CHAPTER 6

Detecting and Correcting Failures

Reporting Failures • Isolating Failures • Correcting Failures • Dynamic Routing Around Failures

The LightStream 2020 enterprise ATM switch lets you detect failures and isolate them to the field replaceable unit (FRU). It also supplies a number of mechanisms such as redundancy and power-on servicing that let you correct failures while the node continues to operate. This chapter describes those features.

Reporting Failures

A LightStream switch or network detects failures in a number of different ways, including:

- Failure to participate in the periodic exchange of messages between cards in a chassis or between cards connected to external devices
- Failure of diagnostic tests or indications of problems from the test and control system (TCS)
- Hardware signals (loss of carrier or parity/checksum failures)
- Sending or receiving illegal messages or poorly timed messages

The LightStream switch provides several mechanisms for reporting these failures:

- Trap messages
- Network statistics
- LEDs

Using these mechanisms, a network controller can determine if the network is having a problem and work to isolate the failure.

Trap Messages

When an error condition or a change in status occurs, software processes generate trap messages, or simply, traps. Traps usually provide the first indication of a problem or a potential problem in your network. Subsequent troubleshooting procedures can be based on the information provided in the trap message. Some trap messages require immediate action; others may provide important information but do not require any action.

The LightStream switch generates the following types of traps:

- SNMP — The SNMP MIB-2 specifications define simple network management protocol (SNMP) traps.
- Operational — A network operator uses operational traps to find and correct problems.

- Informational, trace, and debug -- Your customer support representative generally uses informational, trace, and debug traps to perform advanced troubleshooting and software debugging.

All trap messages are formatted according to SNMP MIB-2 specifications.

You can record trap messages in a log file or display them on a terminal. By default, the LightStream switch records SNMP, operational, and informational traps in a log file on its local network processor (NP) disk and displays SNMP and operational traps on the local console (if one is attached).

The LightStream switch lets you customize the trap log and display. You can also enable a specific trap. You can select which types of traps are reported by setting the trap severity. In addition, you can turn the trap log off, view the trap log from CLI or the LynxOS (the real-time, UNIX-like operating system) shell, or move the trap log to another system and view it there. This lets you display a particular trap without having to display all the traps at that level.

For more details on traps, see the *LightStream 2020 Traps Reference Manual.*

Network Statistics

You can use the statistics facilities provided by the LightStream switch for a variety of purposes. For instance, you can use statistics to evaluate network performance and usage or to troubleshoot a problem.

The LightStream network provides a predefined set of statistics for every port. These per port statistics provide such information as the number of packets sent and received and the number of send and receive errors.

You can tailor statistics collection to your own needs by using the LightStream switch data collection facility (called the Collector). Using the Collector, you can determine which management information base (MIB) variables you want to collect and the collection interval. You can save the collection in a file that can be viewed from a local or remote CLI or moved to another workstation or host and viewed there.

For a more details on statistics collection, see the *LightStream 2020 Network Operations Guide.*

LEDs

There are a number of LEDs on the bulkheads of many cards in a LightStream switch. They serve several purposes:

- LEDs indicate that basic power is available to the card.
- LEDs guide you to a broken card, or to one that has failed its diagnostics.
- LEDs give an informal indication that some traffic is flowing through the node.
- LEDs indicate the status of parts of the TCS that cannot be obtained through the TCS itself. For example, LEDs indicate which TCS hub is primary. (Problems with TCS hub switchover cannot be diagnosed from the TCS itself.)

The switch card, NP, and line card LEDs are visible from the front of the LightStream chassis. The LEDs on the access cards are visible from the rear of the chassis.

The *LightStream 2020 Installation and Troubleshooting Manual* describes the LEDs for each LightStream card.

Isolating Failures

LightStream diagnostics let you isolate hardware failures to a field replaceable unit (FRU). Diagnostics available are power-on self tests (POSTs) that provide a high-level check of the hardware, and diagnostic packages that provide in-depth testing of hardware.

The POST runs automatically, whenever the system or a line card is powered up or when a card is reset. Each NP module, switch card module, and interface module runs POST. If a card passes POST, it has demonstrated a basic level of operability and its green RDY LED goes on. If the card fails, its yellow FLT LED goes on. You can display POST results from the TCS or the CLI using the **show** command. More detailed failure information is available through the TCS. The POST completes in approximately 1 minute.

Note Other failures also light the FLT LED. For more details, see the *LightStream 2020 Installation and Troubleshooting Manual*.

While the POST provides a high-level check of the operability of the cards, diagnostic packages stored on the NP's hard disk provide more in-depth testing; packages are provided for the NP, the switch card module and interface module. These diagnostics can be run remotely, through a telnet or modem connection, or locally from a console connected to the console port.

Most testing can be done on line. You cannot perform switch interface tests or NP tests in a single NP system without taking the switch off line. In all other cases, only the card under test is removed from service.

For a more details on diagnostics, the *LightStream 2020 LightStream 2020 Installation and Troubleshooting Manual*.

Correcting Failures

The LightStream switch is designed to have a low mean time to repair (MTTR). FRUs are easy to access and replace. In addition, the LightStream switch provides hardware redundancy capability and power-on servicing so that portions of the LightStream switch can be serviced while the unit continues to operate.

Hardware Redundancy

The LightStream switch has been designed with full critical element redundancy. Any hardware element that is critical to the operation of the system has a backup that can be brought into service automatically. The critical elements are

- Blowers
- Switch cards
- Network processors (NPs) and associated disk drives
- Power supplies

Every LightStream system has redundant blowers. Redundancy for all other elements is optional. When both blowers are functioning properly, they share the cooling load. If one blower fails, the other one has enough capacity to cool the entire unit.

If a LightStream switch has two switch card modules, one of the switch cards acts as the primary and handles all switch functions. The second switch card acts as a backup. If the primary switch card fails, the secondary switch assumes the role of primary.

If a LightStream switch has two NPs, one of the NPs acts as the primary and handles all of the NP functions for the LightStream switch. The second NP acts as a backup. It is configured exactly like the primary NP; however, it is not part of the active configuration. If the backup NP determines that the primary NP has failed, the backup NP assumes the role of primary, all interface modules perform a warm reboot, and edge interface connections are rerouted.

If the switch has two power supplies, both power supplies are connected to the same 48-volt rail and share the load between them. However, if one power supply fails, the other power supply automatically takes on the entire load without any disruption of power.

Power-On Servicing

Power-on servicing lets you remove and install components while the rest of the system remains up and running. This feature is supported for the following field replaceable units (FRUs):

- Switch card modules
- NPs
- Interface modules (line card and access card)
- Bulk power supplies
- Disk assemblies
- Blowers

The hardware and software supports power-on servicing. The hardware design of the cards and the midplane protects the components during card insertion and removal.

System processes (such as the TCS, self-configuration and network management agents) are also implemented to support power-on servicing. For instance, an NP maintains regular contact with each interface module that it controls. When the NP determines that the interface module is out of service, it updates the topology database to reflect this information and begins the process of rerouting virtual channel connections (VCCs) associated with the down interface module.

The network management system (NMS) also supports the interface module’s graceful removal from service under software control. This provides for clean statistics at the time of shutdown, and supports power-on servicing by controlling the operation of running internal interface module diagnostics, external line and loop tests, and bringing up newly installed cards.

Dynamic Routing Around Failures

Rerouting traffic is an important function of a LightStream network. The LightStream network has the ability to reroute VCCs whenever a failure of one or more communications links interrupts existing traffic flows on configured PVCs or explicitly established VCCs.

VCCs are rerouted using the standard call setup mechanisms to establish new paths. When a trunk fails, each VCC that runs through the failed trunk is recreated over a new path, if one is available. The LightStream switches at each end of a VCC are responsible establishing the new path. Between the time of the failure and the creation of each new circuit, service is temporarily disrupted on each circuit.