Network Management System: Best Practices White Paper

Contents

Introduction

The International Organization for Standardization (ISO) network management model defines five functional areas of network management. This document covers all functional areas. The overall purpose of this document is to provide practical recommendations on each functional area to increase the overall effectiveness of current management tools and practices. It also provides design guidelines for future implementation of network management tools and technologies.

Network Management

The ISO network management model's five functional areas are listed below.

- Fault Management—Detect, isolate, notify, and correct faults encountered in the network.
- Configuration Management—Configuration aspects of network devices such as configuration file management, inventory management, and software management.

- Performance Management—Monitor and measure various aspects of performance so that overall performance can be maintained at an acceptable level.
- Security Management—Provide access to network devices and corporate resources to authorized individuals.
- Accounting Management—Usage information of network resources.

The following diagram shows a reference architecture that Cisco Systems believes should be the minimal solution for managing a data network. This architecture includes a Cisco CallManager server for those who plan to manage Voice over Internet Protocol (VoIP): The diagram shows how you would integrate the CallManager server into the NMS topology.

The network management architecture includes the following:

- Simple Network Management Protocol (SNMP) platform for fault management
- Performance monitoring platform for long term performance management and trending
- CiscoWorks2000 server for configuration management, syslog collection, and hardware and software inventory management

Some SNMP platforms can directly share data with the CiscoWorks2000 server using Common Information Model/eXtensible Markup Language (CIM/XML) methods. CIM is a common data model of an implementation-neutral schema for describing overall management information in a network/enterprise environment. CIM is comprised of a specification and a schema. The specification defines the details for integration with other management models such as SNMP MIBs or Desktop Management Task Force Management Information Files (DMTF MIFs), while the schema provides the actual model descriptions.

XML is a markup language used for representing structured data in textual form. A specific goal of XML was to keep most of the descriptive power of SGML whilst removing as much of the complexity as possible. XML is similar in concept to HTML, but whereas HTML is used to convey graphical information about a document, XML is used to represent structured data in a document.

Cisco's advanced services customers would also include Cisco's NATkit server for additional proactive monitoring and troubleshooting. The NATkit server would either have a remote disk mount (rmount) or file transfer protocol (FTP) access to the data residing on the CiscoWorks2000 server.

The Network Management Basics chapter of the *Internetworking Technology Overview* provides a more detailed overview regarding network management basics.

Fault Management

The goal of fault management is to detect, log, notify users of, and (to the extent possible) automatically fix network problems to keep the network running effectively. Because faults can cause downtime or unacceptable network degradation, fault management is perhaps the most widely implemented of the ISO network management elements.

Network Management Platforms

A network management platform deployed in the enterprise manages an infrastructure that consists of multivendor network elements. The platform receives and processes events from network elements in the network. Events from servers and other critical resources can also be forwarded to a management platform. The following commonly available functions are included in

a standard management platform:

- Network discovery
- Topology mapping of network elements
- Event handler
- Performance data collector and grapher
- Management data browser

Network management platforms can be viewed as the main console for network operations in detecting faults in the infrastructure. The ability to detect problems quickly in any network is critical. Network operations personnel can rely on a graphical network map to display the operational states of critical network elements such as routers and switches.

Network management platforms such HP OpenView, Computer Associates Unicenter, and SUN Solstice can perform a discovery of network devices. Each network device is represented by a graphical element on the management platform's console. Different colors on the graphical elements represent the current operational status of network devices. Network devices can be configured to send notifications, called SNMP traps, to network management platforms. Upon receiving the notifications, the graphical element representing the network device changes to a different color depending on the severity of the notification received. The notification, usually called an event, is placed in a log file. It is particularly important that the most current Cisco Management Information Base (MIB) files be loaded on the SNMP platform to ensure that the various alerts from Cisco devices are interpreted correctly.

Cisco publishes the MIB files for managing various network devices. The Cisco MIB files are located on the cisco.com website, and include the following information:

- MIB files published in SNMPv1 format
- MIB files published in SNMPv2 format
- Supported SNMP traps on Cisco devices
- OIDs for Cisco current SNMP MIB objects

A number of network management platforms are capable of managing multiple geographically distributed sites. This is accomplished by exchanging management data between management consoles at remote sites with a management station at the main site. The main advantage of a distributed architecture is that it reduces management traffic, thus, providing a more effective usage of bandwidth. A distributed architecture also allows personnel to locally manage their networks from remote sites with systems.

A recent enhancement to management platforms is the ability to remotely management network elements using a web interface. This enhancement eliminates the need for special client software on individual user stations to access a management platform.

A typical enterprise is comprised of different network elements. However, each device normally requires vendor-specific element management systems in order to effectively manage the network elements. Therefore, duplicate management stations may be polling network elements for the same information. The data collected by different systems is stored in separate databases, creating administration overhead for users. This limitation has prompted networking and software vendors to adopt standards such as Common Object Request Broker Architecture (CORBA) and Computer-Integrated Manufacturing (CIM) to facilitate the exchange of management data between management platforms and element management systems. With vendors adopting standards in management system development, users can expect interoperability and cost savings in deploying and managing the infrastructure.

CORBA specifies a system that provides interoperability between objects in a heterogeneous, distributed environment and in a manner that is transparent to the programmer. Its design is based on the Object Management Group (OMG) object model.

Troubleshooting Infrastructure

Trivial File Transfer Protocol (TFTP) and system log (syslog) servers are crucial components of a troubleshooting infrastructure in network operations. The TFTP server is used primarily for storing configuration files and software images for network devices. Routers and switches are capable of sending system log messages to a syslog server. The messages facilitate the troubleshooting function when problems are encountered. Occasionally, Cisco support personnel need the syslog messages to perform root cause analysis.

The CiscoWorks2000 Resource Management Essentials (Essentials) distributed syslog collection function allows for the deployment of several UNIX or NT collection stations at remote sites to perform message collection and filtering. The filters can specify which syslog messages will be forwarded to the main Essentials server. A major benefit of implementing distributed collection is the reduction of messages forwarded to the main syslog servers.

Fault Detection and Notification

The purpose of fault management is to detect, isolate, notify, and correct faults encountered in the network. Network devices are capable of alerting management stations when a fault occurs on the systems. An effective fault management system consists of several subsystems. Fault detection is accomplished when the devices send SNMP trap messages, SNMP polling, remote monitoring (RMON) thresholds, and syslog messages. A management system alerts the end user when a fault is reported and corrective actions can be taken.

Traps should be enabled consistently on network devices. Additional traps are supported with new Cisco IOS software releases for routers and switches. It is important to check and update the configuration file to ensure the proper decoding of traps. A periodic review of configured traps with the Cisco Assured Network Services (ANS) team will ensure effective fault detection in the network.

The following table lists the CISCO-STACK-MIB traps that are supported by, and can be used to monitor fault conditions on, Cisco Catalyst local area network (LAN) switches.

Trap	Description
moduleUp	The agent entity has detected that the **moduleStatus** object in this MIB has transitioned to the **ok(2)** state for one of its modules.
moduleDown	The agent entity has detected that the *moduleStatus* object in this MIB has transitioned out of the **ok(2)** state for one of its modules.
chassisAlarmOn	The agent entity has detected that the *chassisTempAlarm*, *chassisMinorAlarm*, or *chassisMajorAlarm* object in this MIB has transitioned to the **on(2)** state. A *chassisMajorAlarm* indicates that one of the following conditions exists:

	• Any voltage failure • Simultaneous temperature and fan failure • One hundred percent power supply failure (two out of two, or one out of one) • Electrically erasable programmable read-only memory (EEPROM) failure • Nonvolatile RAM (NVRAM) failure • MCP communication failure • NMP status unknown A chassisMinorAlarm indicates that one of the following conditions exists: • Temperature alarm • Fan failure • Partial power supply failure (one out of two) • Two power supplies of incompatible type
chassisAlarmOff	The agent entity has detected that the *chassisTempAlarm*, *chassisMinorAlarm*, or *chassisMajorAlarm* object in this MIB has transitioned to the **off(1)** state.

Environmental monitor (envmon) traps are defined in CISCO-ENVMON-MIB trap. The envmon trap sends Cisco enterprise-specific environmental monitor notifications when an environmental threshold is exceeded. When envmon is used, a specific environmental trap type can be enabled, or all trap types from the environmental monitor system can be accepted. If no option is specified, all environmental types are enabled. It can be one or more of the following values:

- voltage—A ciscoEnvMonVoltageNotification is sent if the voltage measured at a given test point is outside the normal range for the test point (such as is at the warning, critical, or shutdown stage).
- shutdown—A ciscoEnvMonShutdownNotification is sent if the environmental monitor detects that a test point is reaching a critical state and is about to initiate a shutdown.
- supply—A ciscoEnvMonRedundantSupplyNotification is sent if the redundant power supply (where extant) fails.
- fan—A ciscoEnvMonFanNotification is sent if any one of the fans in the fan array (where extant) fails.
- temperature—A ciscoEnvMonTemperatureNotification is sent if the temperature measured at a given test point is outside the normal range for the test point (such as is at the warning, critical, or shutdown stage).

Fault detection and monitoring of network elements can be expanded from the device level to the protocol and interface levels. For a network environment, fault monitoring can include Virtual Local Area Network (VLAN), asynchronous transfer mode (ATM), fault indications on physical interfaces, and so forth. Protocol-level fault management implementation is available using an element management system such as the CiscoWorks2000 Campus Manager. The TrafficDirector application in Campus Manager focuses on switch management utilizing mini-RMON support on Catalyst switches.

With an increasing number of network elements and complexity of network issues, an event management system that is capable of correlating different network events (syslog, trap, log files) may be considered. This architecture behind an event management system is comparable to a

Manager of Managers (MOM) system. A well-designed event management system allows personnel in the network operations center (NOC) to be proactive and effective in detecting and diagnosing network issues. Event prioritization and suppression allow network operation personnel to focus on critical network events, investigate several event management systems including the Cisco Info Center, and conduct a feasibility analysis to fully explore the capabilities of such systems. To obtain more information, go to the Cisco Info Center.

Proactive Fault Monitoring and Notification

RMON alarm and event are two groups defined in the RMON specification. Normally, a management station performs polling on network devices to determine the status or value of certain variables. For example, a management station polls a router to find out the central processing unit (CPU) utilization and generate an event when the value hits reaches a configured threshold. This method wastes network bandwidth and can also miss the actual threshold depending on the polling interval.

With RMON alarm and events, a network device is configured to monitor itself for rising and falling thresholds. At a predefined time interval, the network device will takes a sample of a variable and compares it against the thresholds. An SNMP trap can be sent to a management station if the actual value exceeds or falls below the configured thresholds. RMON alarm and event groups provide a proactive method of managing critical network devices.

Cisco Systems recommends implementing RMON alarm and event on critical network devices. Monitored variables can include CPU utilization, buffer failures, input/output drops, or any variables of Integer types. Starting with Cisco IOS Software Release 11.1(1), all router images support RMON alarm and event groups.

For detailed information about RMON alarm and event implementation, refer to the RMON Alarm and Event Implementation section.

RMON Memory Constraints

RMON memory usage is constant across all switch platforms relating to statistics, histories, alarms, and events. RMON uses what is called a *bucket* to store histories and statistics on the RMON agent (which is the switch in this case). The bucket size is defined on the RMON probe (SwitchProbe device) or RMON application (TrafficDirector tool), then sent to the switch to be set.

Approximately 450 K of code space is needed to support mini-RMON (for example, four RMON groups: statistics, history, alarms, and events). The dynamic memory requirement for RMON varies because it depends on the runtime configuration.

The following table defines the runtime RMON memory usage information for each mini-RMON group.

RMON Group Definition	**DRAM Space Used**	**Notes**
Statistics	140 bytes per switched Ethernet/Fast Ethernet port	Per port
History	3.6 K for 50 buckets *	Each additional

		bucket uses 56 bytes
Alarm and Event	2.6 K per alarm and its corresponding event entries	Per alarm per port

*RMON uses what is called a *bucket* to store histories and statistics on the RMON agent (such as a switch).

RMON Alarm and Event Implementation

By incorporating RMON as part of a fault management solution, a user can proactively monitor the network before a potential problem occurs. For example, if the number of broadcast packets received increases significantly, it can cause an increase in CPU utilization. By implementing RMON alarm and event, a user can set up a threshold to monitor the number of broadcast packets received and alert the SNMP platform by means of an SNMP trap if the configured threshold is reached. RMON alarms and events eliminate the excessive polling normally performed by the SNMP platform for accomplishing the same goal.

Two methods are available from which to configure RMON alarm and event:

- Command-line interface (CLI)
- SNMP SET

The following sample procedures show how to set a threshold to monitor the number of broadcast packets received on an interface. The same counter is used in these procedures as is shown in the **show interface** command example at the end of this section.

Command-line Interface Example

To implement RMON alarm and event using the CLI interface, perform the following steps:

1. Find the interface index associated with Ethernet 0 by walking the ifTable MIB.
```
interfaces.ifTable.ifEntry.ifDescr.1 = "Ethernet0"
interfaces.ifTable.ifEntry.ifDescr.2 = "Ethernet1"
interfaces.ifTable.ifEntry.ifDescr.3 = "FastEthernet0"
interfaces.ifTable.ifEntry.ifDescr.4 = "Fddi0"
```
2. Obtain the OID associated with the CLI field to be monitored. For this example, the OID for 'broadcasts' is 1.3.6.1.2.1.2.2.1.12. The Cisco OIDs for specific MIB variables are available from the cisco.com website.
3. Determine the following parameters for setting up thresholds and events.rising and falling thresholdssampling type (absolute or delta)sampling intervalaction when threshold is reachedFor the purpose of this example, a threshold is being set up to monitor the number of broadcast packets received on Ethernet 0. A trap will be generated if the number of broadcast packets received is greater than 500 between 60-second samples. The threshold will be reactivated when the number of input broadcasts does not increase between samples taken.**Note:** For detailed about these command parameters, check the Cisco Connection Online (CCO) documentation for RMON alarm and event commands for your particular Cisco IOS version.
4. Specify the trap sent (RMON event) when the threshold is reached using the following CLI commands (The Cisco IOS commands are displayed in bold):**rmon event 1 trap gateway**

description "High Broadcast on Ethernet 0" owner ciscormon event 2 log description "normal broadcast received on ethernet 0" owner cisco

5. Specify the thresholds and relevant parameters (RMON alarm) using the following CLI commands:**rmon alarm 1 ifEntry.12.1 60 delta rising-threshold 500 1falling-threshold 0 2 owner cisco**
6. Use SNMP to poll these tables to verify that the eventTable entries were made on the device.

```
rmon.event.eventTable.eventEntry.eventIndex.1 = 1

rmon.event.eventTable.eventEntry.eventIndex.2 = 2

rmon.event.eventTable.eventEntry.eventDescription.1 =
"High Broadcast on Ethernet 0"

rmon.event.eventTable.eventEntry.eventDescription.2 =
"normal broadcast received on ethernet 0"

rmon.event.eventTable.eventEntry.eventType.1 = snmp-trap(3)

rmon.event.eventTable.eventEntry.eventType.2 = log(2)

rmon.event.eventTable.eventEntry.eventCommunity.1 = "gateway"

rmon.event.eventTable.eventEntry.eventCommunity.2 = ""

rmon.event.eventTable.eventEntry.eventLastTimeSent.1 =
Timeticks: (0) 0:00:00

rmon.event.eventTable.eventEntry.eventLastTimeSent.2 =
Timeticks: (0) 0:00:00

rmon.event.eventTable.eventEntry.eventOwner.1 = "cisco"

rmon.event.eventTable.eventEntry.eventOwner.2 = "cisco"

rmon.event.eventTable.eventEntry.eventStatus.1 = valid(1)

rmon.event.eventTable.eventEntry.eventStatus.2 = valid(1)
```

7. Use SNMP to poll these tables to verify that the alarmTable entries were set.

```
rmon.alarm.alarmTable.alarmEntry.alarmIndex.1 = 1

rmon.alarm.alarmTable.alarmEntry.alarmInterval.1 = 60

rmon.alarm.alarmTable.alarmEntry.alarmVariable.1 = OID:
interfaces.ifTable.ifEntry.ifInNUcastPkts.2

rmon.alarm.alarmTable.alarmEntry.alarmSampleType.1 = absoluteValue(1)

rmon.alarm.alarmTable.alarmEntry.alarmValue.1 = 170183

rmon.alarm.alarmTable.alarmEntry.alarmStartupAlarm.1 =
risingOrFallingAlarm(3)

rmon.alarm.alarmTable.alarmEntry.alarmRisingThreshold.1 = 500

rmon.alarm.alarmTable.alarmEntry.alarmFallingThreshold.1 = 0

rmon.alarm.alarmTable.alarmEntry.alarmRisingEventIndex.1 = 1

rmon.alarm.alarmTable.alarmEntry.alarmFallingEventIndex.1 = 2

rmon.alarm.alarmTable.alarmEntry.alarmOwner.1 = "cisco"
```

```
rmon.alarm.alarmTable.alarmEntry.alarmStatus.1 = valid(1)
```

SNMP SET Example

In order to implement RMON alarm and event with the SNMP SET operation, complete these steps:

1. Specify the trap sent (RMON event) when the threshold is reached using the following SNMP SET operations:

```
# snmpset -c private 172.16.97.132 1.3.6.1.2.1.16.9.1.1.2.1
  octetstring "High Broadcast on Ethernet 0"
  eventDescription.1 : DISPLAY STRING- (ascii): High Broadcast on Ethernet 0

# snmpset -c private 172.16.97.132 1.3.6.1.2.1.16.9.1.1.3.1
  integer 3 eventType.1 : INTEGER: SNMP-trap

# snmpset -c private 172.16.97.132 1.3.6.1.2.1.16.9.1.1.4.1 octetstring "gateway"
  eventCommunity.1 : OCTET STRING- (ASCII): gateway

# snmpset -c private 172.16.97.132 1.3.6.1.2.1.16.9.1.1.6.1
  octetstring "cisco" eventOwner.1 : OCTET STRING- (ASCII): cisco

# snmpset -c private 172.16.97.132 1.3.6.1.2.1.16.9.1.1.7.1 integer 1
  eventStatus.1 : INTEGER: valid
```

2. Specify the thresholds and relevant parameters (RMON alarm) using the following SNMP SET operations:

```
# snmpset -c private 172.16.97.132 1.3.6.1.2.1.16.9.1.1.2.2
  octetstring "normal broadcast received on ethernet 0"
  eventDescription.2 : DISPLAY STRING- (ASCII): normal broadcast
  received on ethernet 0

# snmpset -c private 172.16.97.132 1.3.6.1.2.1.16.9.1.1.3.2 integer 2
  eventType.2 : INTEGER: log

# snmpset -c private 172.16.97.132 1.3.6.1.2.1.16.9.1.1.6.2 octetstring "cisco"
  eventOwner.2 : OCTET STRING- (ASCII): cisco

# snmpset -c private 172.16.97.132 1.3.6.1.2.1.16.9.1.1.7.2 integer 1
  eventStatus.2 : INTEGER: valid
```

3. Poll these tables to verify that the eventTable entries were made on the device.

```
% snmpwalk -v 1 172.16.97.132 private .1.3.6.1.2.1.16.9.1

# snmpset -c private 172.16.97.132 1.3.6.1.2.1.16.3.1.1.2.1 integer 60
  alarmInterval.1 : INTEGER: 60

# snmpset -c private 172.16.97.132 1.3.6.1.2.1.16.3.1.1.3.1
 objectidentifier .1.3.6.1.2.1.2.2.1.12.2
  alarmVariable.1 : OBJECT IDENTIFIER:
.iso.org.dod.internet.mgmt.mib2.interfaces.ifTable
  ifEntry.ifInNUcastPkts.2

# snmpset -c private 172.16.97.132 1.3.6.1.2.1.16.3.1.1.4.1 integer 2

alarmSampleType.1 : INTEGER: deltaValue

# snmpset -c private 172.16.97.132 1.3.6.1.2.1.16.3.1.1.7.1 integer 500
  alarmRisingThreshold.1 : INTEGER: 500

# snmpset -c private 172.16.97.132 1.3.6.1.2.1.16.3.1.1.8.1 integer 0
```

```
  alarmFallingThreshold.1 : INTEGER: 0

# snmpset -c private 172.16.97.132 1.3.6.1.2.1.16.3.1.1.9.1 integer 1
  alarmRisingEventIndex.1 : INTEGER: 1

# snmpset -c private 172.16.97.132 1.3.6.1.2.1.16.3.1.1.10.1 integer 2
  alarmFallingEventIndex.1 : INTEGER: 2

# snmpset -c private 172.16.97.132 1.3.6.1.2.1.16.3.1.1.11.1 octetstring
"cisco"
  alarmOwner.1 : OCTET STRING- (ASCII): cisco

# snmpset -c private 172.16.97.132 1.3.6.1.2.1.16.3.1.1.12.1 integer 1
  alarmStatus.1 : INTEGER: valid
```

4. Poll these tables to verify that the alarmTable entries were set.

```
% snmpwalk -v 1 172.16.97.132 private .1.3.6.1.2.1.16.3.1
```

show interface

This example is a result of the **show interface** command.

gateway> **show interface *ethernet 0***

```
Ethernet0 is up, line protocol is up
Hardware is Lance, address is 0000.0c38.1669 (bia 0000.0c38.1669)
Description: NMS workstation LAN
Internet address is 172.16.97.132/24
MTU 1500 bytes, BW 10000 Kbit, DLY 1000 usec, rely 255/255, load 1/255
Encapsulation ARPA, loopback not set, keepalive set (10 sec)
ARP type: ARPA, ARP Timeout 04:00:00
Last input 00:00:00, output 00:00:00, output hang never
Last clearing of "show interface" counters never
Queueing strategy: fifo
Output queue 0/40, 27 drops; input queue 0/75, 0 drops
5 minute input rate 1000 bits/sec, 2 packets/sec
5 minute output rate 1000 bits/sec, 1 packets/sec
21337627 packets input, 3263376846 bytes, 0 no buffer

Received 7731303 broadcasts , 0 runts, 0 giants
0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
0 input packets with dribble condition detected
17328035 packets output, 2824522759 bytes, 0 underruns
174 output errors, 44368 collisions, 4 interface resets
0 babbles, 0 late collision, 104772 deferred
174 lost carrier, 0 no carrier
0 output buffer failures, 0 output buffers swapped out
```

Configuration Management

The goal of configuration management is to monitor network and system configuration information so that the effects on network operation of various versions of hardware and software elements can be tracked and managed.

Configuration Standards

With an increasing number of network devices deployed, it is critical to be able to accurately

identify the location of a network device. This location information should provide a detailed description meaningful to those tasked with dispatching resources when a network problem occurs. To expedite a resolution if a network problem occurs, make certain to have available contact information of the person or department responsible for the devices. Contact information should include telephone number and the name of the person or department.

Naming conventions for network devices, starting from device name to individual interface, should be planned and implemented as part of the configuration standard. A well defined naming convention provides personnel with the ability to provide accurate information when troubleshooting network problems. The naming convention for devices can use geographical location, building name, floor, and so forth. For the interface naming convention, it can include the segment to which a port is connected, name of connecting hub, and so forth. On serial interfaces, it should include actual bandwidth, local data link connection identifier (DLCI) number (if Frame Relay), destination, and the circuit ID or information provided by the carrier.

Configuration File Management

When you add new configuration commands on existing network devices needs, you must verify the commands for integrity before actual implementation takes place. An improperly configured network device can have a disastrous effect on network connectivity and performance. Configuration command parameters must be checked to avoid mismatches or incompatibility issues. It is advisable to schedule a thorough review of configurations with Cisco engineers on a regular basis.

A fully functional CiscoWorks2000 Essentials allows for backing up configuration files on routers and Cisco Catalyst switches automatically. The security feature of Essentials can be used to perform authentication on configuration changes. A change audit log is available to track changes and the user name of individuals issuing changes. For configuration changes on multiple devices, two options are available: the web-based NetConfig in the current version of CiscoWorks2000 Essentials or the **cwconfig** script. Configuration files can be downloaded and uploaded using CiscoWorks2000 Essentials utilizing the predefined or user-defined templates.

These functions can be accomplished with the configuration management tools in CiscoWorks2000 Essentials:

- Push configuration files from the Essentials configuration archive to a device or multiple devices
- Pull the configuration from the device to the Essentials archive
- Extract the latest configuration from the archive and write it to a file
- Import configuration from a file and push the configuration to devices
- Compare the last two configurations in the Essentials archive
- Delete configurations older than a specified date or version from the archive
- Copy the startup configuration to the running configuration

Inventory Management

The discovery function of most network management platforms is intended to provide a dynamic listing of devices found in the network. Discovery engines such as those implemented in network management platforms should be utilized.

An inventory database provides detailed configuration information on network devices. Common

information includes models of hardware, installed modules, software images, microcode levels, and so on. All these pieces of information are crucial in completing tasks such as software and hardware maintenance. The up-to-date listing of network devices collected by the discovery process can be used as a master list to collect inventory information using SNMP or scripting. A device list may be imported from CiscoWorks2000 Campus Manager into the inventory database of CiscoWorks2000 Essentials to obtain an up-to-date inventory of Cisco Catalyst switches.

Software Management

A successful upgrade of Cisco IOS images on network devices requires a detailed analysis of the requirements such as memory, boot ROM, microcode level, and so on. The requirements are normally documented and available on Cisco's web site in the form of release notes and installation guides. The process of upgrading a network device running Cisco IOS includes downloading a correct image from CCO, backing up the current image, making sure all hardware requirements are met, and then loading the new image into the device.

The upgrade window to complete device maintenance is fairly limited for some organizations. In a large network environment with limited resources, it might be necessary to schedule and automate software upgrades after business hours. The procedure can be completed either using scripting language such as Expect or an application written specifically to perform such a task.

Changes to software in network devices such as Cisco IOS images and microcode versions should be tracked to assist in the analysis phase when another software maintenance is required. With a modification history report readily available, the person performing the upgrade can minimize the risk of loading incompatible images or microcode into network devices.

Performance Management

Service Level Agreement

A service level agreement (SLA) is a written agreement between a service provider and their customers on the expected performance level of network services. The SLA consists of metrics agreed upon between the provider and its customers. The values set for the metrics must be realistic, meaningful, and measurable for both parties.

Various interface statistics can be collected from network devices to measure the performance level. These statistics can be included as metrics in the SLA. Statistics such as input queue drops, output queue drops, and ignored packets are useful for diagnosing performance-related problems.

At the device level, performance metrics can include CPU utilization, buffer allocation (big buffer, medium buffer, misses, hit ratio), and memory allocation. The performance of certain network protocols is directly related to buffer availability in network devices. Measuring device-level performance statistics are critical in optimizing the performance of higher-level protocols.

Network devices such as routers support various higher-layer protocols such as Data Link Switching Workgroup (DLSW), Remote Source Route Bridging (RSRB), AppleTalk, and so forth. Performance statistics of wide-area network (WAN) technologies including Frame Relay, ATM, Integrated Services Digital Network (ISDN), and others can be monitored and collected.

Performance Monitoring, Measurement, and Reporting

Different performance metrics at the interface, device, and protocol levels should be collected on a regular basis using SNMP. The polling engine in a network management system can be utilized for data collection purposes. Most network management systems are capable of collecting, storing, and presenting polled data.

Various solutions are available in the marketplace to address the needs of performance management for enterprise environments. These systems are capable of collecting, storing, and presenting data from network devices and servers. The web-based interface on most products makes the performance data accessible from anywhere in the enterprise. Some of the commonly deployed performance management solutions include:

- InfoVista VistaView
- SAS IT Service Vision
- Trinagy TREND

An evaluation of the above products will determine if they meet the requirements of different users. Some vendors support integration with network management and system management platforms. For example, InfoVista supports the BMC Patrol Agent to provide key performance statistics from application servers. Each product has a different pricing model and capabilities with the base offering. Support for performance management features for Cisco's devices such as NetFlow, RMON, and Cisco IOS Service Assurance Agent/Response Time Reporter (RTR/SAA CSAA/RTR) is available on some solutions. Concord has recently added support for Cisco's WAN switches that can be used to collect and view performance data.

The CSAA/RTR Service Assurance Agent (SAA)/Response Time Reporter (RTR) feature in Cisco IOS can be utilized for measuring the response time between IP devices. A source router configured with CSAA configured is capable of measuring the response time to a destination IP device that can be a router or an IP device. The response time can be measured between the source and the destination or for each hop along the path. SNMP traps can be configured to alert management consoles if the response time exceeds the predefined thresholds.

Recent enhancements to Cisco IOS extends the capabilities of CSAA to measure the following:

- HyperText Transfer Protocol (HTTP) service performanceDomain name system (DNS) lookupTransmission control protocol (TCP) connectHTTP transaction time
- Interpacket delay variance (jitter) of Voice over IP (VoIP) traffic
- Response time between end points for a specific quality of service (QoS)IP type of service (ToS) bits
- Packet loss using CSAA generated packets

Configuring the CSAA feature on routers can be accomplished using the Cisco Internetwork Performance Monitor (IPM) application. The CSAA/RTR is imbedded in many but not all feature sets of the Cisco IOS software. A release of the Cisco IOS software release that supports CSAA/RTR must be installed on the device that IPM uses to collect performance statistics. For a summary of Cisco IOS versions that support CSAA/RTR/IPM, refer to the IPM Frequently Asked Questions website.

Additional information regarding IPM includes:

- Overview of IPM
- Service Assurance Agent

Performance Analysis and Tuning

User traffic has increased significantly and has placed a higher demand on network resources. Network managers typically have a limited view on the types of traffic running in the network. User and application traffic profiling provides a detailed view of the traffic in the network. Two technologies, RMON probes and NetFlow, provide the ability to collect traffic profiles.

RMON

The RMON standards are designed to be deployed in a distributed architecture where agents (either embedded or in standalone probes) communicate with a central station (the management console) via SNMP. The RFC 1757 RMON standard organizes monitoring functions into nine groups to support Ethernet topologies, and adds a tenth group in RFC 1513 for Token Ring-unique parameters. Fast Ethernet link monitoring is provided in the framework of the RFC 1757 standard, and Fiber-Distributed Data Interface (FDDI) ring monitoring is provided in the framework of both RFC 1757 and RFC 1513.

The emerging RFC 2021 RMON specification drives remote monitoring standards beyond the Media Access Control (MAC) layer to the network and application layers. This setup enables administrators to analyze and troubleshoot networked applications such as Web traffic, NetWare, Notes, e-mail, database access, Network File System (NFS), and others. RMON alarms, statistics, history, and host/conversation groups can now be used to proactively monitor and maintain network availability based on application-layer traffic-the most critical traffic in the network. RMON2 enables network administrators to continue their deployment of standards-based monitoring solutions to support mission-critical, server-based applications.

The following tables list the functions of the RMON groups.

RMON Group (RFC 1757)	**Function**
Statistics	Counters for packets, octets, broadcasts, errors, and offers on the segment or port.
History	Periodically samples and saves statistics group counters for later retrieval.
Hosts	Maintains statistics on each host device on the segment or port.
Host Top N	A user-defined subset report of the Hosts group, sorted by a statistical counter. By returning only the results, management traffic is minimized.
Traffic Matrix	Maintains conversation statistics between hosts on the network.
Alarms	A threshold that can be set on critical RMON variables for proactive management.

Events	Generates SNMP traps and log entries when an Alarms group threshold is exceeded.
Packet Capture	Manages buffers for packets captured by the Filter group for uploading to the management console.
Token Ring	Ring station—detailed statistics on individual stations Ring station order—an ordered list of stations currently on the ring Ring station configuration—configuration and insertion/removal per station Source routing—statistics on source routing, such as hop counts, and others

RMON2	**Function**
Protocol Directory	Protocols for which the agent monitors and maintains statistics.
Protocol Distribution	Statistics for each protocol.
Network Layer Host	Statistics for each network layer address on the segment, ring, or port.
Network Layer Matrix	Traffic statistics for pairs of network layer addresses.
Application Layer Host	Statistics by application layer protocol for each network address.
Application Layer Matrix	Traffic statistics by application layer protocol for pairs of network layer addresses.
User-definable History	Extends history beyond RMON1 link-layer statistics to include any RMON, RMON2, MIB-I, or MIB-II statistics.
Address Mapping	MAC-to-network layer address bindings.
Configuration Group	Agent capabilities and configurations.

NetFlow

The Cisco NetFlow feature allows detailed statistics of traffic flows to be collected for capacity planning, billing, and troubleshooting functions. NetFlow can be configured on individual interfaces, providing information on traffic passing through those interfaces. The following types of information are part of the detailed traffic statistics:

- Source and destination IP addresses
- Input and output interface numbers
- TCP/UDP source port and destination ports
- Number of bytes and packets in the flow
- Source and destination autonomous system numbers

- IP type of service (ToS)

NetFlow data gathered on network devices is exported to a collector machine. The collector performs functions such as reducing the volume of data (filtering and aggregation), hierarchical data storage, and file system management. Cisco provides NetFlow Collector and NetFlow Analyzer applications for gathering and analyzing data from routers and Cisco Catalyst switches. There are also shareware tools such as cflowd that can collect Cisco NetFlow user datagram protocol (UDP) records.

NetFlow data is transported using UDP packets in three different formats:

- Version 1—The original format supported in the initial NetFlow releases.
- Version 5—A later enhancement that added Border Gateway Protocol (BGP) autonomous system information and flow sequence numbers.
- Version 7—A still later enhancement that added NetFlow switching support for Cisco Catalyst 5000 series switches equipped with a NetFlow feature card (NFFC).

Versions 2 through 4 and Version 6 were either not released or are not supported by FlowCollector. In all three versions, the datagram consists of a header and one or more flow records.

For more information, refer to the NetFlow Services Solutions Guide white paper.

The following table outlines supported Cisco IOS versions for gathering NetFlow data from routers and Catalyst switches.

Cisco IOS Software Release	**Supported Cisco Hardware Platform(s)**	**Supported NetFlow Exported Version(s)**
11.1 CA and 11.1 CC	Cisco 7200, 7500, and RSP7000	V1 and V5
11.2 and 11.2 P	Cisco 7200, 7500, and RSP7000	V1
11.2 P	Cisco Route Switch Module (RSM)	V1
11.3 and 11.3 T	Cisco 7200, 7500, and RSP7000	V1
12.0	Cisco 1720, 2600, 3600, 4500, 4700, AS5800, 7200, uBR7200, 7500, RSP7000, and RSM	V1 and V5
12.0 T	Cisco 1720, 2600, 3600, 4500, 4700, AS5800, 7200, uBR7200, 7500, RSP7000, RSM, MGX 8800 RPM, and BPX 8600	V1 and V5
12.0(3)T and	Cisco 1600*, 1720,	V1, V5, and V8

later	2500**, 2600, 3600, 4500, 4700, AS5300*, AS5800, 7200, uBR7200, 7500, RSP7000, RSM, MGX8800 RPM, and BPX 8650	
12.0(6)S	Cisco 12000	V1, V5, and V8
—	Cisco Catalyst 5000 with NetFlow Feature Card (NFFC)***	V7

* Support for NetFlow Export V1, V5, and V8 on Cisco 1600 and 2500 platforms is targeted for Cisco IOS Software Release 12.0(T). NetFlow support for these platforms is not available in the Cisco IOS 12.0 mainline release.

** Support for NetFlow V1, V5, and V8 on the AS5300 platform is targeted for Cisco IOS Software Release 12.06(T).

*** MLS and NetFlow data export is supported in Catalyst 5000 series supervisor engine software release 4.1(1) or later.

Security Management

The goal of security management is to control access to network resources according to local guidelines so that the network cannot be sabotaged (intentionally or unintentionally). A security management subsystem, for example, can monitor users logging on to a network resource, refusing access to those who enter inappropriate access codes. Security management is a very broad subject; therefore this area of the document only covers security as related to SNMP and basic device access security.

Detailed information on advanced security include:

- Increasing Security on IP Networks
- OpenSystems

A good security management implementation starts with sound security policies and procedures in place. It is important to create a platform-specific minimum configuration standard for all routers and switches that follow industry best practices for security and performance.

There are various methods of controlling access on Cisco routers and catalyst switches. Some of these methods include:

- Access Control Lists (ACL)
- User IDs and passwords local to the device
- Terminal Access Controller Access Control System (TACACS)

TACACS is an Internet Engineering Task Force (RFC 1492) standard security protocol that runs between client devices on a network and against a TACACS server. TACACS is an authentication mechanism that is used to authenticate the identity of a device seeking remote access to a

privileged database. Variations of TACACS include TACACS+, the AAA architecture that separates authentication, authorization, and accounting functions.

TACACS+ is used by Cisco to allow a finer control over who can access the Cisco device in non-privileged and privileged mode. Multiple TACACS+ servers can be configured for fault tolerance. With TACACS+ enabled, the router and switch prompts the user for a user name and password. Authentication can be configured for login control or to authenticate individual commands.

Authentication

Authentication is the process of identifying users, including login and password dialog, challenge and response, and messaging support. Authentication is the way a user is identified prior to being allowed access to the router or switch. There is a fundamental relationship between authentication and authorization. The more authorization privileges a user receives, the stronger the authentication should be.

Authorization

Authorization provides remote access control, including one-time authorization and authorization for each service that is requested by the user. On a Cisco router, the authorization level range for users is 0 to 15 with 0 being the lowest level and 15 the highest.

Accounting

Accounting allows for the collecting and sending of security information used for billing, auditing, and reporting, such as user identities, start and stop times, and executed commands. Accounting enables network managers to track the services that users are accessing as well as the amount of network resources they are consuming.

The following table lists basic sample commands for using TACACS+, authentication, authorization, and accounting on a Cisco router and a Catalyst switch. Refer to the Authentication, Authorization, and Accounting Commands document for more in-depth commands.

Cisco IOS Command	**Purpose**
Router	
aaa new-model	Enable Authentication, Authorization, Accounting (AAA) as the primary method for access control.
AAA accounting {*system* | *network* | *connection* | *exec* | *command level*} {*start-stop* | *wait-start* | *stop-only*} {*tacacs+* | *radius*}	Enable accounting with the global configuration commands.
AAA authentication login default tacacs+	Set up the router so that connections to any terminal line configured with the login default will be authenticated with

	TACACS+, and will fail if authentication fails for any reason.
AAA authorization exec default tacacs+ none	Set up the router to check if the user is allowed to run an EXEC shell by asking the TACACS+ server.
tacacs-server host *tacacs+ server ip address*	Specify the TACACS+ server that will be used for authentication with the global configuration commands.
tacacs-server key *shared-secret*	Specify the shared secret that is known by the TACACS+ servers and the Cisco router with the global configuration command.
Catalyst Switch	
set authentication login tacacs enable [*all* | *console* | *http* | *telnet*] [*primary*]	Enable TACACS+ authentication for normal login mode. Use the console or Telnet keywords to enable TACACS+ only for console port or Telnet connection attempts.
set authorization exec enable {*option} fallback option*} [*console* | *telnet* | *both*]	Enable authorization for normal login mode. Use the console or Telnet keywords to enable authorization only for console port or Telnet connection attempts.
Set tacacs-server key *shared-secret*	Specify the shared secret that is known by the TACACS+ servers and switch.
Set tacacs-server host *tacacs+ server ip address*	Specify the TACACS+ server that will be used for authentication with the global configuration commands.
Set accounting commands enable {*config* | *all*} {*stop-only*} *tacacs+*	Enable accounting of configuration commands.

For more information on how to configure AAA to monitor and control access to the command-line interface on the Catalyst enterprise LAN switches, refer to the Controlling Access to the Switch Using Authentication, Authorization, and Accounting document.

SNMP Security

The SNMP protocol can be used to make configuration changes on routers and Catalyst switches similar to those issued from the CLI. Proper security measures should be configured on network devices to prevent unauthorized access and change via SNMP. Community strings should follow

the standard password guidelines for length, characters, and difficulty of guessing. It is important to change the community strings from their public and private defaults.

All SNMP management host(s) should have a static IP address and be explicitly granted SNMP communication rights with the network device by that predefined by IP address and Access Control List (ACL). Cisco IOS and Cisco Catalyst software provides security features that ensure that only authorized management stations are allowed to perform changes on network devices.

Router Security Features

SNMP Privilege Level

This feature limits the types of operations that a management station can have on a router. There are two types of privilege level on routers: Read-Only (RO) and Read-Write (RW). The RO level only allows a management station to query the router data. It does not allow for configuration commands such as rebooting a router and shutting down interfaces to be performed. Only the RW privilege level can be used to perform such operations.

SNMP Access Control List (ACL)

The SNMP ACL feature can be used in conjunction with the SNMP privilege feature to limit specific management stations from requesting management information from routers.

SNMP View

This feature limits specific information that can be retrieved from routers by management stations. It can be used with the SNMP privilege level and ACL features to enforce restricted access of data by management consoles. For configuration samples of SNMP View, go to snmp-server view.

SNMP Version 3

SNMP version 3 (SNMPv3) provides secure exchanges of management data between network devices and management stations. The encryption and authentication features in SNMPv3 ensure high security in transporting packets to a management console. SNMPv3 is supported in Cisco IOS Software Release 12.0(3)T and later. For a technical overview of SNMPv3, go to SNMPv3 documentation.

Access Control List (ACL) on Interfaces

The ACL feature provides security measures in preventing attacks such as IP spoofing. The ACL can be applied on incoming or outgoing interfaces on routers.

Catalyst LAN Switch Security Feature

IP Permit List

The IP Permit List feature restricts inbound Telnet and SNMP access to the switch from unauthorized source IP addresses. Syslog messages and SNMP traps are supported to notify a management system when a violation or unauthorized access occurs.

A combination of the Cisco IOS security features can be used to manage routers and Catalyst switches. A security policy needs to be established that limits the number of management stations

capable of accessing the switches and routers.

For more information on how to increase security on IP networks, go to Increasing Security on IP Networks.

Accounting Management

Accounting management is the process used to measure network utilization parameters so that individual or group users on the network can be regulated appropriately for the purposes of accounting or chargeback. Similar to performance management, the first step toward appropriate accounting management is to measure the utilization of all important network resources. Network resource utilization can be measured using the Cisco NetFlow and Cisco IP Accounting features. Analysis of the data gathered through these methods provides insight into current usage patterns.

A usage-based accounting and billing system is an essential part of any service level agreement (SLA). It provides both a practical way of defining obligations under an SLA and clear consequences for behavior outside the terms of the SLA.

The data can be collected via probes or Cisco NetFlow. Cisco provides NetFlow Collector and NetFlow Analyzer applications for gathering and analyzing data from routers and Catalyst switches. Shareware applications such as cflowd are also used to gather NetFlow data. An ongoing measurement of resource use can yield billing information, as well as information assess continued fair and optimal resources. Some commonly deployed accounting management solutions include:

- Evident Software

NetFlow Activation and Data Collection Strategy

NetFlow (network flow) is an input side-measurement technology that allows for capturing the data required for network planning, monitoring, and accounting applications. NetFlow should be deployed on edge/aggregation router interfaces for service providers or WAN access router interfaces for Enterprise customers.

Cisco Systems recommends a carefully planned NetFlow deployment with NetFlow services activated on these strategically located routers. NetFlow can be deployed incrementally (interface by interface) and strategically (on well chosen routers), rather than deploying NetFlow on every router on the network. Cisco personnel will work with customers to determine on which key routers and key interfaces NetFlow should be activated based on the customer's traffic flow patterns, network topology, and architecture.

Key deployment considerations include:

- NetFlow services should be utilized as an edge metering and access list performance acceleration tool and should not be activated on *hot* core/backbone routers or routers running at very high CPU utilization rates.
- Understand application-driven data collection requirements. Accounting applications may only require originating and terminating router flow information whereas monitoring applications may require a more comprehensive (data intensive) end-to-end view.
- Understand the impact of network topology and routing policy on flow collection strategy. For

example, avoid collecting duplicate flows by activating NetFlow on key aggregation routers where traffic originates or terminates and not on backbone routers or intermediate routers which would provide duplicate views of the same flow information.
- Service providers in the *transit carrier* business (carrying traffic neither originating nor terminating on their network) may utilize NetFlow Export data for measuring transit traffic usage of network resources for accounting and billing purposes.

Configure IP Accounting

Cisco IP accounting support provides basic IP accounting functions. By enabling IP accounting, users can see the number of bytes and packets switched through the Cisco IOS software on a source and destination IP address basis. Only transit IP traffic is measured and only on an outbound basis. Traffic generated by the software or terminating in the software is not included in the accounting statistics. To maintain accurate accounting totals, the software maintains two accounting databases: an active and a check-pointed database.

Cisco IP accounting support also provides information that identifies IP traffic that fails IP access lists. Identifying IP source addresses that violate IP access lists signals possible attempts to breach security. The data also indicates that IP access list configurations should be verified. To make this feature available to users, enable IP accounting of access list violations using the **ip accounting access-violations** command. Users can then display the number of bytes and packets from a single source that attempted to breach security against the access list for the source destination pair. By default, IP accounting displays the number of packets that have passed access lists and were routed.

To enable IP accounting, use one of the following commands for each interface in interface configuration mode:

Command	Purpose
ip accounting	Enable basic IP accounting.
ip accounting access violations	Enable IP accounting with the ability to identify IP traffic that fails IP access lists.

To configure other IP accounting functions, use one or more of the following commands in global configuration mode:

Command	Purpose
ip accounting-threshold *threshold*	Set the maximum number of accounting entries to be created.
ip accounting-list *ip-address wildcard*	Filter accounting information for hosts.
ip accounting-transits *count*	Control the number of transit records that will be stored in the IP accounting database.

Refer to Cisco Technical Tips Conventions for information on conventions used in this document.