GIGAswitch/ATM Firmware Version 2.7
Release Notes

July 9, 1998

This document contains release notes for GIGAswitch/ATM firmware Version 2.7. Version 2.7 of the firmware provides the following enhancements:

- Web-based management with Java extensions
- Increased VPI ranges with the new DAGGL-BC line card
- Improved CBR granularity
- Improved security
- Timing on OC-12 interfaces
- Year 2000 compliance with Version 1.93 of the CMM firmware

Warning

Upgrading to this version requires CMM firmware Version 1.59 or higher. If your CMM firmware version pre-dates Version 1.59, you **MUST** upgrade the CMM firmware to Version 1.59 before upgrading to Version 1.93.

Failure to upgrade the CMM firmware properly will render the CMM card unusable !

To determine the CMM firmware version currently in use, type B at the CLK> prompt. If the firmware version is earlier than Version 1.59, you must first upgrade to Version 1.59. The Version 1.59 CMM firmware is contained in GIGAswitch/ATM firmware Version 1.2. Follow the instructions in the GIGAswitch/ATM firmware Version 1.2 Release Notes. You can find Version 1.2 in the /pub/DEC/GIGAswitchATM directory at the DIGITAL FTP site (ftp.digital.com).

Hardware Requirements

The master line card must be a QLC 2.0 (DAGGL-BA), a QLC 2.1 (DAGGL-BB), or a QLC 2.2 (DAGGL-BC) with at least a 16M SIMM (DAGME-AB) — that is, with at least 16 Mbytes of DRAM. OC-12 (DAGGL-CA/DAGGL-CB) and QLC 1.5 (DAGGL-AA) line cards may be used as slaves only.

Where to Get GIGAswitch/ATM Documentation

The Version 2.7 GIGAswitch/ATM documentation is available at the following locations:

- Web site
  `http://www.networks.digital.com/dr/gigaatm/manuals/`
- Anonymous FTP
  Host: `www.networks.digital.com`
  Directory: `/pub/networks/gigaatm/manuals`
- File names:
  qcv7c-te.ps *GIGAswitch/ATM Installation and Service* (14-slot system)
  qcv8e-te.ps *GIGAswitch/ATM System Management*
  dagbc-te.ps *ATM 4-Port Modular Line Card* (DAGGL-BC)

How to Get a Copy of GIGAswitch/ATM Firmware

The Version 2.7 firmware kit is located in a release area at the DIGITAL FTP site, ftp.digital.com. Copy the image to your system using the following commands:

```
# ftp ftp.digital.com

username: anonymous
password: (your Internet address)

ftp> cd /pub/DEC/GIGAswitchATM
ftp> bin
ftp> get AN3V2_7.tar AN3V2_7.tar
ftp> bye
```

To unpack the new image, use the following UNIX command:

```
# tar -xvf AN3V2_7.tar
```

This command creates a subdirectory within your current working directory named AN3V2_7. The following files are unpacked into the AN3V2_7 subdirectory:

```
AN3_VER27          (sample control file)
LC15V27.APP        (DAGGL-AA/AB/CA/CB application image)
LC15V27.ROM        (DAGGL-AA/AB kernel image)
LC20V27.APP        (DAGGL-BA/BB/BC application image)
LC20V27.SLV        (DAGGL-BA/BB/BC slave application image)
LC20V27.ROM        (DAGGL-BA/BB/BC kernel image)
CMM1_93.X          (CMM firmware)
```

Upgrading Firmware

The GIGAswitch/ATM should be upgraded using the following process:

Upgrade the clock card firmware to 1.93 *before* updating the line card firmware to Version 2.7. If you do not currently have, at least, CMM Version 1.59, upgrade to Version 1.59 first. Note that the new clock firmware (1.93) is backward compatible with the older revisions of clock firmware. Procedures to update the clock card firmware are in *GIGAswitch/ATM Installation and Service*.

The firmware package for Version 2.7 contains a new clock card firmware image Version 1.93. Version 1.93 of the clock firmware has a minor, but important, fix that provides year 2000 support. We recommend that you update the clock firmware image to 1.93 before you upgrade the line card firmware to Version 2.7.

To upgrade the firmware, change the "Line card Start up mode" flag in the clock to "L", or *force_image_reload* from the master. Reboot the switch to load the new line cards with the new firmware. The boot control file should be modified to point to the new images.

NOTE

If you encounter problems with loading new firmware, make sure that the boot file is correct and all the firmware images are in place as specified in the boot file.

Upgrading From Version 2.6 to Version 2.7

When you upgrade from Version 2.6 to Version 2.7, the system displays or logs the following informational messages:

```
flashApp_verify_checksum: error reading flash. offset: 200004 len: 400

TftpPut: Error reading flashduring file transfer. offset: 200200 size: 200

[slot# > ERROR: problem in downloading application.]
```

These expected messages are due to the application image size restrictions in earlier versions of the firmware.

The switch automatically reboots to run new images after the bootrom load to all slaves is complete. If any slave line card has a problem downloading the application image, the switch reboots itself or must be rebooted manually. If any slave line card in the switch is a QLCv2.0 or greater with extended memory, the switch automatically reboots itself after timing out waiting at a barrier to run new images. If none of the QLCv2.0 or greater slave line cards in the switch have extended memory, the switch must be rebooted manually to run new images.

For detailed instructions on how to download firmware images to your GIGAswitch/ATM system, see the Upgrading the Firmware section of *GIGAswitch/ATM Installation and Service*.

Upgrading From Versions Prior to Version 2.5 to Version 2.7

The procedures to upgrade the line card firmware are the same as in previous versions. However, there are some caveats when you upgrade from firmware revisions 2.1.*x* (or older) to Version 2.7. The ROM image of Version 2.7 uses larger flash widths that are available with the QLC 2.0/2.1/2.2 hardware. This is necessary to fit in the larger image sizes of the application. The first attempt to load the new firmware will produce error messages about the application being too big. After the first load procedure, reload the switch with the new firmware again. The second time, the switch uses the new ROM image, which will be able to accommodate the larger application images.

Note that Version 2.7 firmware does not require manual intervention to reboot the switch after the initial load when upgrading the switch firmware. In the past, this was necessary when the old ROM image and the new application image were incompatible. During future upgrades of the switch firmware, after loading the flash with the new ROM and application images, the switch automatically reboots itself to run the new images.

New Features

Using Web-Based Management

Browser Requirements

GIGAswitch/ATM web-based management requires a Java-enabled browser:

- Netscape Navigator Version 3.01 or later
- Internet Explorer Version 3.02 or later

Note that if you use Internet Explorer Version 4.0 or later, the browser enforces security constraints that do not allow you to perform web-based management. To modify the security settings of your browser so you can run the web-based management interface:

1. From the **View** menu, click **Internet Options….**
2. Click the **Security** tab.
3. Select **Trusted sites zone** in the **Zone** drop-down list.
4. Click the **Custom** radio button.
5. Click the **Settings…** button.
6. Scroll down the **Security Settings** tree to **Java**.
7. Under **Java permissions**, click **Custom**.
8. Click **Java Custom Settings….**
9. Click the **Edit Permissions** tab.
10. Under Unsigned Content and then **Run Unsigned Content**, click **Run in sandbox**.
11. Under **Additional Unsigned Permissions** and then **Access to All Network Addresses**, click **Enable**.
12. Click **OK**.
13. In the **Security Settings** dialog, click **OK**.
14. In the **Internet Options** dialog (make sure **Trusted sites zone** is selected in the **Zone** drop-down list), click **Add Sites**….
15. Type the URL for each GIGAswitch/ATM that you plan to access by web-based management and click Add. You can use the symbolic name (for example, http://atmswx.abc.company.com) or the numeric name (for example, http://123.45.67.89). If you plan to use both symbolic and numeric names to access the same switch, add both names. Note that you can add more GIGAswitch/ATM URLs to the trusted sites zone later.
16. Disable the checkbox labeled **Require server verification (https:) for all sites in this zone**.
17. Click **OK**.
18. In the **Internet Options** dialog box, click **OK**.

Accessing Web-Based Management

To monitor and configure your GIGAswitch/ATM system with web-based management:

1. Invoke a Java-enabled web browser.
2. Enter your switch's URL and the nonstandard HTTP port number (8080), separated by a colon (:). For example:

```
http://atmswx.abc.company.com:8080
http://123.45.67.89:8080
```

3. At the **User Name** prompt, enter the user name for either the manager account or the user account for your switch.
4. At the **Password** prompt, enter the corresponding manager account or user account password.
5. Click any tree item on the navigation tree in the left frame.

6. At the **Community Name for *agent*:** prompt, enter either the read-only or read-write SNMP community name. If you plan to modify SNMP values in your switch, you must use the read-write community name.
7. You are now ready to monitor and configure your GIGAswitch/ATM system. The right frame provides information about how to use the web-based management applet.

Accessing the Help Pages

The web management help pages are not installed on your switch. You must access them from a web server. You can copy them to your server from the following web site:

```
http://www.networks.digital.com/dr/gigaatm/manuals/
```

Supported Features

The GIGAswitch/ATM web-based management interface lets you view numerous SNMP-configurable tables that are supported by your switch. However, in this release, the RMON Event Table and the RMON Alarm Table are the only SNMP-configurable tables that you can configure from the web-based management interface.

Virtual Path (VP) Termination

VP termination support is available only on DAGGL-BA/BB/BC line cards. The DAGGL-BC line card supports 4 shaped VPs and 6 unshaped VPs with a VPI range of 0 to 127.

Use the `switch` command to configure a line card to run in VP mode. For the VP mode to take effect, reboot the switch or remove and reinsert the slave line card. Install a DAGVP-AA PHY module into port 4 only after upgrading the firmware to Version 2.5 or greater. The older firmware fails diagnostics of the (new) unrecognized DAGVP-AA module. Presence of the DAGVP-AA is necessary for shaped CBR VPs to be created. Absence of the DAGVP-AA PHY module results in firmware treating configured virtual ports (2, 3, and 4) as unshaped VBR VPs (that is, all VPs including VP 0 will be given an equal share (SCR) of the link bandwidth). This context of unshaped VBR VPs does not imply that every VP can burst up to the maximum available bandwidth on the link. It means only that the total link capacity is divided equally among all the configured VPs, and the VPs will be shaped to an SCR = link BW/*n*, *n* = number of VPs on the link.

The following are the limitations of the VP termination functionality in Version 2.7, which are likely to be irrelevant to the typical usage of VP shaping through public network tunnels:

- *Unshaped VPs*

  Version 2.7 supports 6 unshaped VPs that share the buffer/FIFO resources of VP 0. Thus, the bandwidth available for VP 0 on a link is shared between VP 0 and any unshaped VPs that are created on the link. This hardware design introduces a restriction in the use of point-to-multipoint VCCs over unshaped VPs. It is possible to have *only one branch* of a PMP VCC on VP 0 and the unshaped VPs. In other words, it is not possible to have a PMP branch on VP 0, as well as an unshaped VP. This poses restrictions on the usage of LANE on links carrying unshaped VPs.

  Therefore, we recommend that you do not use LANE on links that support unshaped VPs. This scenario occurs when LANE services are configured on a switch connected via a link carrying unshaped VPs to another switch to which LECs (trying to join ELANs configured on the remote switch) are attached. The manifestation of the problem is that the remote LECs may not be able to join an ELAN successfully. Note that the problem does not exist for LECs connected directly to the LANE services switch, or remote LECs connected to the LANE services switch via a non-VP link, or a link carrying only shaped VPs.

Note that shaped VPs are not subject to the hardware restrictions described above. Shaped VPs have their own buffer pools, FIFO resources, and CBR point-to-multipoint mapping tables. Hence, it is possible to create branches of the same PMP VCC on all the shaped VPs as well as VP 0.

- *VCI Ranges*

  VPs mapped to virtual ports 1, 2, and 3 support up 4 K VCs (0 < VCI < 4096). VPs mapped to virtual ports 4 through 10 support only up to 2K VCs (0 < VCI < 2048).

- *MBS*

  The maximum burst size of VPs must be less than 241.

Fine-Grained CBR

Version 2.7 supports CBR VCs with granularity down to approximately 100 kilobytes per second, which is an eightfold increase in granularity over previous versions of the firmware. This new feature supports low-speed voice circuits over CBR VCs in GIGAswitch/ATM. Fine-grained CBR also allows closer reservation to required bandwidths when provisioning circuits with a service provider. Coarse-grained CBR causes slight overallocation of bandwidth, which can lead to policing by the conservative leaky bucket traffic policers that some service providers use.

Refer to *GIGAswitch/ATM System Management* for information about configuration of fine-grained CBR. Restrictions and caveats in usage of finer-grained CBR are described in the Usage Guidelines section of this document.

Security

Password protection is disabled if the keyswitch on the GIGAswitch/ATM system is set to positions 2, 3, and 4. The security keyswitch must be set to position 1 at all times to provide system security. Refer to *GIGAswitch/ATM Installation and Service* and *GIGAswitch/ATM System Management* for more information about system security and password protection.

Timing on OC-12 (DAGGL-CA/DAGGL-CB) Interfaces

The default setting for OC-12 interfaces has been changed from loop timing to local timing. For environments in which OC-12 links serve as NNI links, each end of the link must source its own clock. In previous versions, both ends of the links were set to loop timing (that is, the transmit clock is derived from the receive clock), which caused intermittent problems in OC-12 links.

Known Issues

MultiChassis Manager

MultiChassis Manager can occasionally report a `no response from agent` error message when attempting to communicate with the switch. The workaround for this is to increase the MultiChassis Manager retry timer.

MultiChassis Manager displays false adpReceiveBuffer information for an inactive QLC V2 multimode fiber ATM port.

Signaling

UNI auto-sensing on links attached to a Sun ATM-622/MMF S-bus adapter V2.1 does not work because of ILMI interoperability problems. The workaround is to set the UNI version manually.

The switch does not correctly auto-sense the UNI version of an OpenVMS ATM end system that is connected to a QLC V2 line card. The workaround is to set the UNI version manually by using the `sig` command from the switch console.

SLIP

Rebooting the switch while the CMM is in SLIP mode causes the switch to reboot multiple times. The workaround is as follows:

1. During switch reboot, enter CMM local mode by entering a BREAK on a direct terminal connection or (~#) from TIP.
2. Put the CMM into console forwarding mode by typing Ctrl/O.
3. Wait for the message `switch initialization complete`.
4. Enter SLIP mode using the normal procedure.

Attempting to “ping” the switch through the host/SLIP port while the switch is NOT in SLIP mode will cause the switch to reboot. To prevent the switch from rebooting, ensure that the switch is in SLIP mode before attempting this operation.

SNMP

E1 ModPHY MIB objects `dsx1CurrentDMs`, `dsx1IntervalDMs`, and `dsx1TotalDMs` are not supported in this version. Displayed values will always be zero, regardless of the actual number of degraded minutes.

User Interface

While using some OBM and CLI menu options on a GIGAswitch/ATM 5-Slot chassis, the switch displays 14 slot positions. Ignore slot positions greater than 5.

For point-to-multipoint CBR circuits, the maximum bandwidth for an OC-3 link is 126 Mb/s. An OBM session incorrectly allows a maximum bandwidth of 133 Mb/s.

VP Link Failures

If a link goes down at the remote end of a VP tunnel (that is, the logical neighbor in an ILMI link is not physically adjacent), you need to wait at least 45 seconds before the link is restored. Likewise, if a VP is deleted at the remote end (which is similar to taking down a logical link), a 45-second wait at least is necessary before the same VP is created again. The wait allows both sides to timeout on the ILMI and keep alive, thereby ensuring an ILMI variable exchange when the logical link is restored. A problem occurs if a COLD-TRAP is not sent from the remote side and ILMI data exchange does not take place.

Hence, if the remote end of an ILMI link (tunneled through an ATM cloud) goes down temporarily, the local side timesout and gets into the TRAP, GET_NEXT cycle. However, if the TRAP sent by the remote side when it comes back up does not make it through the cloud to the local switch, there is a problem. Thus, if the remote side happens to come back alive when the local switch is doing a GET_NEXT, the remote side would reply to the GET_NEXT as usual. The local switch goes on and does the GET_ADDR and transitions to REGISTERED without ever syncing up the ILMI variables. Hence, there is a UNI mismatch scenario. This has been fixed to trigger an ILMI data exchange even if a COLD_TRAP is not received. However, this happens only if the local side times out on the keep-alive.

Documentation Errata

DIGITAL ATM Virtual Path Modular PHY Card (EK-DAGVP-IN) refers to the AMT 4-port modular line card, DAGGL-BA. This document also pertains to the DAGGL-BB and DAGGL-BC line cards.

Under Configuring Line Cards in VP Mode, the second sentence should read:

> Use the `switch -vpmode` command to configure and display VP mode parameters.

Usage Guidelines

Downgrading the Switch Firmware

After downgrading the switch to any version prior to Version 2.7 of the GIGAswitch/ATM firmware, you *must* perform an "nvdataErase" from the switch console.

If you downgrade from Version 2.7 to Version 2.*x*, make sure that your configuration file specifies the name of the application image for the master and slave on the QLCV2 line. If you do not specify the information, the console displays the following error message and the downgrade fails:

```
ERROR: config file error at line#n
```

If you downgrade QLCv1.5 and/or LC622 line cards, the console displays the following error message:

```
[slot# > ERROR: in decompressing application image from flash.]
```

This error message is expected and you can ignore it.

Rebooting to Restart Initialization

Occasionally during bootup (about once in 50 to 100 reboots), a switch can stop initialization with a message `waiting at barrier 2 … .` This is harmless and the workaround is a reboot of the switch.

Configuration Recommendations

To allow access to the LECS from other switches, we recommend that a static route to the well-known LECS address be created on those switches to the LECS-enabled switch. GIGAswitch/ATM and ATMswitch 900 systems run DEC-NNI, which is a link state routing protocol that exchanges routing/topology information. A network of switches that runs DEC-NNI automatically shares static route information if instructed to do so, as described below. Note that you need to configure switches from third-party vendors individually with the static routes (IISP). For a network of DIGITAL switches, set the static route's *forwarding slot* to the master line card's slot number and the *forwarding port* to 0. Also, the route must be **exported** (for example, `decnni -sr -conf 47 -partial -port 1:0 -exp`).

When you configure constant bit rate (CBR) circuits for E1, E3, or T3 links, set the CBR to a value that is less than 70% of the total allowable link rate. The following are the limits to CBR bandwidth allocations imposed by the switch:

- The maximum allowed CBR reservation on an OC-3 link is 126 Mbs of payload bandwidth.
- The maximum allowed CBR reservation on an OC-12 link is 510 Mbs of payload bandwidth.
- The maximum allowed CBR reservation per VC on an OC-12 link is 360 Mbs of payload bandwidth.

If you use a set of DAGGL-BB/BC line cards, all bandwidth allocations are in multiples of 100 kilobits per second. If you use a mixed card set, all bandwidth allocations are in multiples of 800 kilobits per second, which is the minimum CBR granularity in a GIGAswitch/ATM system. Note that UBR/ABR VCs can burst up to link rate, therefore they are not subject to the bandwidth limitations described in this section.

CBR Granularity

Do not hot swap a DAGGL-AA or OC-12 line card into a chassis using fine-grain CBR with DAGGL-BB/BC line cards. The firmware requires a reboot to drop back to using coarse CBR schedules.

Do not use fine granularity CBR with line cards running in VP mode.

When the switch is set for CBR fine granularity, note that an OBM session indicates the incorrect amount of CBR PVC bandwidth allocation, though the utilization is correct.

The CBR schedule computation time increases when you use fine-grained CBR; consequently the call setup time is longer. Therefore, we do not recommend you use fine-grained CBR in configurations running LANE. You need to consider the performance penalty incurred in SVC before using this feature.

Fine-grained CBR can cause increased cell delay variation on low-speed circuits.

LAN Emulation

LAN Emulation must not be used over T1/E1 links.

The `maxfwdrate` option of the `bus` command is unavailable for rate limiting. However, this option determines the amount of CBR bandwidth reserved for BUS multicast forward VCCs. Note that processor limitations prevent the BUS throughput from exceeding 2.5 Mbs. This option is useful only if the multicast traffic requirements are minimal, but the number of ELANs that need to be supported is large. Thus bandwidth reservation per ELAN can be reduced, thereby allowing support of a larger number of ELANs per switch.

Do not use the `lecs -assign -mac` command with hosts that use the ATMworks 350 and 351 adapters. These adapters do not put MAC addresses in their Configure requests, so the LECS never finds a MAC address match. If you want to assign such hosts to ELANs by address, use the `lecs -assign -atm` command.

LANE multicast VCCs are created as CBR point-to-multipoint VCCs (because of better hardware support for CBR multicast) in GIGAswitch/ATM systems. Thus, the total number of ELANs traversing a link is limited to the bandwidth of the link. Likewise, the bandwidth of the VP limits the number of ELANs on a Virtual Path. The bandwidth reservation per ELAN is approximately 4 Mbs (3 Mbs for BUS, and 1 Mbs for LES) on the switch hosting LANE services, and approximately 6 Mbs on all other switches on a network.

Redundant LANE

Version 2.7 supports 72 LECs per switch for redundant LANE.

Failure of a switch (serving as the Master LANE service for a large number of LECs) forces all the LECs to rejoin the switch providing Standby LANE service. This process may result in the following (recoverable) error messages:

- `Add failed for host 47:00:79:00:00:00:00:00:00:00:00:00:00:..`
- `STATUS: rrp_exactly_once: max timeout reached on lc 2 procid 101 name tNCCd task`
- `Failed to locate export route w/ rtentryp = …`
- `Delete failed for host 39:99:99:00:00:00:08:00…`

Note that these error messages occur because of the volume of messages that the processor has to process in such an overload condition. The system, however, will recover and all LECs will join successfully.

PVCs

Specify both the root and branch when you delete a branch of a point-to-multipoint circuit.

Network Configuration

To access the GIGAswitch/ATM switch from outside of its IP subnet, a default gateway address can be set using BOOTP via the “gw” field, as documented in *GIGAswitch/ATM Installation and Service*. However, if the switch’s IP address/netmask information is configured statically, that is, using the OBM interface,

the default gateway address must be set using the `setRoute("gateway_address")` command from the switch console. For example, if the default gateway is 192.20.0.1, use the following command:

```
ATMswitch-> setRoute("192.20.0.1")
```

To delete the default gateway, use the following command:

```
ATMswitch-> clearRecordType(832)
```

Signaling

To configure signaling parameters for VP 0 (VPI 0), specify *slot:port* only. For example, to set the UNI version of VPI 0 of line card 7, enter:

```
sig –uni 31 –port 7:1
```

To configure signaling parameters for nonzero VP, specify *slot:port:vpi*. For example, to set the UNI version of VPI 60 of line card 7, enter:

```
sig –uni 30 –port 7:1:60
```

Some implementations of UNI 3.0 signaling reject calls that are sent with the Default DEC Address Prefix. You can avoid this if you modify the prefix using the `decnni` command or by an OBM session as shown below:

1. Enter OBM.
2. Access OBM menu 6.1.4 (*Set/Show DEC Switch ID*).
3. Show the current ID using option 3 (*Show Configured DEC Switch ID*).
4. Using option 1 (*Configure DEC Switch ID*), enter the new 6-byte ID substituting the first byte 08 with 00.
5. Access OBM menu 6 to save the new DEC Switch ID to Flash using option 4.

Telnet

Invoking a Telnet session to a switch that is already engaged by TIP will cause the initial session to be locked out until the second session is terminated.

If you invoke OBM menus from TIP (Console Port), Telnet access is rejected and the following message is displayed:

```
shell is locked
```

When you exit OBM menus from TIP, access through Telnet is allowed.

Modifying VP Attributes

Traffic on all VPs terminating in a port must be stopped before you attempt to change the attributes of any of the VPs on that port. That is, use the `pvp -mod` command to change SCR/MBS of a VP terminating in line card *L* only when no traffic is flowing through the VP port in line card *L*.

VP 0 Functionality

VP 0 is also implemented as a shaped VP. VP 0 is assigned an SCR equal to the link bandwidth minus the sum of the SCRs of all the nonzero VPs. Traffic on VP 0 is shaped to its assigned SCR as well. When nonzero VPs are created on a link, VP 0 is assigned an SCR equal to the remaining bandwidth on the link. Creation of CBR VCs in VP 0 with bandwidth close to the SCR of the VP may have undesirable side effects. CBR traffic in the switch is governed by a schedule that guarantees the requested bandwidth, and shapes traffic to that bandwidth. If there exist other UBR/ABR VCs in VP 0, then the total bandwidth consumed by the UBR and CBR VCs may be more than the computed SCR of VP 0 (depending on the

incoming traffic pattern, since the different VPs throughput bandwidth of a shaped nonzero VP). Note that the CBR VCs will be restricted to the reserved bandwidth of the VCs, it is only the UBR VCs that burst beyond the SCR of the VP. If only UBR VCs exist on VP 0, or if the bandwidth of the CBR VCs on VPI 0 is less than half the available SCR, then the problem does not exist.

CLIP

The CLIP client in the switch allows access to the IP stack over an ATM link. You can use the CLIP connection to create a Telnet session over ATM and access all console commands. You can also use the CLIP connection to manage the switch over ATM with clearVISN. However, flooding the switch's CLIP client with IP data is not allowed and may cause the switch to crash. For example, flooded pings from a workstation to the switch over CLIP for an extended period of time would lead to such a predicament.

Simultaneous Telnet sessions on the ATM path (via CLIP) and the Ethernet path are not possible. However, connectivity is still available on one path (i.e., ping) when Telnetted via the other path. However, sustained simultaneous pinging on both the Ethernet and the CLIP paths to the switch can lead to lockout of the switch console. A reboot of the switch is required to get out of this state. We recommend that you keep activity on the Ethernet path to a minimum when the CLIP client is in use and receiving data.

Disabling and re-enabling the CLIP client in the switch causes it to pick a new ATM address for the CLIP end point. Ensure that the ARP entries in the end stations have aged out (or have been deleted) if there is no communication between the end stations and the switch's CLIP client after the disable/enable sequence.

CLIP PVCs

To delete a PVC connection to a CLIP client, use following commands:

```
ATMswitch-> clip -dis
ATMswitch-> clearRecordType(866)
```

Changing the IP address of the CLIP client requires a switch reboot to re-enable the CLIP for the new IP address to take effect.

Internal Error Messages

Internal error messages displayed during switch failures use a zero-based numbering scheme for both slot and port. Furthermore, the CMM module is not included, so slot 8 is reported as 6, slot 9 as 7, etc.