Running the minimal level test initiates a set of checks on the supervisor's policy feature card (PFC) and a set of loopback tests on all ports. The complete level test runs all available checks.

Diagnostics can also be performed at runtime. Health diagnostics can be run in the background on a continued basis to scan the system for any faults that might arise during the operation of the switch. On-demand based diagnostics can be initiated by the administrator to help troubleshoot possible hardware faults. There is also an option to schedule diagnostics to be run at a specific time or on a periodic basis.

It is recommended that background checks be enabled on the Catalyst 6500 to proactively and promptly identify problems if they arise.

More details on GOLD can be found in a white paper located on Cisco.com at

http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps708/prod_white_paper0900aecd801e659f.shtml

4.2 Minimize Downtime from Fiber Patching Errors or STP Loops

Where a network has an extensive fiber cabling plant, problems have arisen whereby incorrect patching can lead to a fiber switch port having its RX fiber strand connected to one switch port and its TX fiber strand connected to another. In this manner, the port may be deemed active if it is receiving an optical signal on its RX strand. However, the problem is that this link does not provide a valid communications path at Layer 2 (and above).

Unidirectional Link Detection (UDLD) is a Cisco designed solution to solve this problem. UDLD is a Layer 2 protocol that exchanges periodic hellos to establish the neighbor's identity. It also confirms that traffic is flowing in a bidirectional manner between the correct neighbors. In the problem scenario described above, UDLD would detect that the fiber link is not terminated at the other end by the same neighbor if used in conjunction with L1 detection mechanisms such as Ethernet autonegotiation and Far End Fault Indication (FEFI). This problem would result in the erroneous patch port being shut down. In addition, UDLD has also been very popular as a means to detect cases of malfunctioning links that become unidirectional and cause loss of BPDUs, thus initiating dangerous STP loops.

It is recommended that UDLD be enabled on all ports terminating links where a potential miswiring can take place or on ports where STP is running. UDLD needs to be enabled globally as well as on a per interface basis. The commands to enable UDLD globally and on an interface are as follows:

```
6K-LV2-CL3(config)# udld enable
6K-LV2-CL3(config)# interface g1/2
6K-LV2-CL3(config-if)# udld port
```

UDLD supports two modes of operation: normal and aggressive. The recommended mode to use in the vast majority of the network designs is the former, whereas the latter should be used with great care in those cases where a loss of connectivity with a neighbor represents a serious connectivity problem (for example, when a bundled port in an EtherChannel[®] connection loses connectivity with its peer and EtherChannel is not running a protocol such as PAgP or LACP).

While UDLD can be used to detect "directionality" problems on any Ethernet link, another feature should be used in synergy with UDLD to guarantee greater STP robustness: the feature in question is called STP Loopguard. The main advantage of Loopguard is that it is part of the STP algorithm and so it proactively acts and reacts to STP events at the same speed of the STP protocol. Its main disadvantage is that it runs on logical links, as STP does, whereas other features like UDLD and PAgP/LACP run on the individual physical links and so have a lower level type of visibility into connectivity problems. As a generic principle, feature combinations can offer a better protection than the individual features alone, and it is therefore recommended to fully exploit feature synergies when possible.

Loopguard can be enabled either globally or on a per-port basis thus:

```
6K-LV2-CL3(config)# spanning-tree loopguard default
```

```
6K-LV2-CL3(config-if)# spanning-tree guard loop
```

4.3 Speed Resolution to Cable Management Problems

Cable faults can be problematic and hard to resolve at the best of times. Cisco Systems has introduced a range of Catalyst 6500 line cards that support Time Domain Reflectometry (TDR), which is designed to assist in the resolution of copper based cabling problems. TDR is provided on the following copper based Ethernet line cards:

- WS-X6748-GE-TX (48 port CEF720 10/100/1000 line card)
- WS-X6548-GE-TX/45AF (48 port CEF256 10/100/1000 line card)
- WS-X6148A-GE-TX/45AF (48 port Classic 10/100/1000 line card)
- WS-X6148-GE-TX/45AF (48 port Classic 10/100/1000 line card)
- WS-X6148A-RJ-45/45AF (48 port Classic 10/100 line card)

TDR can be used to assist in diagnosing a range of major and minor cable issues such as loose connectors, cable cuts (for example, broken cable), incorrect cabling pin-outs, shorted cable (for example, cables touching each other through damaged insulation), shortened conductors, and more. TDR is run from the command line for a given port. An example of the output is shown below.

```
6K-LV2-CL3# test cable-diagnostics tdr interface g1/3
```

```
TDR test started on interface Gi1/3
```

```
A TDR test can take a few seconds to run on an interface
```

```
Use 'show cable-diagnostics tdr' to read the TDR results.
```

```
6K-LV2-CL3# show cable-diagnostics tdr interface g1/3
```

```
TDR test last run on: March 5 10:22:06
```

Interface	Speed	Pair	Cable length	Distance to fault	Channel	Pair status
Gi1/3	1000	1-2	1 +/- 6 m	N/A	Pair B	Terminated
		3-4	1 +/- 6 m	N/A	Pair A	Terminated
		5-6	1 +/- 6 m	N/A	Pair C	Terminated
		7-8	1 +/- 2 m	N/A	Pair D	Short

The output from this command shows pair 7-8 as having been shorted. It is recommended that customers consider using the TDR-enabled 10/100 or 10/100/1000 line cards for host and server connectivity (particularly the 6748 modules, which were designed for high performance environments). This can improve resolution times for cable problems and also provide a verification tool for correctly installed cabling.

5. BEST PRACTICES: PERFORMANCE

Performance is one of the key requirements for most network switches to help ensure that attached clients can obtain a good service level from the network. The Catalyst 6500 is architected to support a number of processing elements in hardware, which can maximize performance. Features like Cisco Express Forwarding are the cornerstone to switching packets in hardware. Additional processing elements have also been integrated into the hardware and include but are not limited to policing, Generic Route Encapsulation (GRE), IPv6, Multiprotocol Label Switching (MPLS), Network Address Translation (NAT), and more.

The following provides a list of some best practices to further enhance the overall system performance of the Catalyst 6500.

5.1 Improving Server Performance

For many customer networks, server performance is critical in meeting service level agreements. There are a multitude of tweaks that can be applied to the server hardware to yield improved performance levels. From a network perspective, one major feature that can be enabled on the Catalyst 6500 to improve server performance is jumbo frames.

Enabling the jumbo frames feature allows the transmission of large packets between the switch and the servers. The maximum standard Ethernet frame size is 1518 bytes (1522 bytes with an 802.1Q header). This consists of a 1500-byte data payload (whose length is the standard Maximum Transmission Unit [MTU]) and an 18-byte Ethernet header + trailer. The 802.1Q header adds 4 extra bytes of header information. Considerable performance advantage can be gained from increasing the amount of data transferred in the payload of each frame. The main benefit that can be derived is a reduction in the number of interrupts needed on the end stations to complete a data transfer or a large transaction. The Catalyst 6500 supports jumbo frames on its range of Ethernet line cards (with the notable exception of the 6148-GE-TX/45AF and the 6548-GE-TX/45AF modules). Jumbo frames usually range in size from 8192 bytes to 9216 bytes (this is dependent on the specific generation of line card in use). In parallel with enabling jumbo frames on the switch, the server must also be primed with a suitable NIC that also supports jumbo frames.

When jumbo frames are enabled on the Catalyst 6500, an MTU range from 1500 bytes to 9216 bytes can be chosen. The system's default MTU is 9216 bytes, but this can be changed with the following command:

```
6K-LV2-CL3(config)# system jumbomtu ?
<1500-9216> Jumbo mtu size in Bytes, default is 9216
```

With the command above, a global MTU size different from the default can be set to dictate the maximum size that a given interface on the switch can use. Following this, the MTU of a specific interface can be reset as shown in the following example:

```
6K-LV2-CL3(config-if)# mtu ?
<1500-9216> MTU size in bytes
```

It is recommended that consideration be given to enabling jumbo frames on appropriate servers in the network to gain a performance advantage.

5.2 Eliminate Excess Background Traffic

Networks traffic falls in three categories: broadcast, multicast, and unicast traffic. While broadcast traffic is a necessary component of many protocols that run on today's networks, broadcast data does generate CPU overhead on each network node that receives and processes broadcast packets.

The Catalyst 6500 implements the storm control feature (also known as broadcast suppression) on selected line cards and is designed to allow the specification of a limit of how much broadcast (and multicast) traffic can be transmitted. This provides some control to the administration team as to how much broadcast traffic they deem acceptable for a given node to process. Storm control works by monitoring the amount of storm traffic (broadcast, multicast, and so on) that is sent every one-second interval on a set interface. It allows the administrator to specify how much of that storm traffic can be sent as a percentage of the total bandwidth of that interface. An example of a configuration is shown below:

```
6K-LV2-CL3(config-if)# storm-control broadcast level 70
```

The example above states that, once the broadcast traffic load goes beyond 70% of the available interface bandwidth within a one-second time window, the subsequent broadcast packets for the remainder of that one-second interval will be dropped.

Storm control is implemented in hardware and has no effect on the overall performance of the switch.

It is recommended that storm control be enabled on key nodes in the network to avoid them from having to process excess traffic.

5.3 Reduce Overhead of Multicast Traffic

Multicast traffic on networks has become more common since the introduction of numerous video-streaming technologies from the likes of Real Audio, Apple, and Microsoft. Like broadcast traffic, multicast traffic is seen as a broadcast packet from a Layer 2 switching perspective. The result on attached hosts is the same as that seen when it processes a broadcast packet. Multicast traffic forwarded to hosts that do not request it still causes a CPU interrupt on the end host.

The Catalyst 6500 supports hardware assist for the IGMP Snooping feature, which works by inspecting (snooping) IGMP packets that pass through the switch. By snooping IGMP packets, the switch can learn which hosts wish to receive multicast streams. This information is loaded into the hardware forwarding table(s) and enables the switch to forward multicast data only to those hosts requesting it. The result is to minimize the unnecessary propagation of multicast traffic through a Layer 2 switch domain.

IGMP Snooping can be enabled on the switch as follows:

```
6K-LV2-CL3(config)# ip igmp snooping
```

It is recommended that IGMP snooping be enabled on the switch to benefit from the improvements that can be achieved with the reduction of the propagation of the multicast traffic.

Reductions in multicast forwarding can be further enhanced with the implementation of PIM Snooping. PIM Snooping works at Layer 2 and snoops on PIM control messages between PIM devices. By default, a PIM router will flood IP multicast packets to a downstream PIM router even if that PIM router has no multicast receivers downstream of it. This can result in an unnecessary waste of bandwidth between PIM-enabled devices. PIM Snooping determines if there are downstream receivers and, if none is found, it will not unnecessarily flood the link with the multicast traffic.

PIM Snooping can be enabled as follows:

```
6K-LV2-CL3(config)# ip pim snooping
```

Like with IGMP Snooping, the recommendation is to enable PIM snooping if multicast pruning is necessary within a certain Layer 2 domain.

6. SUMMARY

Customers investing in network infrastructure can gain a significant edge through the deployment of the Catalyst 6500 technology. Beyond the major marketed features of the Catalyst 6500 lies a plethora of less well-known yet important functionalities that when implemented can yield significant benefits even for the smallest of networks.

The following list summarizes the best practices highlighted in this paper.

Category	Recommendation	Gold	Silver	Bronze
Manageability	Configure Hostname	√		
	Configure Login Banner		√	
	Set System Clock		√	
	Use NTP to set system Clock	√		
	Set up base SNMP Parameters		√	
	Backup Cisco IOS Software and Configuration files	√		
	Use SPAN and RSPAN	√		
	Implement NDE for Traffic Analysis	√		
	Implement Power Redundancy		√	
	Understand Power Shutdown Sequence			√
Security	Disable Cisco Discovery Protocol	√		
	Protect against Common Attacks (SYN, etc.)	√		
	Secure Switch CLI	√		
	Implement SSH for remote access	√		
	Implement Control Plane Protection	√		
	Protect against Address Spoofing using uRPF	√		
	Implement Port Security	√		
	Implement Dynamic ARP Inspection	√		
	Remove VLAN 1 as management VLAN	√		
	Implement STP Root Guard			√
	Implement BPDU Guard	√		
	Secure with Private VLANs		√	
Availability	Implement GOLD	√		
	Implement UDLD and Loopguard	√		
	Implement TDR			√
Performance	Use Jumbo Frames			√
	Implement Broadcast Suppression		√	
	Implement IGMP Snooping		√	
	Implement PIM Snooping			√



Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the **Cisco.com Website at www.cisco.com/go/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus • Czech Republic
Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel • Italy
Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal
Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2006 Cisco Systems, Inc. All rights reserved. CCSP, CCVP, the Cisco Square Bridge logo, Follow Me Browsing, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, Packet, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0601R)

Printed in USA

C11-366433-00 09/06