GIGAswitch/FDDI System

Release Notes Version 3.0

Part Number: AA–PZT9E–TE

GIGAswitch Firmware Baselevels:

- Switch Control Processor (SCP) OP 3.00, SCP BB 2.00, SCP DL 1.00
- Two-port FDDI GIGAswitch Line Card (FGL-2) 3.00
- Four-port FDDI GIGAswitch Line Card (FGL-4) 3.00
- Two-port ATM GIGAswitch Line Card (AGL-2) 3.00
- Two-port ATM GIGAswitch Line Card (AGL-2+) 3.00
- Clock card (CLK) 3.00
- Power System Controller (PSC) 2.00
- Management Information Base (MIB) 3.00

Digital Equipment Corporation
Maynard, Massachusetts

March 1996

This document was prepared using VAX DOCUMENT, Version 2.1.

Contents

Figures

Preface

This document describes new features, documentation changes, bug fixes, problems and restrictions that pertain to the GIGAswitch/FDDI V3.0 firmware release.

Intended Audience

The *GIGAswitch/FDDI System Release Notes* are intended for customers and Digital Service personnel. Read the release notes before you install, service, or use the GIGAswitch/FDDI System.

GIGAswitch/FDDI Firmware 3.0

New Code Images for V3.0

This release provides the following **new** firmware:

- FGL-2 version 3.00

  fgl2_op_300.rsx
  (fg2_30.rsx)

- FGL-4 version 3.00

  fgl4_op_300.rsx
  (fg4_30.rsx)

- AGL-2 version 3.00

  agl_op_300.rsx
  (ag2_30.rsx)

- AGL-2+ version 3.00

  aglp_op_300.rsx
  (ag2p_30.rsx)

- CLK version 3.00

  clk_op_300.rsx
  (clk_30.rsx)

- SCP version 3.00

  scp_op_3.00.tftp
  (scp_30.ftp)

New MIBs

The following new MIBs are provided:

- GIGAswitch MIB version 3.00

  mib_gs_300.txt
  (mib_30.txt)

  mib_gs_300_and_elan_270.txt
  (e_mib_30.txt)

Version 3 Overview and Installation

GIGAswitch/FDDI version 3.0 firmware provides several new features, updated documentation and bug fixes.

New Features

The new features are:

- New OBM Menus
- 24K Translation Table
- Demand Learning
- Support for AGL-2+ linecard
- Hunt Groups

These features are described in the sections that follow.

New Documentation

The updated documentation is:

- *GIGAswitch/FDDI System Out-Of-Band Management (OBM) Guide* part number EK-GOBMG-MG. B01
- *GIGAswitch/FDDI System AGL Reference Guide* part number EK-GAGL2-MG-MG. B01

Installation

HUBwatch V4

HUBwatch V4 does **not work** with this firmware release. clearVISN MultiChassis Manager V5.0 will work with V3 firmware.

Certain cautions should be observed when upgrading to version 3.0 firmware:

1. Due to the size of the new SCP firmware image, it cannot be downloaded in a switch running SCP V2.10 or earlier. **SCP V2.20 or later must be running in order to load SCP V3.0 firmware**.
2. The SCP firmware that supports hunt groups is not compatible with pre version 3.0 linecard firmware. In order to avoid compatibility problems, **linecards must be upgraded to V3.0 before upgrading the SCP**. If this is not done, all linecards not upgraded will come up in a "FW Rev Mismatch" state when slot configuration is displayed in OBM. Ports in such linecards will not be able to participate in bridging; these linecards will be limited to accepting a new firmware upgrade.

If, by mistake, the SCP is upgraded before any of the linecards, the switch will be inaccessible by the management station. To restore bridging capability, use the OBM Extended Option Menu (choice 12 from the Main Menu). Choice 3 from the Extended Options Menu allows you to disable/enable Hunt Group support. If you choose to disable hunt group support the switch will reboot and will function normally, except that hunt groups cannot be configured. With hunt groups disabled ports on linecards running pre-version 3.0 firmware can participate in bridging. After the linecards have been upgraded, the same menu option may be used to restore hunt group capability.

Note

It is not recommended that the GIGAswitch/FDDI System continue operation with the SCP and linecards running incompatible firmware - even with hunt groups disabled.

3. In version 3.0 of SCP code the way that MAC addresses are assigned to GIGAswitch/FDDI ports changes. This should have little observable effect - except at the time the V3 image is first loaded. The ARP cache in the network management station (NMS) will have the GIGAswitch/FDDI port's old MAC address assigned to its IP address. If your NMS maintains a database of MAC addresses, you may have to clear the old contents. For more information on how to clear the contents refer to the documentation for the NMS. If the GIGAswitch /FDDI system is again loaded with pre V3 code the NMS database should be cleared again.

New Features

New OBM Menus

The OBM menus have been changed in this firmware release. They have been simplified and made to follow consistent formats.

OBM functions are performed via a series of menus which appear on the OBM terminal. Each menu has a name and a list of choices, preceded by numbers. The user selects the number of the desired choice.

The following conventions are used:

- A choice which ends with the word "menu" results in an additional menu of choices.
- A choice that begins with the word "Show" results in a display, followed by the same menu (or a subset, in the case of the Main Menu).
- A choice that begins with "Show/set" results in a display, followed by a menu of choices related to that display.
- All other choices result in some action, perhaps preceded by input requests or a warning.

  When multiple input values are requested in a single prompt the values supplied should be separated by spaces.

  When one or more port number is required as one of the inputs, the FPPN numbering scheme should be used in one of the following formats:

  single port (10.4)
  list of ports (10.1,10.2,12.2,14.1)
  range of ports (10.1-12.4)
  "all" - to indicate all ports
- The final choice for all menus (except for the Main Menu) returns to a previous menu—often to the Main Menu. The final choice in the Main Menu ends the OBM session.
- Pressing <CR> causes a previous menu or the main menu to reapear, except as indicated below.
- A display that cannot be shown on a single screen will have "(<cr>=More)" following the prompt. Press <CR> to continue the display.
- Ctrl/D will abort the OBM session at any time.

The first menu presented by the OBM is called the Main Menu. The Main Menu consists of the System, Port, Bridge, MIBs, Extensions and OBM sections.

```
Main Menu

  System      1. Show box configuration
              2. Show/set slot configuration
              3. Show/set system time
              4. Clear management memory
              5. Reboot menu

  Port        6. Show LAN address assignments
              7. Show/set privileged and bootserver ports
              8. Show/set delayed ports
              9. IP menu

  Bridge      10. Bridge menu

  MIBs        11. MIB viewer menu

  Extensions  12. Extended options menu

  OBM         13. OBM menu
              14. End OBM Session

Choice:
```

Choices made from the Main menu provide similar functionality to corresponding choices in the OBM menus in previous versions of GIGAswitch/FDDI firmware. For more information on the new OBM features refer to the *GIGAswitch/FDDI System Out-Of-Band Management (OBM) Guide* - EK-GOBMG-MG. B01

24K Translation Table Size

Version 3.0 allows one to choose a translation table (TT) size of up to almost 24,000 MAC addresses. There are now 4 choices of maximum table size: 3,737, 7,737, 15,737 and 23,993. The desired size is selected using the OBM Bridge Menu, which is choice 10 on the Main Menu.

The choice of translation table size determines the maximum number of MAC addresses the SCP will send to any port. If the firmware running on a linecard does not support the chosen maximum size the SCP will not allow its ports to be brought on line. It will indicate a "FW Rev Mismatch" status in the OBM slot display. This condition can be resolved by either downloading appropriate firmware to the linecard in question or by resetting the Translation Table size to a lower number.

The AGL-2 does not support 24K table size. It will always show a FW Rev Mismatch in the OBM slot display when booted with 24K size in effect. Hence the TT size must be set to a lower number in order to have an active AGL-2 port. The AGL-2+ linecard does support 24K table size.

It is recommended that users employ the smallest table size that will support their network. This will reduce switch processor overhead, and minimize learning delays. The default TT size is 7,737.

Demand Learning

Demand Learning is a feature that reduces the learning activity on certain ports. Ordinarily when a new source address (SA) is seen on a port, the SCP informs every GIGAswitch/FDDI port of the association between that address and the port on which it was seen. Thus the entry takes up space in every port's forwarding table.

When demand learning is enabled the SCP initially only notifies the port on which the address was seen. This is critical, since that port is responsible for aging the address, and it needs to know that packets destined for that address should not be forwarded through the crossbar.

Other ports are not notified until they "need" to be. If a packet arrives at a port, destined for that particular address, this port will send the packet to the SCP for flooding (since, for this port, it is an unknown address). The SCP will then realize that this port has a "need" to know that address. So, in addition to flooding the packet, it will convey the information about the address to this port. The next time this port sees the address it will know which port to send it to.

This feature has the advantage of possibly reducing the number of entries in each port's forwarding table, which conserves space and reduces the overhead of maintaining the table. It has the disadvantage of causing an extra flooding event for each port that sends to each address. If every port will eventually send packets to most addresses, then demand learning saves little, and consumes extra overhead. If most ports only send to a limited number of addresses, then the overhead of this feature may well be worth expending for the improved capacity/performance effects.

It will probably require some analysis, or even experimentation to determine whether this feature is appropriate at a given site.

The demand learning feature is set up using the following MIB object:

```
...gigaversion1
   .gigaBridge.gigaStp
    .gigaStpDemandLearningEnable
```

Set this object to **True** (1) to enable demand learning, and to **False** (2) to disable demand learning. It is set to **False** (2) by default.

AGL-2+ Module

The AGL-2+ module is supported in this release. Operationally, the AGL-2+ has the same features as its predecessor, the AGL-2. It differs only in the following ways:

- The AGL-2+ will support translation table size up to 23,993. Its predecessor only supports 15,737 addresses.
- The AGL-2+ uses new modPHY daughter cards instead of the daughter cards used on the AGL-2. These modPHYs can be installed without removing the module from the chassis. **However they should not be installed while the module is powered on**. Instructions for installation are included with the modPHY.

For more information on the new AGL-2+ features refer to the *GIGAswitch/FDDI System AGL Reference Guide* part number EK-GAGL2-MG. B01.

Hunt Groups

The hunt group feature allows a pair of GIGAswitch/FDDI Systems to communicate over two or more active links. This is accomplished by configuring two or more physical ports as a hunt group. When this is done the switch treats that group of ports as a single bridge port. Spanning tree, learning, aging and filtering all see the hunt group as a single port. A source address seen on one of these ports is associated with the hunt group, not the physical port; spanning tree places the hunt group bridge port, not the individual ports, into forwarding or blocking state; addresses are aged on the entire hunt group, not on an individual port; and filters are applied to the hunt group, not the individual ports.

In Figure 1 a 3-member hunt group joins GS1 to GS2, and a 2-member hunt group joins GS1 to GS3.

Configuring several physical ports as a hunt group offers two advantages:

- It allows a higher rate of traffic flow between two GIGAswitch/FDDI Systems.
- It provides quick failover in the event of a link or port failure.

The following section contains more detailed information about hunt groups.

Figure 1 Hunt Groups

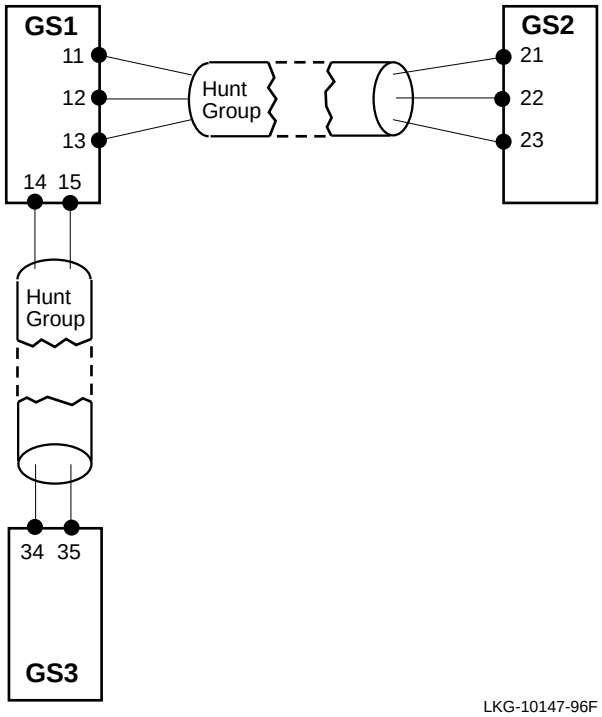

Hunt Groups

Hunt Group Member Ports

A physical port configured in a hunt group should:

- be connected in a point-to-point link.
- run in full duplex mode.
- be connected to a port (on another GIGAswitch/FDDI System) which is also configured as part of a hunt group.

Whenever a physical port is configured in a hunt group the SCP regularly sends proprietary protocol messages through that port. Once the two switches at opposite ends of the link agree that they are at opposite ends of a hunt group the hunt group is established. As additional ports are identified as belonging to the same hunt group, the hunt group is reconfigured. "Hunt group member numbers" are assigned as members join the hunt group. Hunt group member number 1 is assigned to the member with the lowest physical port number, 2 to the next lowest, etc.

Hunt Group Port Numbers

Hunt groups appear to the GIGAswitch/FDDI System as new ports, with different numbers than ordinary ports. There are two numbering schemes used to refer to ordinary ports, SPN (sequential port number) and FPPN (front panel port number). Both of these schemes have been extended to refer to hunt groups as well. The SPN of an ordinary port can range from 1 to 36. The SPNs of hunt groups range from 37 to 64. The FPPN of an ordinary port can range from 1.1 to 14.2. The FPPNs of hunt groups range from 99.37 to 99.64. In hunt groups the FPPN does not have any physical meaning, as it does for ordinary switch ports.

With the addition of hunt group ports it becomes necessary to distinguish the different uses for port numbers. Port numbers refer to both physical and logical ports (or bridge ports). Prior to the existence of hunt groups the physical port number and the logical port number were identical for a given port. For hunt groups that will no longer be the case. Two physical ports may have different media or other characteristics, even though they belong to the same hunt group. The logical port number is used for all bridge operations: spanning tree, learning, aging, filtering. Note that SPNs or FPPNs can **both** be used to refer to either logical or physical ports. The following discussion uses SPNs.

Logical Ports

Every bridge entity in the GIGAswitch (e.g., learning, aging, filtering, and spanning tree process) deals with logical ports. The GIGAswitch/FDDI has 64 logical ports, with SPN 1 to 64. Logical ports are also known as bridge ports. The ports that are actually placed in the box are called physical ports. The switch can access up to 36 physical ports, with SPNs from 1 to 36.

Logical ports can operate only after some physical ports are assigned to them. Since there are more logical ports than physical ports, some logical ports must have no physical port assigned to them. These logical ports are called "empty" ports. Empty ports do not participate in bridge functions.

In the default system configuration, each logical port between 1 and 36 is associated with one and only one physical port, and the mapping from physical ports to logical ports is the identity mapping. Namely, physical port n is assigned to logical port n, where 1 <= n <= 36. A logical port to which only one physical port is assigned is called a singleton bridge port. Logical ports 37 to 64 are empty ports in the default configuration. Hunt groups are created by assigning 1 or more physical ports to logical ports in the range 37-64. Any physical port which is not so assigned retains (or reverts to) its default assignment.

Hunt Group Example

To create a hunt group consisting of physical ports 1, 3 and 5, choose a logical port number by which to refer to this hunt group. The logical port number must be in the range 37-64. Suppose logical port 45 is chosen. Assign physical ports 1, 3 and 5 so they belong to logical port 45 (45:{1,3,5}). When this is done the logical ports 1, 3 and 5 become empty logical ports.

The three physical ports that are now logically assigned to port 45 must next be connected to three ports on a second GIGAswitch/FDDI System. The ports they're connected to must be assigned to a hunt group as well. There is no requirement that the logical port numbers be the same on both switches.

Learning

A packet entering the switch through a physical port configured in a hunt group has its source address (SA) learned on the hunt group logical port. In the above example a packet entering the switch through physical port 1 would have its SA learned on port 45. The forwarding tables in the SCP and on all linecards will list port 45 as the home of that MAC address.

Aging

With several physical links connecting two switches the aging process becomes more complicated. This complexity has been addressed by having the two GIGAswitch/FDDI Systems communicate aging information explicitly. A GIGAswitch/FDDI System will not by itself age out addresses seen on a hunt group port. It will age them out only after the switch at the other end of the hunt group links has aged them out.

This will result in an additional delay in aging, but it will not affect network operation significantly.

Filtering

Filters (both dynamic and management-set) apply to logical ports. It is important to note that previously defined filter matrices (including the default filter matrix) may prevent traffic from traversing a hunt group port, since hunt group ports all have SPN greater than 36. So **be sure to examine, and, if necessary, modify existing filter matrices to be certain they account for new logical ports introduced by hunt groups**.

Single-Path and Multi-Path Packets

When a multi-member hunt group exists there are multiple paths between certain points on the network. Since a packet may traverse any of the hunt group's member links, it has more than one way of going from a station on one side of the hunt group to a station on the other side. Hence there is the possibility that a stream of packets from a given source to a given destination could arrive out of order.

For some protocols and applications out-of-order packet arrival is all right. For others it is unacceptable. The GIGAswitch/FDDI System can differentiate packets as single-path or multi-path, based on incoming physical port and protocol type. In a later section we explain how a network manager can specify which packets are identified as single-path and which as multi-path. As the name suggests, single-path packets from a given source to a given destination will be guaranteed to traverse a single path. Hence they will arrive in order. Multi-path packets may traverse different paths between a source and a destination. Hence they may arrive out of order. In later sections we describe how load balancing is performed for single-path and multi-path packets.

Out-of-Order Packets

While most transport protocols are designed to handle out of order packet delivery, there are some implementations of such protocols that have been observed to fail when significant numbers of packets arrive out of order. Furthermore, there are some protocols, Local Area Transport (LAT) for instance, that are not designed to handle out of order packet delivery at all. The network manager can designate a protocol to be "single-path" in order to adjust to such circumstances.

TCP and Out-of-Order Packets

Some implementations of TCP may be tolerant of, but sensitive to out-of-order packet delivery. One such example is Digital UNIX TCP/IP. This implementation employs a "fast retransmit" algorithm wherein the receiver of an out-of-order packet immediately sends a (duplicate) ACK for the last in-sequence packet received. When the number of duplicate ACKs exceeds a certain threshold, the sender considers a packet to have been dropped, and retransmits, even if the rexmt timer has not yet expired.

The threshold is kept in a kernel variable called **tcprexmtthresh**. The default value for this variable is 3, which means that after receiving 3 duplicate ACKs caused by 3 out of order packets, the sender retransmits. If many such retransmissions occur, application performance or available bandwidth could be adversely affected. It is recommended that the value of this variable be increased for Digital UNIX systems if a large number of retransmissions are observed. Increasing it to 100 should eliminate such unnecessary retransmissions caused by out-of-order packet arrival.

The following commands will accomplish this:

Commands	Comments
% dbx -k /vmunix	Invoke the dbx debugger
(dbx) p tcprexmtthresh 3	Print the current value Current value is 3
(dbx) assign tcprexmtthresh=100 100	Assign new value of 100
(dbx) patch tcprexmtthresh=100 100	Patch image on disk, so value 100 will be used the next time the system reboots

Dynamic Load Balancing

Referring back to our example of a three member hunt group, recall that logical port 45 has three physical ports (1, 3 and 5) assigned to it. When a multi-path packet arrives at the switch, destined for an address seen on port 45, it will be sent to the first of the ports 1, 3, or 5 which is available to take traffic from the crossbar. This results in the most efficient use of the crossbar bandwidth. But, as noted above, it can result in packets being delivered out of order.

Note that actual load balancing need occur only when the crossbar connection to one of the hunt group members is busy. In the absence of crossbar congestion all traffic may flow through a single physical port.

But lack of congestion at the crossbar does not necessarily imply lack of congestion on the outbound links. Dynamic load balancing cannot ensure that the external links are load balanced.

As long as all ports in the hunt group can deliver frames to the outbound link at the crossbar speed (100 Mb/s), there should not be a problem. But a DS3 port, for example, cannot keep up with the crossbar bandwidth. Hence when a DS3 port is a hunt group member, dynamic load-balancing may not work efficiently. It is not recommended that DS3 links be used in hunt groups as a means of increasing bandwidth for multi-path packets. But static load-balancing (described below) works properly for DS3 links.

And DS3 links can be effectively configured as part of a hunt group to provide a redundant data path with enhanced failover capabilities. The reason we require hunt group member ports to operate in Full Duplex mode is to assure they can provide 100 Mb/s output bandwidth.

Static Load Balancing

Static load balancing is applied to single-path packets, and works as follows: Under stable conditions, all single-path packets received on any one GIGAswitch/FDDI inbound physical port are transmitted on the same outbound physical port of any given hunt group. However, single-path packets received on different inbound physical ports may be configured to use different outbound physical ports in a hunt group. In this way the load may be spread over the hunt group members. The network manager controls the mapping of inbound physical ports to outbound hunt group members, as indicated in Figure 2.

Figure 2 Static Load Balancing

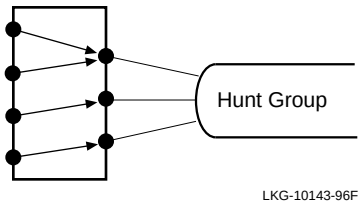

The advantage of static load balancing is that packets are delivered between any source address and any destination address in the order in which they are sent, except in the rare event of network reconfigurations. This contrasts with dynamic load balancing, in which packets are routinely delivered out of order.

The disadvantages of static load balancing are these: First, the load balancing is not as efficient as dynamic load balancing, since the assignment of a packet to a hunt group transmit link does not take into account dynamic port loading. Second, to make the best use of the feature, the network manager must know switch traffic patterns and must configure each hunt group carefully to spread the load evenly over its members.

Finally, when hunt group member links come up or down there is a possibility that a few packets may be delivered out of order. In-order packet delivery may be absolutely guaranteed, by use of **fixed** traffic category described below, which eliminates load balancing entirely.

Traffic Groups

When a packet arrives at a GIGAswitch/FDDI port the receiving linecard classifies it as single-path or multi-path, according to its service class and port number. A later section describes how to configure service class and ports to properly classify packets.

The network manager may assign a physical port to one of 16 traffic groups (numbered 1-16). By default, all ports (1-36) are assigned to traffic group 1. A single-path packet entering a port is assigned that port's traffic group number. It is said to belong to that traffic group.

All single-path packets belonging to the same traffic group and destined to the same hunt group will exit the switch through the same physical port. For each hunt group the network manager designates a traffic category and a hunt group member number per traffic group. (Recall that the physical links of a hunt group are numbered from 1 to the current hunt group size by "hunt group member number".)

The designated traffic category and hunt group member number determine the physical out bound link as follows:

- If the traffic category is **reconfig**, the packet is transmitted on the link which corresponds to the designated hunt group member number.
- If the traffic category is **fixed**, the packet is always transmitted on one selected hunt group link. A different link is selected only if the currently selected link leaves the hunt group (for example, if it fails). The hunt group member number is ignored in this case.

If a designated hunt group member number, n, exceeds the number of active hunt group members, k, the number (n mod k)[1]. is used instead. Whenever a physical port joins or leaves a hunt group, the numbers assigned to members often change, and the the mapping of traffic groups to members will change accordingly. But recall, that this only applies to **reconfig** traffic category.

[1] n mod k = remainder after dividing n by k

Illustration of Static Load-Balancing

Consider a GIGAswitch/FDDI system configured by the network manager as follows:

- Ports 11, 12 and 13 comprise a hunt group with bridge port number 41 (41:{11,12,13}).
- Ports 21,22, 23 and 24 comprise a hunt group with bridge port number 42 (42:{21,22,23, 24}).

Figure 3 Hunt Groups 41 and 42

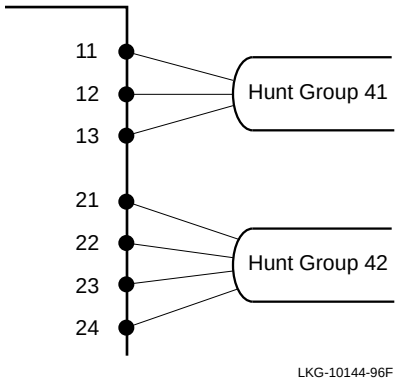

- The memberships of traffic groups 1 through 6 are configured according to the following table. Traffic groups 7-16 have no ports.

Traffic Group	Ports
1	1, 2, 3, 4, 5, 6
2	7, 8, 9, 10, 11, 12
3	13, 14, 15, 16, 17, 18
4	19, 20, 21, 22, 23, 24
5	25, 26, 27, 28, 29, 30
6	31, 32, 33, 34, 35, 36

Figure 4 Traffic Groups

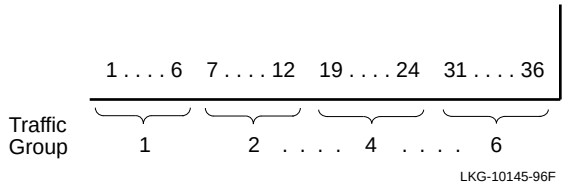

- Traffic groups are mapped by the network manager to member numbers in hunt group logical ports 41 and 42 as follows:

Traffic Group	Hunt Group 41 Member Number	Hunt group 42 Member number
1	1	1
2	2	2
3	3	2
4	1	3
5	2	3
6	4	4

Packets from traffic group 1 leave through hunt group member 1 of both hunt groups, 41 and 42.

Packets from traffic group 2 leave both hunt groups through member 2.

Packets from traffic group 3 leave hunt group 41 through member 3, and leave hunt group 42 through member 2.

Packets from traffic group 4 leave 41 through member 1, and leave 42 through member 3.

Packets from traffic group 5 leave 41 through member 2, and leave 42 through member 3.

Packets from traffic group 6 leave 41 and 42 through member 4. But hunt group 41 has only 3 members. So this traffic group leaves hunt group 41 through member 1, which is (4 mod 3).

Figure 5 Static Load Balancing on Hunt Group 41

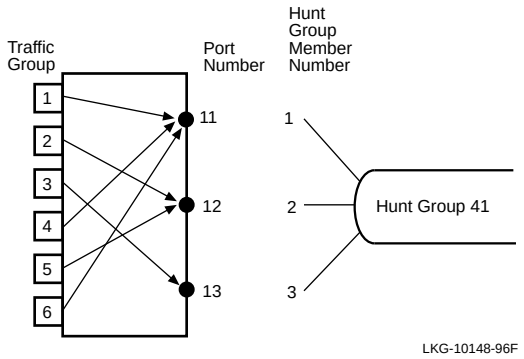

Single-path traffic leaving on bridge port 41:

Traffic Group	Inbound Physical Ports	Hunt Group Member Number	Physical Port Number
1,4,6	1-6 19-24 31-36	1	11
2,5	7-12 25-30	2	12
3	13-18	3	13

Figure 6 Static Load Balancing on Hunt Group 42

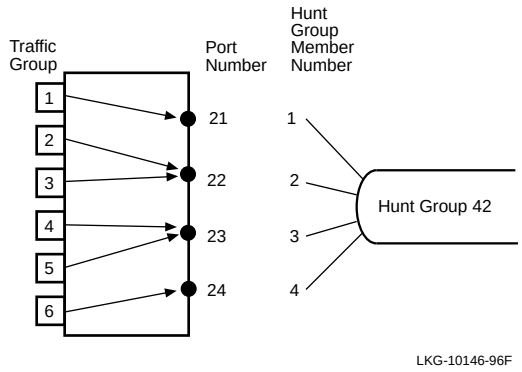

Single-path traffic leaving on bridge port 42:

Traffic Group	Inbound Physical Ports	Hunt Group Member Number	Physical Port Number
1	1-6	1	21
2,3	7-12 13-18	2	22
4,5	19-24 25-30	3	23
6	31-36	4	24

If physical port 22 were to fail, there would be only three hunt group members in hunt group 42 (ports 21, 23, and 24). Port 23 would become member number 2, and port 24 would become member number 3. All the traffic that had been going to member number 4 would now be handled by member number 1 (4 mod 3), which is port 21. The new traffic pattern would be as follows:

Figure 7 New Static Load Balancing on Hunt Group 42

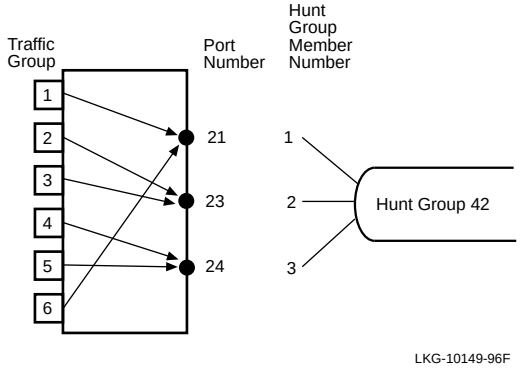

Single-path traffic leaving on bridge port 42 after failure of port 2:

Traffic Group	Inbound Physical Ports		Hunt Group Member Number	Physical Port Number
1,6	1-6	31-36	1	21
2,3	7-12	13-18	2	23
4,5	19-24	25-30	3	24

Hunt Group MIB Objects

An updated revision of the GIGAswitch/FDDI MIB is required to manage hunt groups. The following objects, part of the **gigaSets** branch of the MIB, are used to manage hunt groups.

portGroupMembershipTable

portGroupBridgePort
portGroupMembership
portGroupMembershipWorkBuf
portGroupPortType
portGroupPortTypeWorkBuf
portGroupPortOperStatus

portGroupMembershipFppnTable

portGroupFppnBridgePort
portGroupFppnMembership
portGroupFppnMembershipWorkBuf
portGroupFppnPortType
portGroupFppnPortTypeWorkBuf
portGroupFppnPortOperStatus

portGroupStatusTable

portGroupStatusBridgePort
portGroupStatusPortNumber
portGroupStatusPortType
portGroupStatusOperStatus

portGroupAction

portGroupActionStatus

The **portGroupMembershipTable** is used to configure hunt groups. It is indexed by bridge port numbers, in the range 37-64. The **portGroupMembership** field contains a specification of the physical ports assigned to the indexed hunt group. The **portGroupPortType** field should always have value **huntGroup** (1).

The **portGroupMembership** objects are not directly settable. They are read-only. Instead the **portGroupMembershipWorkBuf** objects are set using SNMP. When all the **WorkBuf** objects have the desired values, set the **portGroupAction** to **Update** (2), to cause the changes in **WorkBuf** fields to take effect in a single operation.

The following status indicators can be used to find out various states of each port:

- **portGroupPortOperStatus** indicates whether a hunt group port is initialized and works as a logical port.
- **portGroupStatusTable** provides various information about physical ports such as hunt group memberships and operational states.
- **portGroupActionStatus** indicates the status of the last **portGroupAction** operation.

Hunt Group Configuration Example

This section demonstrates how to use these MIB objects to create the hunt group of the earlier example: hunt group logical port number 45 has 3 members, physical ports 1, 3 and 5. The hunt group is set up in steps 1 through 4 below. It is torn down in step 5.

Note: To perform the same operations using FPPN notation for ports, use the **portGroupMembershipFppnTable**.

1. Set **portGroupMembershipWorkBuf** to a string describing the members.

```
set portGroupMembershipWorkBuf_45     (1,3,5)
```

The instance variable of the object is logical port number 45. It is the same as the hunt group port number. The string is of the form (<list>), where <list> is a port list specification. A port list specification is a sequence of port numbers (spn) or port ranges, separated by commas. Specifying a null list, (), removes all ports from the hunt group.

2. Get **portGroupMembershipWorkBuf** object and see that the value is correctly set. If not, repeat step (1).

```
get portGroupMembershipWorkBuf_45
```

(Object **portGroupMembership_45** has the value that is currently effective. This will not show any change until the next step.)

3. Set **portGroupAction** to **doUpdate** (2) to make the set effective. After the successful completion of this SNMP set operation, the system will execute the following actions:

   - Copy **portGroupMembershipWorkBuf_45** to **portGroupMembership_45**
   - Permanently record the changes in the system management memory.
   - Initialize hunt group port 45 (i.e., update physical ports 1, 3 and 5 as members of logical port 45).

   As a side effect, logical ports 1, 3, and 5 become empty ports.

   Note: Do not set **portGroupAction** to **doUpdate** until all hunt groups have been specified as desired. The above copy operation will performed on all indices for which **portGroupMembershipWorkBuf** is different than **portGroupMembership**.

4. Get **portGroupMembership** and see if it is identical with its corresponding work buffer.

```
get portGroupMembership_45
```

   Hunt group 45 has now been created.

5. To tear down hunt group port 45, the following steps must be taken:

```
set portGroupMembershipWorkBuf_45     ()
    portGroupAction doUpdate
```

After the successful completion of these 2 sets, hunt group 45 has been torn down. Logical port 45 becomes an empty port and logical ports 1, 3 and 5 become singleton bridge ports.

MIB Objects for Static Load Balancing

The following objects, in the **gigaSets** branch of the MIB, are used to configure static load balancing on hunt groups:

trafficGroupMembershipTable

trafficGroupNumber
trafficGroupMembership

trafficGroupAttributeTable

trafficGroupNum
trafficGroupHgNumber
trafficGroupHgMember
trafficGroupCategory

trafficGroupMembershipTable is a table which specifies the physical ports that are assigned to each of the 16 traffic groups. It is indexed by the **trafficGroupNumber**, a number in the range 1-16. Each entry has an object, **trafficGroupMembership**, whose value is a string that specifies the physical ports assigned to that traffic group. The string is of the form (<list>), where <list> is a sequence of port numbers or ranges of port numbers, separated by commas. A port number is a number between 1 and 36 inclusive. A range is of the form a-b, where a and b are port numbers, and a < b. This range specifies all numbers between and including a and b. Thus (1-3,5,7-36) specifies every physical port except ports 4 and 6. When a traffic group is specified any physical port included in the specification is removed from its previous traffic group. Any physical port that is not explicitly defined to be in a traffic group defaults (or reverts) to traffic group 1.

trafficGroupAttributeTable is a table that specifies how traffic groups are allocated to hunt group members. It is indexed by **trafficGroupNum** (a number in the range 1-16), and **huntGroupHgNumber**, a number in the range 37-64. Each entry has two associated objects:

trafficGroupCategory, which has one of the following values:

fixed - causes all traffic to exit the switch through a single hunt group member, which changes only if the current one leaves the hunt group. This will guarentee in-order packet delivery in all cases, but provides no load balancing.

reconfig - assigns traffic to hunt group members, as specified by the **trafficGroupHgMember**. Traffic is spread among hunt group members, and in-order delivery is guaranteed, except during a hunt group reconfiguration.

trafficGroupHgMember, which specifies the hunt group member number of the port which will transmit packets from this traffic group. This is a number in the range 1-16.

Static Load Balancing Example

Setting up the load balancing for hunt group 41, with member ports 11, 12, 13 are described in the examples starting on page 15.

First we define traffic groups 1, 2, 3, 4, 5 and 6.

```
set trafficGroupMembership_1     (1-6)
    trafficGroupMembership_2     (7-12)
    trafficGroupMembership_3     (13-18)
    trafficGroupMembership_4     (19-24)
    trafficGroupMembership_5     (25-30)
    trafficGroupMembership_6     (31-36)
```

Next, in the **trafficGroupAttributeTable**, for hunt group 41, and each traffic group, 1, 2, 3, 4, 5 and 6, set the value of **trafficGroupCategory** to be **reconfig**.

```
set trafficGroupCategory_1.41     reconfig
    trafficGroupCategory_2.41     reconfig
    trafficGroupCategory_3.41     reconfig
    trafficGroupCategory_4.41     reconfig
    trafficGroupCategory_5.41     reconfig
    trafficGroupCategory_6.41     reconfig
```

Finally assign the values of **trafficGroupHgMember** according to the table on page 16.

```
set trafficGroupHgMember_1.41     1
    trafficGroupHgMember_2.41     2
    trafficGroupHgMember_3.41     3
    trafficGroupHgMember_4.41     1
    trafficGroupHgMember_5.41     2
    trafficGroupHgMember_6.41     4
```

MIB Objects to Specify Single-path Packets

The following MIB objects are used to specify which packets are to be classified as single-path packets. Packets are classified as single-path or multi-path according to the protocol type. So a packet's SAP or SNAP value is the key to its classification. Recall single-path packets are the packets which are governed by static load balancing across hunt group member ports. All of these objects are part of the **.ServiceClassAssignments** branch of the MIB.

ebrNportSnapSvcTable

ebrNportSnapSvc
ebrNportSnapSinglePath
ebrNportSnapSvcStatus

ebrNportSapSvcTable

ebrNportSapSvc
ebrNportSapSinglePath
ebrNportSapSvcStatus

These tables are indexed by SNAP and SAP values, respectively. A particular protocol may be established as a single-path protocol by setting the value of its **ebrNportSapSvc** or **ebrNportSnapSvc** to be 0. Setting the value to 1 makes this a multi-path protocol. The GIGAswitch/FDDI System comes with certain protocols preconfigured as multi-path. They are: IP, IPX, NISCA, ARP, DECnet Phase IV, ISO CNLS. All other protocols are, by default, single-path. Use the above objects to change any of these default settings.

The **ebrNportSnapSinglePath** and **ebrNportSapSinglePath** fields of the table may be set with a port list. On these ports the corresponding protocol will be treated as single-path, even though it may be defined (using the **...SvcTable**) as a multi-path protocol. In short, this is a way to override a multi-path designation on a select list of ports.

Single-Path Protocol Example

To change the service class of a protocol to be different from the default values one must create an entry in the **erbNportSnapSrvTable** MIB object. This table is indexed by SNAP value. For example, the SNAP value for IP is 00-00-00-08-00. So the entry corresponding to IP will be indexed by 0.0.0.8.0. By default IP is set as a multi-path protocol. If one wishes to set IP to be a single-path protocol, set the MIB object:

```
...gigaBridge
   .ServiceClassAssignments
    .ebrNportSnapSvcTable
     .ebrNportSnapSvcEntry
      .ebrNportSnapSvc.0.0.0.8.0
```

to have value 0 (for single-path). Then set the status of this entry to be "permanent" by setting the MIB object:

```
...gigaBridge
   .ServiceClassAssignments
    .ebrNportSnapSvcTable
     .ebrNportSnapSvcEntry
      .ebrNportSnapSvcStatus.0.0.0.8.0
```

to be 2 (permanent).

Setting the status of an entry to be "invalid" (value = 1) will invalidate that entry. The service class of the corresponding protocol will revert to the default setting.

To set the service class of a protocol to be "multi-path" use the value 1, in place of 0, above.

Problems Resolved in v3.0

The following problems, seen in previous releases, are fixed in v3.0, as described below.

AGL-2 Fast Firmware Download

The time required for firmware download of an AGL-2 module has been reduced from 15-20 minutes to just over 2 minutes. The new download process requires that the latest versions of firmware be running on both SCP and on AGL-2. Hence the time savings will not be observed until the succeeding firmware version is loaded. The released firmware for AGL-2+ includes the faster download process, but field test and pre-released firmware versions do not.

Full Duplex Operation

In previous versions a port would occassionally drop out of full duplex mode. Such drops will no longer occur.

Flooding Performance

A shared buffer scheme and other optimizations have been added to improve flooding performance.

FGL-4 in FGL-2 Slot

Fixed a problem which caused an FGL-4 to crash when placed in an FGL-2 slot.

ifOperStatus

This SNMP object now reports the proper value.

Backup SCP

Fixed problem which caused SCP to to go through inititialization twice when it became a backup SCP.

Short Entry Age

Fixed bug whereby SCP would sometimes not switch to the short entry time when required (eg, if a link goes down).

Flooding Counters

Changed flooding counters type from INTEGER to COUNTER, to prevent it from appearing to be a negative number.

M-port Failover

Speeded up the wrongSA scan to reduce the time required for M-port failover.

Linecard Error Logs

Linecard error logs now only display valid entries, and the entries include the firmware version.

PMD Leds

When an FDDI port is disabled via **snmpFddiPortAction** MIB object the left PMD led will now blink green to indicate the port has been disabled by management.

IP Packet Reassembly

IP packets directed to the GIGAswitch/FDDI System are now reassembled correctly.

Downline Load Indicator

While the FGL-4 FLASH is being written during a firmware download the module led will blink green.

FGL Port Indexing

Several problems related to the reporting of port information in snmp have been corrected.

Forward Delay

The value of **ForwardDelay** was reported by SNMP as twice the actual value. This is now reported correctly.

Spanning Tree Parameters Change

If the spanning tree root's parameters were changed during a preforwarding delay, the change did not take effect until after the preforwarding delay. Now it properly takes effect immediately.

Improper TCN

Topology change notifications would occasionally be sent at inappropriate times. This no longer occurs.

Cut Through

Cut through is now properly enabled and reported on all ports.

Forwarding State of M-ports

Once an M-port enters the forwarding state it will not leave it - except during a firmware download or when disabled by the SCP.

Known Problems and Restrictions

Except for improper BPDU problem, the rest of this section pertains to known problems and restrictions having to do with AGL.

Improper BPDU

When spanning tree parameters are changed improper BPDUs may be transmited.

OAM Disabled for AGL

This MIB object **aglInterfaceATMOAMStatus** is permanently set to "disabled" in this release. It cannot be set to "enabled".

Non Zero VPI

Non zero values at VPI are not supported in this release.

STS3-c vs STM1 Mode

If the two ends of a SONET link are set to different mode values, the PHY LED on the STM1 side will blink green indicating that the PHY is down and the link is not functional. In addition, some or all of the following counters will increment on the STM1 side:

aglsonetSectionCurrentESs
aglsonetSectionCurrentSESs
aglsonetSectionCurrentSEFSs

aglsonetLineCurrentESs
aglsonetLineCurrentSESs
aglsonetLineCurrentUASs

aglsonetPathCurrentESs
aglsonetPathCurrentSESs
aglsonetPathCurrentUASs

The STS3-c side of the link may appear to be normal except that the following counters may increment:

- **aglsonetLineCurrentUASs**
- **aglsonetPathCurrentUASs**

Counters

The counters for AAL-5 CRC errors and ATM Header Error Check (HEC) errors are not accessible in this release.

The counters for per VCC cells transmitted are not accessible in this release.

UNI 3.0

The current release of AGL does not support the ILMI portion of the UNI 3.0 specification.

Traffic Shaping

Traffic shaping is not supported in this release.