CHAPTER 6

Network Management

Cisco multiservice management tools are standards-based and fully compatible with multivendor network environments and existing management systems. They are delivered in a layered, modular framework, with open interfaces at each layer.

The Cisco management framework is consistent with the Telecommunication Management Network (TMN). Like TMN, Cisco uses a five-layer model that defines both the logical division and the communication between areas of the service provider's business operations and management processes. Consistent with this architecture, Cisco has developed a suite of service, network, and element management solutions.

Network and Element Management

Cisco service management solutions integrate with network and element management solutions to enable continuous management and control of the entire network or individual elements, such as hubs, routers, switches, probes, and data collection devices.

Open, Standard Interfaces

The Cisco enhanced TMN-based management architecture allows service providers to readily support standards as they are defined and adopted. The architecture uses the protocol best suited to the application's needs. For example, Simple Network Management Protocol (SNMP) represents and performs operations on management objects; Trivial File Transfer Protocol (TFTP) transfers large volumes of data; and telnet provides direct access to and control of network elements via the command-line interface.

CiscoView

At the element layer, equipment and device management functions are performed by CiscoView, a GUI-based application that provides dynamic status, statistics, and comprehensive configuration information for Cisco internetworking products (switches, routers, concentrators, and adapters). CiscoView provides the following core functions:

- Graphically displays the MGX 8250 switch from a network management location, giving network managers a complete view of all Cisco products in the network without physically checking each device at remote sites
- Provides exception reporting, enabling users to grasp essential inquiry information quickly

- Displays a continuously updated physical picture of the MGX 8250 Service Modules.
- Simultaneously supports multiple switches, routers, hubs, or access servers through multiple invocation within the same session

Embedded Management Functions

The following paragraphs describe the embedded management functions.

Configuration Management

Configuration on the PXM will be done through an SNMP manager or CLI interface. All configuration information will be kept in a Management Information Base (MIB). SNMP will be used between an external management system and the platform to retrieve configuration data, provision logical and physical interfaces, distribute alarm information, gather real-time counters, and invoke diagnostic functions.

Resource Management

There are multiple network controller software modules (i.e., MPLS, PNNI) that talk with the platform software via the Virtual Switch Interface (VSI) protocol to deal with their respective network topology routing and signaling. Every connection added will take away from the Edge Concentrator's pool of limited resources (for example, total number of connections, bandwidth, or connection IDs). These controllers compete for the following resources:

- Card-level resources: number of connections.
- Port-level resources: connection identifiers (for example, DLCI, VPI, VCI, etc.) and bandwidth.

In parallel with this concept, two types of resource partitioning can be performed

- Card resource partitioning

  When a card is first brought up, the card partition consists of each controller sharing the maximum number of connections for the card. These values are enforced as the maximum number of connections against the card for that particular controller. These values are also inherited by the port resource partition when a port is created.

- Port resource partitioning

  When a port is added, the port partition contains the connection identifier, bandwidth, and number of connections space per controller. By default, the port resources are again fully shared among the controllers and the connection space values are inherited from the card partition. The values specified in the port partition are advertised to the controllers and they are bound to these limits when adding connections for that port.

Customers who have multiple controller types should perform both card and port resource partitioning. Customers who have multiple controller types and also want to do card-level resource partitioning, would have to perform card resource partitioning first. Port resource partitioning would then be performed to divide the port resources and to further divide the card resources at the port level if so desired.

Provisioning

In a feeder mode, the MGX 8250 interfaces with a BPX ATM backbone network and acts as a shelf of the BPX. In a standalone mode, the MGX 8250 interfaces with a third party ATM network.

- To add an end-to-end connection, from a local feeder to a remote feeder, add three connections that span three segments.
- To add an end-to-end connection, from a local feeder to a routing node, add two connections that span two segments.
- To add an end-to-end connection in a standalone scenario, add the local connection from the Service Module to the outbound user port.

From that point, you would add a connection in the third-party network to the desired terminating device.

The number of steps to add an end-to-end connection can dramatically be reduced by using Cisco WAN Manager (CWM). CWM will allow the user to specify the originating and terminating end and the connection parameters using a GUI interface. All segment connections are added transparently.

The Connection Manager on CWM can be used to create and maintain end-to-end connections or Permanent Virtual Circuits (PVCs). A connection consists of a source (localEnd), a destination (remoteEnd) and a set of connection parameters required for the routing.

The Connection Template Manager feature is used to define a set of parameters so that they can be reused in the Connection Manager to define connections. Templates can be saved to files and then used to create or modify connections.

The Multiple Users Security feature, in which each CWM user has their own access profile, is used to determine whether you have the permissions to use each option in the CWM Connection Manager. The security mapping for the CWM Connection Manager is

- Read Permission:List connections and view multicast connections and templates.
- Create Permission: Configure connections and do association backup.
- Modify Permission: Modify connections. In addition, you have Read Permission with the Modify Permission.
- Delete Permission: Delete connections. You can also have read permission in combination with relete permission.

Fault Management

The Cisco WAN solution is a distributed intelligent system design. All edge concentrators run independently of the Cisco WAN Manager. Should the Cisco WAN Manager become disabled, there is no effect on the network and its operations. The Cisco WAN Manager is essentially a window to the network and not the controlling “mind” of it. You can back up gateway configurations and store them on the Cisco WAN Manager.

The MGX 8250 service transparently supports end-to-end F4 and F5 OAM flows.

OAM cell processing uses a combination of hardware and software functionality within the MGX 8250.

There are several types of OAM cells implemented by the MGX 8250 itself that are used in the detection and restoration process combined with distributed network intelligence.

The MGX 8250 supports F1 flows for ATM over DS3/E3.

Connection-failure handling means supporting Alarm Indication Signal (AIS), Remote Detection Indicator (RDI), A-bit Alarm, test delay, and cc mechanism (continuity checking).

When there is a failure, AIS is generated by CPE or by the Service Module (SM). If AIS is generated by CPE, the SM will forward AIS to the network (to the other end). The other CPE, when it receives AIS, will send RDI. RDI is used as AIS acknowledgment.

For some SMs such as FRSM, the communication between the CPE and the SM will be A-bit and not AIS. In this case, if the SM receives A-bit alarm from CPE, then it will generate AIS to the network. If there is an interface failure, SM will generate an A-bit alarm and send it to the CPE.

CC mechanism is used to indicate to the far end that the connection is still alive. If the SM supports cc, then the detection of connection failure will be at the card level and not at a controller level (for example, the SM can send an OAM cell every second and if it does not receive an OAM cell, it will fail the connection generating AIS or A-bit to the CPE. The cc OAM is not the OAM loopback.

OAM loopback is configurable at a card level and is supported on RPM/B. RPM/B polls for all the connections. The timer and polling frequency is configurable per connection. RPM/B sends an OAM loopback cell every 1 second (if use default value) and detects other end failure in 3 seconds (if using the default value), which are the same as the cc timer and frequency.

Test delay is a CLI command that is used to test continuity of the connection. It uses OAM loopback cell.

Table 6-1 gives a status of the support of the cc feature for each card.

Table 6-1 cc(Continuity Check) Support

Card	cc(Continuity Check) Support
FRSM	Supports cc through OAM loopback cells
AUSM	Doesn't support cc
CESM	N/A (it doesn’t receive traffic then declares the failure)
VISM	Doesn’t support cc.

Table 6-2 summarizes the connection failure handling for the different Service Modules.

Table 6-2 Service Modules Connection Failure Handling

	FRSM	AUSM	CESM	RPM/B	VISM
Supports cc	Supports cc through OAM loopback cells	Doesn’t support cc	—	Doesn’t support cc. However it has OAM loop back that is very similar to cc	Doesn’t support cc.
cc doesn’t receive OAM cell after timeout	Notifies PXM of failure (PXM sends VSI con trap to the controller). Generates A-bit to CPE.	Notifies PXM of failure (PXM sends VSI con trap to the controller). Generates AIS to CPE	Notifies PXM of failure (PXM sends VSI con trap to the controller). Generates AIS to CPE	RPM/B DOESN’T notify PXM. RPM/B DOESN’T send anything to its Ethernet port. RPM/B itself is considered as CPE.	—

Table 6-2 Service Modules Connection Failure Handling (continued)

	FRSM	AUSM	CESM	RPM/B	VISM
Connection fails due to receiving AIS from CPE	In this case it receives A-bit from CPE, not AIS. Notifies PXM of failure Generates AIS to network	Notifies PXM of failure. Generates AIS to network.	Notifies PXM of failure. Generates AIS to network.	N/A since RPM/B considers itself as CPE.	Notifies PXM of failure. Generates AIS to network.
Connection receives AIS from network	Notifies PXM of failure. Generates RDI to network Generates A-bit to CPE.	Notifies PXM of failure. Generates RDI to network Generates AIS to CPE.	Notifies PXM of failure. Generates RDI to network Generates AIS to CPE.	—	Notifies PXM of failure. Generates RDI to network Generates AIS to CPE.
Interface failure	Notifies PXM of failure	Notifies PXM of failure	Notifies PXM of failure	DOESN'T send AIS to network for all connections belonging to failed interface. PXM in this case has to send AIS to network for all connections belonging to failed interface	Notifies PXM of failure
Card failure	Does nothing. If other end card supports cc, then other end will eventually detect failure. Otherwise, controller will set AIS on the trunk side (via VSI commit).	Does nothing. If other end card supports cc, then other end will eventually detect failure. Otherwise, controller will set AIS on the trunk side (via VSI commit).	Does nothing. If other end card supports cc, then other end will eventually detect failure. Otherwise, controller will set AIS on the trunk side (via VSI commit).	Does nothing. If other end card supports cc, then other end will eventually detect failure. Otherwise, controller will set AIS on the trunk side (via VSI commit).	Does nothing. If other end card supports cc, then other end will eventually detect failure. Otherwise, controller will set AIS on the trunk side (via VSI commit).
Test delay	Supports test delay	Supports test delay	Supports test delay	Does not support test delay	Supports test delay

OAM Loopback

The **tstcon** command uses end-to-end loopback OAM cell.If there are ATM switches in between, all switches will just pass the cell to its neighbor. (This OAM cell will be turned around at the router if the ATM connection is till the router.) If the connection is between two VISMs with VoAAL2, other-side (terminating) VISM will loop the OAM loopback cell back.

So, if you have VoIP connection, the router, being CPE device for that ATM link, will terminate the OAM loopback cell at its CPE interface. (the other side is an IP cloud on some physical medium,it could be Ethernet, Token Ring, Frame Relay, or ATM). This method is the way ATM-forum has defined the OAM F5 flow. The user can use this facility to trouble shoot ATM connection until router.

Alarms

All alarms and events generated by the applications will be converted into traps and funneled through the central SNMP agent interface. The central alarm collector will create the actual SNMP Trap PDU. It will log the alarms and forward them to all registered SNMP managers. As part of the robust trap mechanism, the alarm distributor assigns a sequence number to each trap and saves them in a circular buffer. Managers who receive trap out of sequence have the option of retrieving the missing traps from the circular buffer using SNMP.

Traps are generated by applications on the PXM cards. Traps are also generated by legacy Service Modules. Traps are recorded on the PXM in the Robust Trap Mechanism MIB, which is used to provide a concept of robustness to SNMP Traps. Certain types of traps are also classified as alarm events. An alarm has a trigger trap that leads to the alarm state and a corresponding clear trap that exits the alarm state. Alarms are logged as events. A hierarchy is maintained to reflect the current highest alarm level1 on shelf and card.

To support VISM and RPM cards, the proxy tasks has been modified to process and forward legacy-style traps generated by those Service Modules.

Fault detection includes hardware and software error malfunctions in all the MGX 8250 components, including Service Modules and common equipment. In addition, the edge concentrator provides a detailed alarm presentation, indicating alarm severity, type, date and time, alarm source, cause and state.

The edge concentrator records all log activity and maintains a historical log. For user-requested commands, alarm messages, and status messages, the disk on the switch will hold 72 hours of information under normal conditions. In addition, the information can be downloaded to the Cisco WAN Manager workstation on an as-needed basis. Performance information is stored in user-definable buckets of 5, 10, 15, 30, or 60 minutes. These performance counters can be aggregated by the Cisco WAN Manager for report generation and analysis.

Alarm data is stored in a circular buffer. If the buffer fills, the oldest entries will be overwritten first. The alarm data file may be manually transferred via TFTP to a designated workstation. This capability could be automated by writing appropriate scripts on a UNIX station.

The MGX 8250 hardware faults are identified by a combination of network element name, shelf number, slot number, port number, and front or back cards. These faults helps the operators or craft personnel to easily locate the hardware unit and to perform diagnostics or hardware replacement actions.

Alarms and problem reports are sent to the following devices or logs.

- Local Craft terminal
- Problem log
- Alarm log

- Alarm panel
- Remote peripheral device

Performance Management

The switch fabric provides real-time counters for performance monitoring as well as debugging. The real-time statistics are collected for four object types.

- Connections
- Ports
- Interfaces (lines)
- Trunks

For each object there can be several sub-objects (types of lines, ports, and so on), and for each sub-object type there are several statistics.

PXM 1

The following counters are provided for PXM1:

- Sonet Line and Trunk Counters
  - Section Counter LOSs
  - Section Counter LOFs
  - Path Counter AISs
  - Path Counter RFIs
  - Line Counter AISs
  - Line Counter RFIs
- PLCP Counters
  - dsx3PlcpRcvOOFCount
  - dsx3PlcpRcvRAICount
  - dsx3PlcpFECount
  - dsx3PlcpFEBECount
  - dsx3PlcpFEBESecCount
  - dsx3PlcpSEFEBESecCount
  - dsx3PlcpHECCount
  - dsx3PlcpHECSecCount
  - dsx3PlcpSEHECSecCount
- DS3 Counters
  - dsx3LCVCurrent
  - dsx3LESCurrent
  - dsx3LSESCurrent
  - dsx3PCVCurrent

  - dsx3PESCurrent
  - dsx3PSESCurrent
  - dsx3SEFSCurrent
  - dsx3AISSCurrent
  - dsx3UASCurrent
  - dsx3PlcpRcvOOFCount
  - dsx3PlcpRcvRAICount
  - dsx3PlcpFECount
  - dsx3PlcpFEBECount
  - dsx3PlcpFEBESecCount
  - dsx3PlcpSEFEBESecCount
  - dsx3PlcpHECCount
  - dsx3PlcpHECSecCount
  - dsx3PlcpSEHECSecCount
  - dsx3RcvLOSCount
  - dsx3RcvOOFCount
  - dsx3RcvRAICount
  - dsx3FECount
  - dsx3PlcpBip8CVCurrent
  - dsx3PlcpBip8ESCurrent
  - dsx3PlcpBip8SESCurrent
  - dsx3PlcpSEFSCurrent
  - dsx3PlcpUASCurrent
- ATM Counters—Ingress
  - Number of cells received with CLP=0 on a connection
  - Number of cells received with CLP=1 on a connection
- ATM Counters—Egress
  - Number of cells received on a connection
  - Number of cells transmitted on a connection
  - Number of cells received on a connection with EFCI bit set
  - Number of cells transmitted on a connection with EFCI bit set
- On the broadband interfaces on PXM1, the counters available are
  - Number of cells received from the port
  - Number of valid OAM cells received
  - Number of RM cells received
  - Number of cells received from the port with CLP=0
  - Number of cells received from the port with CLP=1
  - Number of cell with CLP=0 discarded

  - Number of cell with CLP = 1 discarded
  - Number of OAM cells transmitted
  - Number of RM cells transmitted
  - Number of cells transmitted for which CLP bit was set
  - Number of cells transmitted for which CLP bit was not set
- For each connection on the PXM1, the counters available are
  - Number of cells received from the port with CLP = 0
  - Number of cells received from the port with CLP = 1
  - Number of cells that were non-conforming at the GCRA-1
  - Number of cells that were non-conforming at the GCRA-2
  - Number of cell with CLP = 0 received from port and discarded
  - Number of cell with CLP = 1 received from port and discarded
  - Number of cells transmitted (to cell bus or towards trunk card)
  - Number of cells transmitted for which EFCI was not set
  - Number of cells transmitted for which EFCI was set
  - Number of cells with CLP = 0 toward port that were discarded
  - Number of cells with CLP = 1 toward port that were discarded
  - Number of EOF cells received

SRM-3T3/B

The following counters are provided for SRM-3T3/B:

- dsx3LCVCurrent
- dsx3LESCurrent
- dsx3LSESCurrent
- dsx3PCVCurrent
- dsx3PESCurrent
- dsx3PSESCurrent
- dsx3CCVCurrent
- dsx3CESCurrent
- dsx3CSESCurrent
- dsx3SEFSCurrent
- dsx3AISSCurrent
- dsx3UASCurrent
- dsx3RcvLOSCount
- dsx3RcvOOFCount
- dsx3RAICount
- dsx3FECount

- dsx3RcvFEBECounter
- dsx3RcvEXZCounter

High-Speed FRSM

The following counters are provided for high-speed FRSM cards:

- DS1 Alarm Stats
  - statDsx1LCVCurrent
  - statDsx1LESCurrent
  - statDsx1LSESCurrent
  - statDsx1CRCCurrent
  - statDsx1SEFSCurrent
  - statDsx1AISSCurrent
  - statDsx1UASCurrent
- DS1 Counter Stats
  - statDsx1RcvLOSCount
  - statDsx1RcvOOFCount
  - statDsx1RcvRAICount
  - statDsx1RcvFECount
- DS3 Alarm Stats
  - statDsx3LCVCurrent
  - statDsx3LESCurrent
  - statDsx3LSESCurrent
  - statDsx3PCVCurrent
  - statDsx3PESSCurrent
  - statDsx3PSESSCurrent
  - statDsx3SEFSCurrent
  - statDsx3AISSCurrent
  - statDsx3UASCurrent
- DS3 Counter Stats
  - statDsx3RcvLOSCount
  - statDsx3RcvOOFCount
  - statDsx3RcvRAICount
  - statDsx3RcvFECount

Frame Relay

The following counters are provided for frame relay servicess:

- Port Counters
  - statPortRcvFrames
  - statPortRcvBytes
  - statPortRcvFramesDiscCRCError
  - statPortRcvFramesDiscIllegalHeader
  - statPortRcvFramesDiscAlignmentError
  - statPortRcvFramesDiscIllegalLen
  - statPortRcvFramesUnknownDLCI
  - statPortRcvFramesDiscXceedDEThresh
  - statPortXmtFrames
  - statPortXmtBytes
  - statPortXmtFramesFECN
  - statPortXmtFramesBECN
  - statPortXmtFramesDiscXceedQDepth
  - statPortXmtBytesDiscXceedQDepth
  - statPortXmtFramesDuringLMIAlarm
  - statPortXmtBytesDuringLMIAlarm
  - statPortRcvStatusInquiry
  - statPortRcvInvalidRequest
  - statPortRcvUNISeqMismatch
  - statPortXmtStatus
  - statPortXmtAsynchUpdate
  - statPortUNISignallingTimeout
  - statPortXmtStatusInquiry
  - statPortRcvStatus
  - statPortRcvAsynchUpdate
  - statPortRcvNNISeqMismatch,
  - statPortNNISignallingTimeout
- Channel Counters
  - statChanRcvFrames
  - statChanRcvBytes
  - statChanRcvFramesDE
  - statChanRcvBytesDE
  - statChanRcvFramesDiscard
  - statChanRcvBytesDiscard

  - statChanRcvFramesDiscXceedQDepth
  - statChanRcvBytesDiscXceedQDepth
  - statChanRcvFramesDiscXceedDEThresh
  - statChanXmtFrames
  - statChanXmtBytes
  - statChanXmtFramesFECN
  - statChanXmtFramesBECN
  - statChanXmtFramesDE
  - statChanXmtFramesDiscard
  - statChanXmtBytesDiscard
  - statChanXmtFramesDiscXceedQDepth
  - statChanXmtBytesDiscXceedQDepth
  - statChanXmtFramesDiscCRCError
  - statChanXmtFramesDiscReAssmFail
  - statChanXmtFramesDuringLMIAlarm
  - statChanXmtBytesDuringLMIAlarm
  - statChanRcvFramesDiscUPC
  - statChanXmtBytesTaggedDE
  - statChanXmtFramesTaggedDE
  - statChanXmtFramesInvalidCPIs
  - statChanXmtFramesLengthViolations
  - statChanXmtFramesOversizedSDUs
  - statChanXmtFramesUnknownProtocols
  - statChanRcvFramesUnknownProtocols
  - statChanSecUpTime
  - statChanRcvBytesTaggedDE
  - statChanRcvFramesTaggedDE
  - statChanRcvBytesTaggedDE
  - statChanRcvFramesTaggedDE

FRSM-T1E1

The following counters are provided for FRSM-T1E1 cards:

- Frame Relay Port Counters
  - Received frames discarded due to Aborts
  - Received frames discarded due to illegal header (EA bit)
  - Received frames discarded due to CRC errors
  - Received frames discarded due to alignment errors

  - Received frames discarded due to unknown DLCI
  - Received frames discarded due to illegal frame length
  - Received frames discarded due to DE threshold exceeded
  - Received frames with DE already set
  - Received frames with FECN already set
  - Received frames with BECN already set
  - Received frames tagged FECN
  - Received frames
  - Received bytes
  - Transmit frames discarded due to underrun
  - Transmit frames discarded due to Abort
  - Transmit frames discarded due to egress Q-depth exceeded
  - Transmit bytes discarded due to egress Q-depth exceeded
  - Transmit frames discarded due to egress DE threshold
  - Exceeded Transmit frames
  - Transmit bytes
  - Transmit Frames with FECN set
  - Transmit Frames with BECN set
  - LMI receive status inquiry request count
  - LMI transmit status inquiry request count
  - LMI invalid receive status count
  - LMI signaling protocol (keep alive time-out count)
  - LMI sequence number error count
  - LMI receive status transmit count (in response to request)
  - LMI transmit status transmit count (in response to request)
  - Transmit frames during LMI alarm
  - Transmit bytes during LMI alarm
  - LMI update status transmit count (in response to configuration changes)
- Frame Relay Channel Counters
  - Number of frames received
  - Number of bytes received
  - Number of frames received with DE already set
  - Number of bytes received with DE already set
  - Number of frames received with unknown DLCI
  - Number of frames received but discarded
  - Number of received bytes discarded
  - Number of received bytes discarded due to exceeded Q-depth
  - Number of frames received and discarded due to: intershelf alarm

  - Exceeded DE threshold
  - Exceeded Q depth
  - Number of frames received with FECN set
  - Number of frames received with BECN set
  - Number of frames received tagged FECN
  - Number of frames received tagged BECN
  - Number of frames transmitted
  - Number of bytes transmitted
  - Number of frames transmitted with DE set
  - Number of frames discarded due to reassembly errors
  - Number of frames transmitted during LMI logical port alarm
  - Number of frames transmitted with FECN set
  - Number of frames transmitted with BECN set
  - Number of transmit frames discarded
  - Number of transmit bytes discarded
  - Number of transmit frames discarded due to: CRC error
  - Egress Q depth exceeded
  - Egress DE threshold exceeded source abort
  - Physical link failure (T1)
- ATM Cell-Related Counters
  - Number of cells transmitted to PXM
  - Number of cells transmitted with CLP bit set
  - Number of OAM AIS cells transmitted
  - Number of OAM FERF cells transmitted
  - Number of BCM cells transmitted
  - Number of OAM end-end loopback cells transmitted
  - Number of OAM segment loopback cells transmitted
  - Number of cells received from PXM
  - Number of cells received with CLP bit set
  - Number of OAM AIS cells received
  - Number of OAM FERF cells received
  - Number of BCM cells received
  - Number of OAM end-end loopback cells received
  - Number of OAM segment loopback cells received
  - Number of OAM cells discarded due to CRC-10 error

AUSM/B

The following counters are provided for AUSM/B:

- Line Counters
  - LoS occurrences
  - OoF occurrences
  - Remote loss of signal/frame (RAI) occurrences
  - All ones received (AIS) occurrences
  - Bipolar violation occurrences
  - Cyclic redundancy check (CRC) error occurrences
  - Line code violation (LCV)
  - Line errored second (LES)
  - Line severely errored second (LSES)
  - Code violation (CV)
  - Errored Second (ES)
  - SES
  - SEFS
  - AISS
  - UAS
- Port Counters (IMA ports)
  - Number of cells received from the port
  - Number of cells received with unknown VPI/VCI
  - Last unknown VPI/VCI received
  - Number of cells discarded due to error in cell header
  - Number of cells received with nonzero GFC field
  - Number of cells transmitted to the port
  - Number of cells transmitted for which EFCI was set
  - Number of egress cells discarded because of service interface physical layer alarm
- Channel Counters—Ingress
  - Number of cells received from the port on the virtual connection (VC)
  - Number of cells received with CLP = 1
  - Number of cells received with EFCI = 1
  - Number of cells received but discarded because queue exceeded queue depth
  - Number of cells received but discarded because queue exceeded CLP threshold
  - Number of cells received for which CLP was set because of UPC violations
  - Peak queue depth
  - Number of cells transmitted to cell bus
  - Number of cells transmitted to cell bus for which EFCI was set

  - Number of cells for transmission to cell bus discarded because of shelf alarm
  - Number of OAM cells received and discarded
  - Number of AIS cells received
  - Number of RDI FERF cells received
  - Number of segment loopback cells received
  - Number of segment loopback cells transmitted to cell bus
- Channel Counters—Egress
  - Number of cells received from cell bus for this virtual circuit
  - Number of cells received with CLP = 1
  - Number of cells discarded because queue exceeded queue depth (per egress queue)
  - Number of cells discarded because queue exceeded CLP threshold (per egress queue)
  - Number of OAM cells discarded
  - Number of AIS cells transmitted to port
  - Number of segment loopback cells transmitted
  - Number of segment loopback cells received from cellbus

CESM-T1E1

The following counters are provided for CESM-T1E1:

- Line Counters
  - FEBE count
  - OoF count
  - LCV count
  - FER count
  - CRC error count
- AAL1 SAR Counters
  - Number of OAM cells received
  - Number of OAM cells dropped FIFO full
  - Number of SN CRCs not correctable
  - Number of cells with SN different from SN+1
  - Number of cells received from UTOPIA interface
  - Number of cells transmitted to UTOPIA interface
- ATM Layer Counters
  - Number of cells transmitted
  - Number of cells transmitted with CLP bit set
  - Number of AIS cells transmitted
  - Number of FERF cells transmitted
  - Number of end-to-end loopback cells transmitted

  - Number of segment loopback cells transmitted
  - Number of cells received
  - Number of cells received with CLP bit set
  - Number of AIS cells received
  - Number of FERF cells received
  - Number of end-to-end loopback cells received
  - Number of segment loopback cells received
  - Number of OAM cells discarded because of CRC-10 error

CESM-T3E3

The following counters are provided for CESM-T3E3:

- DS3 Line Group
  - Dsx3LCVCurrent
  - Dsx3LESCurrent
  - Dsx3LSESCurrent
  - Dsx3UASCurrent
  - Dsx3RcvLOSCount
- Channel Counters
  - CesReassCells
  - CesGenCells
  - CesHdrErrors
  - CesSeqMismatchCnt
  - CesLostCells
  - CesChanSecUpTime
  - XmtCellsFERF
  - RcvCellsFERF
  - XmtCellsAIS
  - RcvCellsAIS
  - XmtCellsSegmentLpBk
  - RcvCellsSegmentLpBk
  - RcvCellsDiscOAM

The MGX 8250 is capable of transmitting status reports to the Element Management layer on CWM/CiscoView. All the information about the element can be maintained in the status reports including information on switching matrix, modules, interfaces, and utilization of the element.

The MGX 8250 can send a full inventory report to the EM layer concerning the modules that make up its structure, including the hardware revisions and serial numbers, and the associated operating software.

The MGX 8250 supports the capability to schedule the performance counters in 5,10, 15, 30, and 60-minute intervals, as far as data collection is concerned. The data-collection intervals (commonly known as polling cycles) can be 15-minutes, 30-minutes, or 60-minutes long. These performance data counters can be aggregated at the Cisco WAN Manager to generate daily reports.

Security Management

When the user logs into a MGX 8250 node, the user is required to supply the user ID and password and the slot to direct input to. When the operator adds a new user, the new user has to specify the user ID and the access level. The choices for the privilege are GROUP1, GROUP2, GROUP3, GROUP4, GROUP5, or ANYUSER.

Each telnet session may be terminated by the user or by a timer, whose timer value is determined when the session is established. The timer signals the telnet connection to be terminated if the user does not provide any input for a certain period of time.

In contrast to every other Service Module in MGX 8250, the RPM will be driven by IOS CLI. RPM also requires the user to log in again using a possibly different user ID and password. This IOS-style authentication provides an initial entry to the router. Further authentication is required if the user needs access to more privileged commands.

On the MGX 8250 platform, the userID and password is stored in the disk database. The user ID and password are not encrypted.

Accounting Management

Statistics are collected by the MGX 8250 periodically. The Cisco WAN Manager allows usage data collection from network connections and interfaces for innovative usage-based billing to customers.

The MGX 8250 will maintain CDRs for PVCs. The following information is contained in the CRDs:

- PVC type CBR, VBR-RT, VBR NRT, ABR, UBR
- Traffic descriptor
- QoS parameters
- Traffic volume number of cells
- Date/time period during which data were collected

Embedded Management Interfaces

This section is divided into two subsections.

- SNMP
- Command Line Interface

SNMP

The Cisco WAN Manager, which integrates with HPOV, provides a complete and robust SNMP network management platform with a graphical user interface (GUI).

The WAN Manager Event Log displays descriptions of network- and operator-generated occurrences. Internally, event descriptions are generated as a result of the trap information, that transpires between the network management system and the network agents. Simple Network Management Protocol (SNMP) processes controls these traps.

An SNMP agent is software that is capable of answering valid queries from an SNMP station (such as the Cisco WAN Manager workstation), about information defined in the Management Information Base (MIB). A network device that provides information about the MIB to Cisco WAN Manager has an SNMP agent. Cisco WAN Manager and the SNMP agents exchange messages over the network's transport layer protocol.

Command Line Interface

The MGX 8250 Control Point Software provides a single and integrated point of control for managing the platform. It provides full-shelf and interface management for all hardware modules, service provisioning, and fault finding/diagnostic support for the complete shelf.

The preferred tools for configuring, monitoring, and controlling an MGX 8250 edge concentrator are the CiscoView and Cisco WAN Manager applications for equipment management and connection management, respectively.

The command line interface (CLI) is highly applicable during initial installation, troubleshooting, and any situation where low-level control is useful.

Each command falls into a range of command privilege levels. When a user ID is created, it is assigned a privilege level and the user can issue commands allowed by that level only.

The MGX 8250 provides the following CLI features:

- CLI access through serial console port on PXM
- CLI access through serial modem port on PXM
- CLI access through Ethernet port on PXM
- Maximum number of simultaneous CLI telnet sessions—10
- Complete CLI support for PXM platform software including SRM-3T3/B functions.
- Complete set of CLI commands for RPM
- Complete set of CLI commands on the following Service Modules: FRSM (8T1/E1, HS1/B, HS2, 2T3/E3, CT3), CESM (8T1/E1, T3/E3), AUSM/B (8T1/E1)

The standard telnet command that is available from both the HPOV's topology map and the CWM topology map supports telnet access to MGX 8250. The telnet session will give the user access to the PXM card. From the PXM card, the user can navigate to the desired Service Module by entering the **cc** command.

Management Tools

This section has been split into the following four sections.

- CiscoView
- Cisco WAN Manager
- Cisco Info Center
- Cisco Provisioning Center

CiscoView

CiscoView is a GUI-based device management software application that provides dynamic status, real-time counters, and comprehensive configuration information for the Cisco internetworking products (switches, routers, concentrators, and adapters). CiscoView graphically displays a real-time physical view of Cisco devices. Additionally, this SNMP-based network management tool provides monitoring functions and offers basic troubleshooting capabilities.

Using CiscoView, users can easily understand the tremendous volume of management data available for internetworking devices, because CiscoView organizes it into graphical device representations presented in a clear, consistent format.

CiscoView will be used as the element management tool for the MGX 8250. CiscoView interacts directly with the edge concentrator agent.

CiscoView software can be integrated with several of the leading SNMP-based network management platforms, providing a seamless, powerful network view. It is also included within CW2000. CiscoView software can also be run on UNIX workstations as a fully functional, independent management application.

The key functions are

- Graphically displays the MGX 8250 from a centralized network management location, giving network managers a complete view of the MGX 8250 and the other Cisco products in the network without physically checking each device at remote sites
- Oriented for exception reporting, allowing users to quickly grasp essential inquiry information
- GUI that shows a continuously updated physical picture of the MGX 8250 Service Modules and other physical components including the routers, hubs, or access servers in the network
- Invoked several times in the same session to simultaneously support multiple switches, routers, hubs, or access servers
- CiscoView displays two primary types of information:
  - Configuration information includes data such as information about a device chassis, controller card and interface cards. It is displayed in CiscoView Configuration windows.
  - Performance information includes data such as the number of Ethernet errors during a given period. It is displayed in the CiscoView Monitor windows, that is also referred to as dashboards.
- Integrated with the following network management platforms to provide a seamless and powerful system to manage Cisco devices:
  - OpenView
  - IBM NetView for AIX

Cisco WAN Manager

Cisco WAN Manager (earlier known as StrataView Plus) is an SNMP-based multiprotocol management software package designed specifically for wide-area multiservice networks. It provides integrated service management and process automation to simplify the management of even the most complex networks. The Cisco WAN Manager allows you to monitor usage, provision connections, detect faults, configure devices, and track network statistics.

Cisco WAN Manager is designed to address the significant demands of managing and operating next-generation wide-area multiservice networks. The multiservice environment is more complex, with a greater number of connections and wider variety of services, making the administration of the network a potentially impossible task without the right tools.

Based on a robust, scalable architecture, Cisco WAN Manager not only meets today's business requirements for network control and operation, but also integrates with other Cisco network management products to provide end-to-end service management of wide-area multiservice networks.

The following features are available with Cisco WAN Manager:

- Scalability—Today’s wide-area multiservice networks may start out as a few-node network but can grow into a several hundreds-node network. Cisco WAN Manager is optimized for scalability and is designed to scale as the network grows in both enterprise and service provider environments.
- Management and Operations—Cisco WAN Manager software provides powerful fault, configuration, and performance management capabilities for the wide-area multiservice network. A user-friendly, graphics-oriented interface running under HP OpenView for Solaris platforms, IBM NetView for AIX, and HPOV for HP-UX platforms lets network managers quickly provision new services and view the entire network at once to identify and isolate network problems.
- Service Management API and Integration—Seamless integration into an existing network management environment is critical in providing end-to-end service management. Cisco WAN Manager Service Agent provides an SNMP interface for network and service layer management views and control. This feature enables automated provisioning and fault management and provides a basis for other higher-level service management applications. These applications can be the Cisco Service Management applications, service provider’s Operations Support Systems (OSS), or other third party vendor value-added applications. With this interface, Cisco WAN Manager can seamlessly integrate into the customer’s network management environment.
- Performance and Capacity Management—As the cost of high-speed wide-area networks (WANs) increase with bandwidth, there is a greater demand for performance and capacity planning, and cost justification and allocation. Cisco WAN Manager Statistics Agent software collects comprehensive network statistics for cost allocation, performance management, and capacity planning. The Statistics Agent uses TFTP, which is optimized for bulk data transfer, and can upload in excess of three million usage statistics per hour per agent. The data is then stored in a standard SQL database for historical reporting and trend analysis.

Layer 2/Layer 3 Connection Management will be enhanced to support RPM as one of the ATM end points in end-end connection setup. CMProxy (Service Agent) will be enhanced to support this functionality. RPM port provisioning will not be supported via PortProxy.

Cisco Info Center

Cisco Info Center is a real-time, high-performance service-level monitoring and diagnostic tool that provides network fault monitoring, trouble isolation, and real-time service-level management for multitechnology, multivendor networks. Cisco Info Center is designed to help operators focus on important network events, offering a combination of filtering, alarm reduction rules, flexible alarm

viewing, and partitioning. It enables service levels to be specified and monitored, providing valuable insight into SLA conformance. Customer, VPN, and administrative partitioning and distribution of information are also supported by Cisco Info Center, further enhancing service providers' ability to manage the network and extend SLA monitoring capabilities to their customers. For example, a fault on an ATM trunk or change in an ATM grooming parameter may affect an IP VPN service. Using Cisco Info Center, a network operator is able to quickly focus on service-affecting alarms and understand both the services and customers affected by the fault. Service providers can also use Cisco Info Center's information partitioning capabilities to make this information available to their customers via the Web as an added service dimension.

Key Benefits

The key benefits of the Cisco Info Center are

- Simplified operations—By integrating alarms and events from multiple technologies and vendors into a single environment, operators have to learn and use only one platform for troubleshooting and diagnostics. Automation helps reduce the amount of manual effort required to resolve problems.
- Scalable and highly distributable architecture—Distributed client/server architecture allows for configuration of distributed multimanagement domains. Data filtering and deduplicating features let operators monitor domains by workgroups, geographies, or customer ID.
- Enhanced service offering—This feature supports Web/Java-based service-level monitoring applications developed by service providers so end-customers have ready access to their portion of the network.
- Integrated layer 2 and layer 3 monitoring—Facilitates intelligent resource and service assurance monitoring with consolidated event and alarm management across the entire network, enabling end-to-end service management support.
- Powerful administrative interface—Highly customizable to suit a network manager’s specific viewing requirements. Java and Web-based front-end support enables operators to use Web technologies for integrated fault management and transfer of information to customers to prove compliance with service-level agreements (SLA).
- Highly customizable event correlation engine—Empowers the operators to interpret data, while event-triggered actions can be configured to respond to certain behaviors automatically. The flexible combination of rules can be further enhanced dynamically, enabling creation of new automation rules (event-triggered actions) based on observed network behavior and combination of events. These rules can then trigger user-defined actions such as execution of auto-diagnostic scripts.
- Information overload—Operators analyze data to determine the status of a network element or a class of service. Cisco Info Center is capable of consolidating, deduplicating, partitioning, and correlating information to present the core fault data in a way that is easy to interpret.
- Multi-vendor networks—Cisco Info Center receives multiple data streams, independent of the underlying network element technology, providing a comprehensive centralized fault management center.

Cisco Provisioning Center

The Cisco Provisioning Center makes delivering services to subscribers quick and easy with a rapid, error-free means of provisioning the network infrastructure. By integrating with a service order management system, Cisco Provisioning Center dramatically reduces the costs and time-to-market issues associated with service deployment by using flow-through service provisioning. For example, Cisco Provisioning Center provides powerful capabilities to automatically map a multitechnology VPN

service to various underlying QoS parameters including weighted fair queue (WFQ) and committed access rate (CAR) in Layer 3 and available bit rate (ABR) and constant bit rate (CBR) services in Layer 2 ATM.

Another unique feature of Cisco Provisioning Center is a service validation step that is tightly integrated into the multiphase commit process. This automated step ensures that a requested service such as premium Internet access can be provided by the network prior to committing it to deployment. This reduces rollbacks and ensures the operational integrity of the service provisioning process while enabling rapid, error-free service deployment. This automated step is essential for "self-service" provisioning by customers through a Web interface.

Key Benefits

Automated, integrated provisioning with CPC offers several key benefits.

- Rapid deployment of integrated L2/L3 services such as VPN and customer network management (CNM) by maintaining a database that associates customers with network elements and the services they are providing
- Higher-quality deployment through automation improves service quality over error-prone, manual deployment methods
- Lower operation costs through improved efficiency
- Reduced training costs because less expertise is required
- Faster time to market through automation features that simplify deployment
- True end-to-end provisioning through future integrated L2/L3 provisioning capability across multiple platforms from multiple vendors

Multilayer, Multivendor Provisioning

As an integrated L2/L3 tool, CPC supports not only provisioning of Cisco equipment end to end but also supports third-party blades for Newbridge and Ascend/Lucent. A blade is a generic interface between CPC and element managers. The support and acquisition of other vendor blades are attainable through the Cisco partner Syndesis, Ltd. Flow-through APIs enable integration with existing service OSS and CNM systems for order management, billing, and capacity planning for lower time to market and reduced cost of service. Operators have a choice of defining a service in technology and equipment-neutral terms for transparent deployment across a variety of equipment, or equipment-specific terms for services that take full advantage of specific element features.

Objects Reflect Services for Activation

CPC offers customer-extensible service. Each service offered by a provider is represented by a unique service object. Service objects allow operators to view the network in terms of end-user or subscriber services, or by a traditional set of nodes, ports, and circuits. Complex configuration changes are grouped into simple units that align with subscriber service orders. This grouping simplifies and accelerates order processing and improves order consistency.

Rapidly Build Service Objects

CPC supports the rapid customization of new services, so providers can quickly develop and deploy new kinds of service by defining new classes of service object. Service objects can be added, deleted, and modified in single global operations which CPC breaks down into elementary actions on individual

subnets or equipment. Decisions about how a service should be laid in are made by CPC and can be viewed by network operators or OSS applications. CPC ensures that the operation is applied successfully to all elements of the network in a coordinated manner. If any elementary action fails, then the entire operation is automatically rolled back and the original configurations are restored.

Centralized Database and Network Model

CPC is based on client/server architecture to support distributed computing through relational database systems. CPC runs under UNIX on Informix Version 7 and above and Solaris 2.5.1 and above. The distributed architecture allows CPC to address a full range of service provider capacity and throughput requirements.

The CPC database contains both the current state of the network configuration plus pending changes in the process of being deployed. A CPC administrator can view these events and decide when to upload topology information, or automated scripts can automatically upload the information.

Open Flow-Through Interfaces

The GUI allows operators to directly interact with service objects through a visual interface.

However, automated configuration is available using the flow-through interface, which allows provisioning and order processing applications to make high-level calls for configuration services. CPC can communicate with other applications via the flow-through interface using UNIX shell scripts, Java applets, or CORBA middleware.

The flow-through interface allows CPC to become an integral component of a service provider's total service creation and management system. Orders can flow directly from an existing order processing or customer care system into CPC for immediate service activation. Operators can view services, components of services, network connections, transactions, network elements, change requests, and logs.

Reliable Service Activation

CPC is based on advanced change management features that provide unprecedented reliability and control over service activation. All configuration changes associated with the same change to a service are applied in a single, network-wide transaction. Each change begins as a change request (CR) and includes an associated audit log.

The CPC database tracks resource allocation so that other system components always know what is available. When a service is created, service threaders use the resource and topology information to find the optimal end-to-end path through the network that satisfies a specified QoS level. Using this generic functionality, CPC-based systems support features such as load sharing among Network-to-Network Interface (NNI) links and failure recovery based on the subscribed class of service (CoS).

As an application, CPC sits in the network and service management layers of the TMN model. Element managers are used by CPC through blades, which take advantage of the element manager as a configuration delivery mechanism.

Element Managers

Element managers such as Cisco WAN Manager, Cisco IP Manager, and Cisco Access Manager provide access to a specific type of equipment such as a suite of switching nodes from a particular vendor. Also called blades, element managers encapsulate specific knowledge about the equipment and translate it into an equipment-neutral representation. A blade can support a specific product, a subset of a vendor’s

entire product set, or the entire product set of a vendor. It can make use of other products such as the equipment manufacturer's own provisioning server to access network elements. CPC third-party blades for Newbridge and Cascade are attainable through Cisco partner Syndesis, Ltd.

To create a complete and working application, blades enable the CPC engine to configure all of the network elements that participate in providing a service. Services that span multiple equipment types require more than one blade.

When blades are installed in a system, subnetwork resources are published to the CPC database so that threaders can construct end-to-end services based on network policies. Threaders choose the best path after considering variables such as QoS requirements, total bandwidth consumption, under-utilized internetworks links, and lowest overall cost.