GIGAswitch/FDDI System

Special Features Guide

part number: EK–GGGSF–UM .B01

This document contains information for managing a GIGAswitch/FDDI System.

Revision/Update Information: This is a revised guide.

Digital Equipment Corporation
Maynard, Massachusetts

Second Printing, July 1995

The postpaid Reader's Comments forms at the end of this document request your critical evaluation to assist in preparing future documentation.

This document was prepared using VAX DOCUMENT, Version 2.1.

Contents

4 ARP Server

5 Other Features

Glossary of GIGAswitch/FDDI Terms

Figures

Tables

Preface

Intended Audience

This guide is intended for network managers who will manage a GIGAswitch/FDDI System in an extended local area network (LAN).

Structure of This Guide

This guide describes several special features of the GIGAswitch/FDDI System. The following table shows where to find information:

Refer to:	For Information About:
Chapter 1	FDDI Features
Chapter 2	Hunt Groups
Chapter 3	ARP Server
Chapter 4	Rate Limiting and Cut-through Forwarding
Glossary	Definitions of GIGAswitch/FDDI System terms

Additional Resources

The following guides provide additional information:

- *Bridge and Extended LAN Reference Manual*, EK–DEBAM–HR
- *FDDI Single-mode Fiber (SMF) modPMD*, AV–QK1PA–TE
- *Configuring the SNMP Agent*, AA–PR84A–TE
- RFC1157: SNMP Standard[1]
- RFC1285: FDDI MIB
- RFC1286: Bridge MIB
- RFC1213: MIB-II
- RFC1407: DS3 MIB
- DEC ELAN Vendor MIB Version 2.7
- GIGAswitch/FDDI MIB Version 1.3

[1] RFCs can be obtained from NIC.DDN.MIL in the rfc directory using anonymous ftp.

- IEEE 802.1d MAC Bridging Standard
- FDDI Standards (ISO 9314-1, 9314-2, 9314-3, 9314-4)
- ATM MIB - (Included in GIGAswitch/FDDI MIB)
- SONET MIB - (Included in GIGAswitch/FDDI MIB)

Documentation for your network management station (NMS) should also be available for regular use.

Conventions

The following conventions are used in this guide:

Bold typeface	A word or phrase is being emphasized. It also indicates a MIB object.
Italic typeface	The complete titles of manuals.
Return	You press the return key on the keyboard.
Ctrl/O	You must hold down the key labeled Ctrl while you press another key or a pointing device button.

1

FDDI Features

The GIGAswitch/FDDI System supports several features directly related to Fiber Distributed Data Interface (FDDI). Such features are applicable to ports on FGL-2 and FGL-4 linecards only, since only those linecards connect the crossbar to FDDI links. These FDDI features are supported:

- FDDI Full Duplex Technology (FFDT)
- Ring Purger
- M-Ports

They are described in the following sections.

FDDI Full Duplex Technology

Ordinarily, stations on an FDDI ring originate data on the ring only when they "own" the ring's token. Data is passed in the "downstream" direction around the ring. Each station retransmits the data it receives. When a station completes its turn originating data it passes the token to its "downstream" neighbor.

Each station can simultaneously send packets on its downstream link and receive packets on its upstream link. But it will only originate packets when it holds the token. That is how the ring's bandwidth is shared by the ring's stations.

FDDI Full Duplex Technology (FFDT) is a proprietary protocol developed by Digital Equipment Corporation to enable both stations on a two-station ring to originate data on the ring at the same time. If a station with FFDT capability notices that its upstream neighbor is the same as its downstream neighbor (that is, there are only two stations on the ring) it will initiate a FFDT communication with the other ring member. Assuming both stations support the FFDT protocol, they will agree to suppress sending a token, and to originate packets on the ring at any time. If a third station joins the ring, normal token passing resumes.

FFDT enables you to theoretically double the bandwidth of a point-to-point FDDI connection over the normal token mode.

Full-duplex operation can be controlled on a per port basis with the **eFDXEnable** MIB object.

The **eFDXEnable** MIB object is indexed by two parameters, MAC # and SMT #. For the GIGAswitch/FDDI System these will always have the same value, the SPN of the port in question. Setting this object to true (1) allows the port to participate in the FFDT protocol. Setting the object to false (2) will inhibit full-duplex operation. Full-duplex operation is ENABLED by default.

The FFDT protocol is licensed by Digital Equipment Corporation to other vendors who wish to support this enhanced functionality.

Ring Purger

Ring Purger is a mechanism developed by Digital Equipment Corporation to ensure that packets are removed from an FDDI ring. Ordinarily, a sending station notices when packets it originates in the downstream direction reappear from upstream.

When the station identifies itself as the originator of a packet, it purges it from the ring (that is, it refrains from retransmitting it). But occasionally a station will fail to purge its own packets. For example, if a station is removed before it purges all its packets those packets may continue on the ring.

Digital Equipment Corporation has developed a technique to periodically purge such unwanted packets from the ring. This facility is called Ring Purger. One of the stations that support this facility becomes the official ring purger, and periodically performs this function. GIGAswitch/FDDI ports support this functionality.

Ring Purger operation can be controlled on a per-port basis with the **eMACRingPurgerEnable** MIB object.

The **eMACRingPurgerEnable** MIB object is indexed by two parameters, MAC # and SMT #. For the GIGAswitch/FDDI System these will always have the same value, the SPN of the port in question. Setting this object to true (1) allows the port to participate in ring purger operation. Setting the object to false (2) will prevent the port from becoming a ring purger. Ring Purger is DISABLED by default.

Dual Homing

Dual homing is an FDDI configuration technique that provides a standby path in case a primary connection fails. This technique permits the A and B ports of a Dual Attachment Station (DAS) or a dual attachment concentrator (DAC) to be attached to two specially configured ports called M-ports. Only one of these M-ports will be active. Usually, the DAS will choose to activate the M-port connected to its B port by default. It will activate the other M-port if the original connection fails for any reason. When both links are operational the A link is in standby mode and runs a link confidence test continuously.

A failover from B to A occurs when the B link ceases to be operational. A failover from A to B occurs when the B link resumes operational status.

A DAS or DAC may be dual homed to one or two GIGAswitch/FDDI Systems. Any FGL port may be configured to be an M-port.

Configuring M–Ports

To configure M-ports, modify the MIB object **dec ema gigabox lineCard mPortTable mPortEnable**. Set it to **true** to create an M-port; or set it to **false** to revert to an ordinary SAS or DAS port.

The following is an example of setting and showing M-port status:

```
set snmp <entity> dec ema sysobjid bridges gigaswitch gigaversion1
gigabox - lineCard mPortTable (36,36) mPortEnable = true {or false}

show snmp <entity> dec ema sysobjid bridges gigaswitch gigaversion1
gigabox - lineCard mPortTable (36,36) all characteristics
```

The **mPortEnable** MIB object is indexed by two parameters MAC # and SMT # (which in the above example is 36,36). For the GIGAswitch/FDDI System these will always have the same value, the SPN of the port in question. Setting this object to true (1) creates an M-port. Setting the object to false (2) reverts to a SAS or DAS port.

When a port changes from SAS or DAS to M-port enabled or from M-port enabled to SAS or DAS the PHY Status LED[1] flashes green for about 1 second during which time the port stops forwarding packets. In this transitional period some frames may be lost. Once a bridge port has M-port enabled, and after a brief listening period, the corresponding PORT LED indicates the forwarding state (solid green)—whether or not an active physical FDDI connection is present.

[1] Formerly called PMD PHY LED

When the M-port connections become active, the PHY Status LED of the GIGAswitch/FDDI M-port connected to the A port will be flashing alternately amber/green. This indicates the M to A connection is in standby mode. At the same time, the PHY Status LED of the GIGAswitch/FDDI M-port connected to the B port is solid green, indicating this FDDI connection is currently active.

Failover

Dual homing provides a powerful failover capability for redundant GIGAswitch/FDDI configurations. It is helpful to understand how the failover process affects the operation of the switch. Failover from one M-Port to another can be caused by:

- Failure of the active M-Port itself (perhaps caused by a linecard failure or even a failure of the switch).

  or

- Failure of the link connected to the active M-Port

Failure of the station dual homed to the M-Port would not result in a failover, since that station is connected to both M-Ports (the active one and the standby one). The next sections describe the ways failover may affect switch operation: address learning and spanning tree reconfiguration.

Address Learning

While an M-port is active it learns the addresses of packets seen on that port. When failover to the standby M-Port occurs these addresses will no longer be seen on the original port, and will begin showing up on another port.

Consider the fate of a packet whose destination address was previously learned on a port that has become inactive due to an M-Port failover. The packet will either:

- traverse the crossbar with no delay—if (a) that address has been seen and relearned on another port.

  or

- be flooded—if that address has not yet been learned on a new port, but has (b) been purged from the switch’s forwarding table.

  or

- be dropped—if the address has not been either (a) relearned or (b) purged.

Waiting for (a) or (b) to occur is part of the delay associated with M-Port failover. Relearning occurs as addresses are seen on new ports. Purging can occur in two ways:

1. If the failure is outside the switch, then the port will age out addresses according to the value of the aging time.

2. If the port itself failed, then the switch will initiate a "cleanup" process which will remove all that port's addresses from the switch's forwarding tables. It is important to note that until this process is complete that port cannot be brought back on line.

The time required for relearning or purging depends on the number of addresses in the forwarding database, and the amount of traffic.

Spanning Tree

Dual homing failover may also trigger a spanning tree change, which would cause a port to enter a non forwarding state for a period of time which depends on the spanning tree parameters of the given network.

A spanning tree change would occur if the failover caused the GIGAswitch/FDDI System to see the spanning tree root bridge on a different port.

Dual Homed Configurations

Three types of configuration are discussed below and depicted in Figure 1–1. Each configuration is shown with two DASs, dual homed to GIGAswitch/FDDI System M-ports. The M-ports in the figure appear on different FDDI linecards. If these linecards are on different GIGAswitch/FDDI there is typically a link connecting these systems.

The configurations are:

1. Each DAS dual homed to a single (DAS) M-Port.
2. Each DAS dual homed to two (DAS) M-Ports. Both DASs to the same two ports.
3. Each DAS dual homed to two M-Ports. Each DAS to its own pair of ports.

Configuration 1 has nominal redundancy. It guards only against failure of a link. It can only be configured on DAS M-ports (on FGL-2 linecards).

Configuration 2 is a compromise. It provides redundancy for all failures, but with reduced bandwidth. In the case of a failure the two DASs will share a ring. But it only requires 2 bridge ports, compared to four in configuration 3. This configuration requires DAS M-ports (on FGL-2 linecards).

Configuration 3 uses two bridge ports for each DAS. It is fully redundant. If one active link (or port) fails the standby will become active, with no loss of bandwidth. It can be configured with ports on FGL-2 or FGL-4 linecards.

Figure 1–1 Dual-Homing Configurations

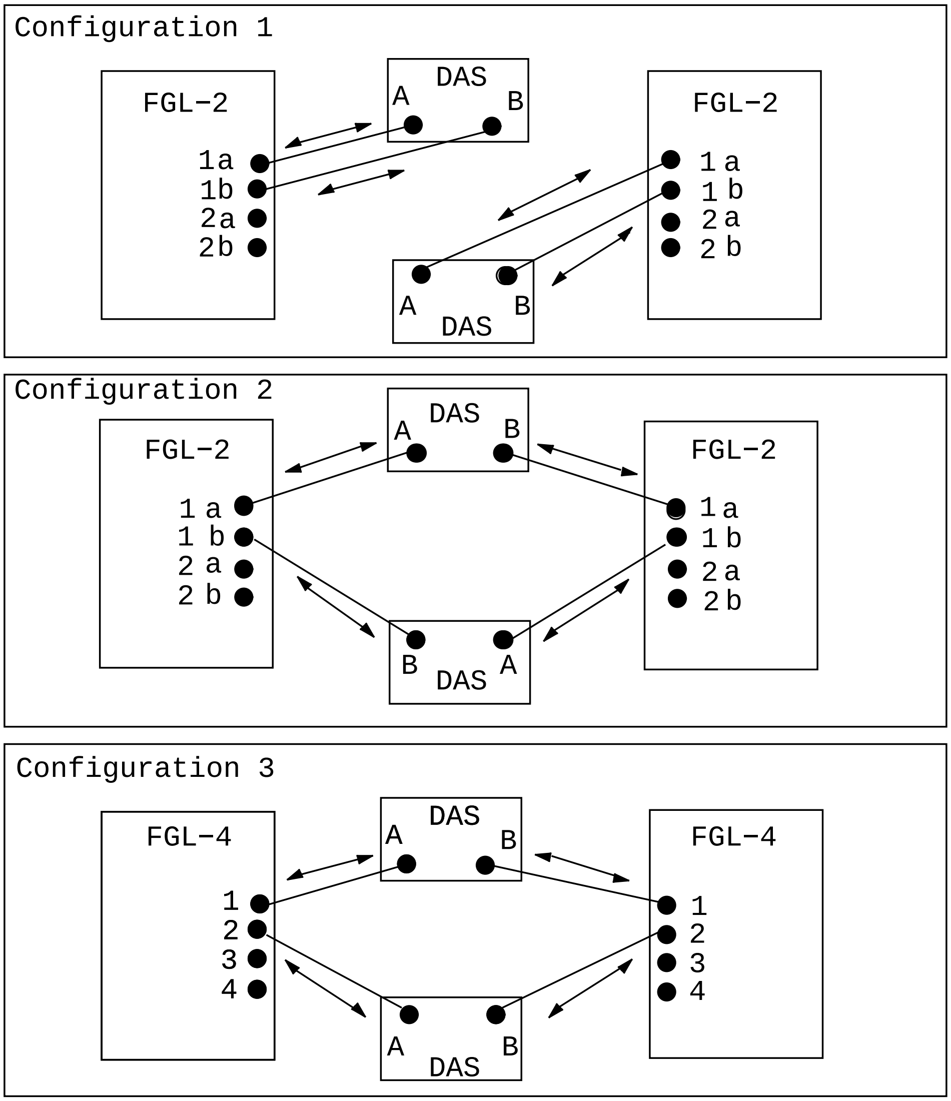

2
Hunt Groups

Hunt Group Overview

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.

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.

A physical port configured in a hunt group should:

- be connected to 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 is assigned to the member with the lowest physical port number, 2 to the next lowest, ect.

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. 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 between 37 to 64 are, in the default configuration, empty ports. 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 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

Commands	Comments
(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, and 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 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.

Figure 2–1 Static Load Balancing

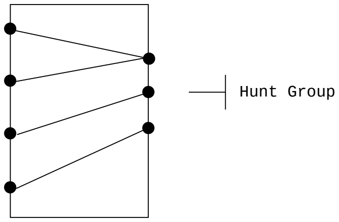

The advantage of static load balancing is that packets are delivered between any source address and any destination address in the order in which they are sent, except in the rare event of network reconfigurations. This contrasts with dynamic load balancing, in which packets are routinely delivered out 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 **static** 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 all hunt group links are numbered (in the order in which they join the hunt group) from 1 to the current hunt group size. This "hunt group member number" thus specifies a particular hunt group link.

The 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 **static**, 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 2–2 Hunt Groups 41 and 42

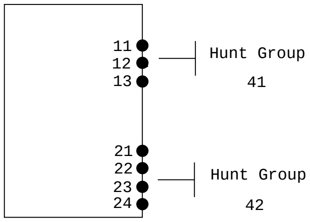

- The memberships of traffic groups 1 through 6 are configured according to the following table. Traffic groups 7-16 have no ports.

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 2–3 Traffic Groups

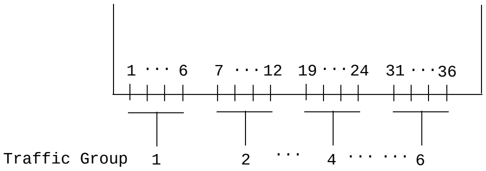

- Traffic groups are mapped by the network manager to member numbers in hunt group logical ports 41 and 42 as follows:

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 2–4 Static Load Balancing on Hunt Groups 41 and 42

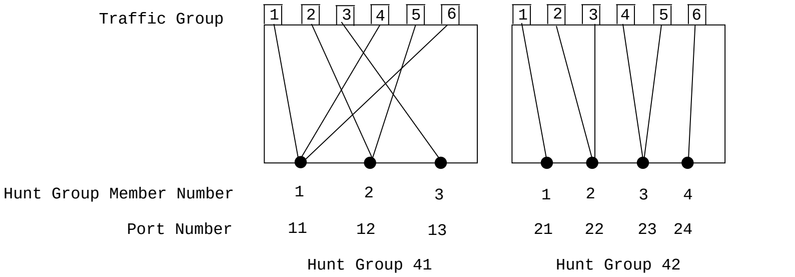

Single-path traffic leaving on bridge port 41:

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

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 on 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:

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** is not directly settable. Instead the corresponding **WorkBuf** fields are set using SNMP. When all the **WorkBuf** fields 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.

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. A port range is of the form a-b, where a and b are port numbers and a < b.

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 either of the objects in **workBuf** are different than the corresponding objects in **MembershipTable**.

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

   ```
   get portGroupMembership_45
   ```

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

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

After the successful completion of these steps, 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
huntGroupNumber
trafficGroupMemberNumber
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 **huntGroupNumber**, a number in the range 37-64. Each entry has two associated objects:

> **trafficGroupCategory**, which has one of the following values:
>
> > **static** - 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 **trafficGroupMemberNumber**. Traffic is spread among hunt group members, and in-order delivery is guaranteed, except during a hunt group reconfiguration.
>
> **trafficGroupMemberNumber**, 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

We show how to set up the load balancing described above example for hunt group 41, with member ports 11, 12, 13.

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

```
set trafficGroupNumber_1     (1-6)
    trafficGroupNumber_2     (7-12)
    trafficGroupNumber_3     (13-18)
    trafficGroupNumber_4     (19-24)
    trafficGroupNumber_5     (25-30)
    trafficGroupNumber_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 **trafficGroup_MemberNumber** according to the table on page 13.

```
set trafficGroup_Member_Number_1.41    1
    trafficGroup_Member_Number_2.41    2
    trafficGroup_Member_Number_3.41    3
    trafficGroup_Member_Number_4.41    1
    trafficGroup_Member_Number_5.41    2
    trafficGroup_Member_Number_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.

3

Learning Domains

Domains Overview

Domains are used to keep certain switch activities separated on the basis of port. A domain is defined as a set of logical ports. Domains must be disjoint; a (logical) port can only be in one domain. There are three types of domains relevant to GIGAswitch/FDDI Systems: Traffic Domains, Learning Domains, and Spanning Tree Domains. These domains are specified differently, but both serve to keep a specific type of activity focused on a subset of the (logical) ports of the GIGAswitch/FDDI System.

Learning Domains

A learning domain is a set of (logical) ports which learn the same set of addresses. At a minimum all ports in a learning domain will learn all addresses seen on any of the domain's member ports. There can be up to eight (8) learning domains, numbered 1 to 8.

- Address Learning

  One of the functions of a bridge, such as GIGAswitch/FDDI, is to note and retain an association between network hardware addresses and ports. This enables frames with a known destination address (DA) to be forwarded directly to the proper port and to no others. This is how a bridge segregates traffic, allowing a high agregate bandwidth on the extended LAN. The activity of noting and retaining these associations is called "learning". It is done by observing the source address (SA) of each incomming frame, and associating it with the port on which it arrives.

  The GIGAswitch/FDDI attains its high throughput with a distributed architecture. The decision concerning the proper destination port of a frame is made at the incomming port. This avoids the need for all frames to pass through a single "sorting" process. This distributed architecture requires that all learning data be transferred to all the ports.

- Restricted Learning

  Learning domains serve to limit the extent of the learning process. Any address that is seen on a port in a learning domain is distributed to all other ports in that learning domain, but to no other ports. It is even possible for the same address to be seen on two separate ports, as long as they are in separate learning domains. While this is unusual it is not impossible, and the GIGAswitch/FDDI System will deal with it correctly. It will associate that address with one port in one learning domain, and with another port in the other learning domain. A learning domain is specified by listing the logical ports to be included. Any logical port that is not included explicitly in a learning domain will be placed in learning domain 1, the default learning domain..

Spanning Tree Domains

A spanning tree domain is a set of (logical) ports over which a spanning tree process operates. There can be up to eight (8) spanning tree processes, each operating on its own spanning tree domain, in a GIGASwitch/FDDI System. Initially all (logical) ports (1-64) are part of spanning tree domain 1. Because of the close connection between the operation of spanning tree domains and learning domains we do not allow them to be defined independently. A spanning tree domain is specified as a union of learning domains

4

ARP Server

This chapter describes the ARP Server functionality available with the GIGAswitch/FDDI System.

Normal ARP Operation

Hosts in an IP subnet must keep track of the mapping between IP network addresses and MAC hardware addresses of other nodes in the subnet. Once an association of IP-MAC address is learned it is kept in an ARP cache for a period of time.

When a host requires the MAC address for an IP address that is not currently in its cache it sends an ARP broadcast request, which ordinarily reaches all nodes in its subnet. The host with the target IP address sees this request and replies with a unicast ARP response, which identifies its IP-MAC address association.

In many IP implementations, when the IP system boots, it sends an ARP reply broadcast message containing its IP-MAC address pair. This information may be placed in the ARP cache of each participating IP host.

Purpose of ARP Server

As the number of IP hosts on the network increases, the ARP broadcast traffic can overload the LAN. When the GIGAswitch/FDDI ARP server is enabled the SCP answers the ARP request broadcast messages when it has the desired information, rather than flooding the ARP request to all its ports. This cuts down on ARP broadcast traffic.

ARP Cache Contents

The GIGAswitch/FDDI ARP server cache only contains IP address mapping information from ARP messages processed by the SCP. These messages include broadcast messages (both request and reply), and ARP unicast messages that are addressed to the SCP.

ARP Server Operation

If the ARP server is enabled:

1. The GIGAswitch/FDDI System floods all non-IP ARP broadcast requests, as permitted by existing filters.
2. The IP ARP broadcast packets are handled as follows:
   - Requests received over GIGAswitch/FDDI System ports with no IP address assigned are neither forwarded nor answered.

- If the IP-MAC translation is unknown, then the request is flooded to all ports except the incoming port.
- If the target MAC address is known to be on the same port as the requestor, then the GIGAswitch/FDDI System neither forwards nor answers the request.
- If the target MAC address is known to be on a different port than the requestor, then the GIGAswitch/FDDI System answers the request.
- If the target MAC address is not known to be on any port (for example aged out), then the request is flooded to all ports except the incoming port—even if the IP-MAC address translation is known.

Dropping ARP Requests

Packets which are in error, such as ARP requests to IP broadcast addresses, are discarded by the GIGAswitch/FDDI System. These are counted in the **arpFramesDiscarded** MIB object.

Enabling ARP Server

The GIGAswitch/FDDI ARP Server may be enabled by the NMS using the SNMP object **arpAgent**. Once enabled the GIGAswitch/FDDI ARP Server remains enabled, even after switch reboot, until disabled by NMS.

Monitoring ARP Server

The maximum size of the ARP server cache is the same as that of the translation table (SNMP object **ttSize**). Contents of the GIGAswitch/FDDI ARP server cache may be determined by using the SNMP object **netToMediaTable**.

Note

Since the GIGAswitch/FDDI ARP server cache is being continually updated even between SNMP GET messages, the information in the table may appear inconsistent when comparing one SNMP GET message to another.

Updating ARP Cache

Each entry in the ARP server is updated periodically. When the timer (**arpTimeoutInSeconds**) expires, the SCP sends a unicast message to the hardware address of the entry expiring in the cache. The SCP waits **arpPeriodBetweenRequests** and sends the request again if no reply was received. The SCP repeats the request the number of times specified by **arpRequestRetries**. If it still receives no answer the entry is removed from the cache. The ARP cache **is not** maintained across reboots or failovers.

Default BL2.1 values

Table 4–1 provides the default GIGAswitch BL2.1 values for the SNMP objects for the GIGAswitch/FDDI ARP server.

Table 4–1 ARP Server SNMP Object Defaults

Object	Default
arpAgent	False
arpPeriodBetweenRequests	1 second
arpRequestRetries	2 retries
arpTimoutInSeconds	600 seconds

5

Other Features

This chapter describes four additional features of the GIGAswitch/FDDI System: Rate Limiting, Cut-Through Forwarding, 24K Translation Table and Demand learning.

Rate Limiting

Since the GIGAswitch/FDDI System is a multiport IEEE 802.1d bridge, it floods all multicast and unknown destination address packets to all ports that do not filter such packets. If there are a lot of packets being flooded at any given time they can affect performance on the LANs connected to the GIGAswitch/FDDI System. The GIGAswitch/FDDI System limits the rate of multicast traffic (which includes broadcast frames, and unicast frames with unknown destination addresses) for the entire switch.

The GIGAswitch/FDDI System is designed to give all connected LANs an equitable share of its resources, and to prevent any "badly-behaved" LAN segment from consuming more than its fair share of throughput resources. Any LAN segment with a high rate of multicast traffic will not consume an unfair amount of the multicast rate limit, since the GIGAswitch/FDDI System guarantees a per-port lower bound on the multicast flooding bandwidth.

Two frame rates are controlled. The first is for multicast and broadcast frames. The second is for frames that are flooded because their DA was unknown. Both numbers are measured by a period and a count (in kilobytes per second). The default for both is 300 kilobytes/second. This default is designed to satisfy a vast majority of network configuration requirements.

Before raising the rate limit value, consider the capacity of the lower bandwidth links in your extended LAN. A higher rate limit value will allow a higher rate of traffic on these LANs, and this may overload a lower capacity LAN. Conversely, if a lower capacity LAN is being saturated by traffic from the FDDI network, you may consider lowering the rate limit value.

The GIGAswitch MIB objects that affect rate limiting are listed in Table 5–1.

Table 5–1 Rate Limiting MIB objects

Group	Object	Definition
gigaBridge	**floodUnknownUnicastRate**	Maximum bytes-per-second bandwidth of packets with unknown DAs.
gigaBridge	**floodMulticastRate**	Maximum bytes-per-second bandwidth of packets with multicast DAs.

There are some MIB objects in the DEC ELAN Vendor MIB that concern rate limiting, but do *not* apply to the GIGAswitch /FDDI System. These objects are **ebrRateLimitSwitch**, and **ebrRateLimit**.

Cut-through Forwarding

The GIGAswitch/FDDI System employs a cut-through forwarding technique that makes it possible for a port to begin transmitting before the entire packet has arrived (provided the outbound port is not busy). On an inbound port transmission to the outbound port (via the crossbar) began as soon as the destination address (DA), source address (SA), and the protocol id have been seen—assuming the outbound port is not busy.

At the outbound port transmission to the datalink can begin as soon as the DA, SA, and protocol id have been seen—assuming the link is available. This provides higher throughput than standard store-and-forward techniques. Cut-through forwarding can be enabled/disabled at the inbound or outbound port. The default state is enabled.

Cut-through is controlled by MIB objects in the GIGAswitch MIB. These objects include: **cutThroughTable**, **cutThroughEntry**, **cutThroughTable**, **cutThroughInbound**, **cutThroughOutbound**, and **cutThroughBridgePort**.

24K Translation Table Overview

The translation table (TT) size can handle up to almost 24K MAC addresses. There are 4 choices of the maximum table size: 3,737, 7,737, 15,737 and 23,993. The desired size is selected from the Bridge Menu, reached via choice 10 in 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 will not support 24K table size. It will always show on the OBM slot display a FW Rev Mismatch when booted with 24K size in effect. Hence the 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 the consequences of certain network anomalies.

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.

Glossary of GIGAswitch/FDDI Terms

Address Resolution Protocol

See ARP.

agent

In the client-server model, the part of the system that prepares and exchanges information on behalf of a client or server application.

alarm

A message sent to operator terminals that are enabled or defined by management software. Alarms are set using the network management station (NMS). *See also NMS.*

American National Standards Institute

See ANSI.

ANSI

American National Standards Institute. A national standards organization with members from computer manufacturers and users in the United States. It is the U.S. member body of ISO and is involved with the development of standards around the OSI Reference Model. ANSI proposes, compiles, and publishes standards for programming languages, databases, telecommunications, and other products.

ARP

Address Resolution Protocol. A protocol that maps a high-level Internet address with a low-level physical hardware address. Limited to networks that support hardware broadcast.

backup

A network device or circuit that is used if the primary device or circuit becomes unavailable. The spanning tree algorithm can put bridges or network branches in backup mode if they are redundant with others and might create loops in the network. *See also spanning tree.*

BOOTP

Boot protocol. A protocol that determines a diskless host Internet address at startup, so that the host can operate in an Internet network. *See also protocol.*

BPDU

An IEEE 802.1d Bridge Protocol Data Unit.

broadcast

Simultaneous transmission of data to more than one destination in a network, so that all broadcast addresses receive the same message.

CBS

Crossbar switch. The switching module that forms the heart of the GIGAswitch/FDDI System.

community name

SNMP password for primitive security. *See also SNMP.*

crossbar switch

See CBS.

cutthrough

A process that enables the GIGAswitch/FDDI System to start forwarding a packet out of a port before the entire packet is received. Inbound cutthrough begins packet transmission through the crossbar switch before it is fully received from the inbound port. Outbound cutthrough begins packet transmission on the outbound port before it is fully received from the crossbar switch.

DA

Destination Address. A unique network address identifying a target system. For filter purposes, this is the 48-bit MAC address. Packets are filtered based on the destination address of the packet.

destination address

See DA.

DAS

Dual Attachment Station. An FDDI station that offers two connections to the dual counter-rotating ring. *See also FDDI.*

dotted-decimal notation

The representation for a 32-bit integer that consists of four 8-bit numbers written in base 10 with decimals separating them. This is used to represent IP addresses on the Internet.

dual attachment station

See DAS.

dual homing

An FDDI method of cabling concentrators and stations that enables an alternate or backup path to the ring if the primary connection fails. *See also FDDI.*

FEU

Front end unit. Power supply for the GIGAswitch/FDDI System.

FDDI

Fiber Distributed Data Interface. A set of ANSI/ISO standards that define a high-bandwidth (100 megabits per second), general-purpose LAN connection between computers and peripheral equipment in a timed-token passing, dual ring of trees configuration.

Fiber Distributed Data Interface

See FDDI.

FGL-2

Fiber GIGAswitch/FDDI Linecard, 2-port.

FGL-4

Fiber GIGAswitch/ FDDI Linecard, 4-port. Can configure only as a single attachment station, it cannot be configured as a dual attachment station.

filtering

The process where a bridge evaluates incoming messages and selects those it needs to process, and those which it blocks from delivery. Filters can be set using management station commands.

forwarding

The ability of a bridge or router to accept messages from one local area network (LAN) segment and retransmit those messages to another LAN segment. *See also LAN.*

frame

A data transmission unit containing data or control information, address information, and a frame check sequence.

front-end unit

See FEU.

full-duplex

Pertaining to a type of data communications system capable of providing simultaneous, independent transmission and reception in both directions.

get

An SNMP request for data command. *See also SNMP.*

get-next

An SNMP command that gets the next data item in the MIB object tree. *See also SNMP, MIB.*

hotswap

The ability to remove and insert a component without powering down the GIGAswitch/FDDI System. This procedure does not interrupt normal operation.

in-band management

A technique for carrying control signals within the same bandwidth as data being carried. In-band management for the GIGAswitch/FDDI System is performed using a network management station (NMS). *See also NMS*.

Internal Protocol

See IP.

IP

Internet Protocol. The network layer protocol for the Internet protocol suite that provides the basis for the connectionless, best-effort packet delivery service. IP includes the Internet Control Message Protocol (ICMP) as an integral part. The Internet protocol suite is referred to as TCP/IP because IP is one of the two most fundamental protocols.

LAN

Local area network. A self-contained group of computers and communications devices (such as modems, routers, servers, and repeaters) that offer a high-speed, reliable communications channel. LANs span a limited distance, such as a building or cluster of buildings, but can be connected to wide area networks (WANs) with bridges or routers.

learning

The process by which a bridge discovers and remembers which ports network devices are connected to.

local area network

See LAN.

Maintenance Operation Protocol

See MOP.

Management Information Base

See MIB.

Management Station for UNIX

See MSU.

MIB

Management Information Base. A collection of objects that can be accessed with a network management protocol.

MMF

Multimode fiber. Used in FDDI networks to support network station connections up to 2 kilometers

MOP

Maintenance Operations Protocol. A network management protocol within DECnet software that handles tasks such as downline loading, upline dumping, and circuit testing.

multicast

A special form of broadcast transmission where copies of the packet are only delivered to a subset of all destinations.

network management station

See NMS.

NMS

Network management station. The system responsible for managing a network. The NMS talks to network management agents, which reside in the managed nodes, using a network management protocol (such as SNMP). *See also SNMP*.

OBM

Out-of-band management. In network management, a technique for carrying control signals over a separate channel rather than within the main signal bandwidth. Out-of-band management for the GIGAswitch/FDDI System is performed with a local terminal connected directly to the system with an RS-232 cable.

out-of-band management

See OBM.

PDU

Protocol data unit. The data units (messages or blocks of data) passed between peer entities on different open systems. PDUs consist of both Protocol Control Information (PCI) and user data.

Physical media device

See PMD.

PAID

Process identification. A binary value that uniquely identifies a process. Each process has a process identification and a process name.

PID

Protocol ID (*not process ID*).

PM

Presentation module. An interaction method for use with DECmcc.

PMD

Physical layer media dependent. The GIGAswitch/FDDI System supports two types of PMDs: multimode fiber and single mode fiber.

POLYCENTER

POLYCENTER network management software monitors, controls, and tests entities in the DECnet, DECnet/OSI, and multivendor distributed environment. The GIGAswitch/FDDI System can be managed by the POLYCENTER Network Manager (formerly DECmcc) with the SNMP access module, or by the POLYCENTER SNMP Manager (formerly DECmcc Management Station for ULTRIX). *See also SNMP*.

port

An individual connector on the GIGAswitch/FDDI System that connects a LAN to the GIGAswitch/FDDI System. *See also LAN*.

presentation Module

See PM.

privileged port

A port that can perform SNMP set operations on a secure GIGAswitch/FDDI System. Privileged ports are defined by network management. *See also SNMP*.

process identification

See PI.

protocol

A set of rules for the implementation of a network communication system. Protocols cover options such as signaling methods, coding, packaging of messages, and methods of preventing and correcting errors.

protocol data unit

See PDU.

PSA

Power system assembly.

PSC

Power system controller.

rate limiting

Limits imposed on multicast traffic and traffic with unknown destination addresses (DAs). This reduces the risk of overloading ports with traffic. *See also DA*.

SA

Source address. The unique network address indicating the originator of a message.

Simple Network Managing Protocol

See SNMP.

SAP

Service access point. The point at which an entity provides a service to a user entity in the layer above it. The SAP is named according to the layer providing the services (transport services are provided at a transport SAP, or TSAP, at the top of the transport layer).

SAS

Single attachment station. An FDDI station that offers one S port for attachment to the FDDI ring, usually via a concentrator. *See also FDDI.*

SCP

Switch control processor.

service access point

See SAP.

SET

An SNMP command that can set (alter) an SNMP object. *See also SNMP.*

single attachment station

See SAS.

slot

A groove where a module or card can be installed.

SMF

Single Mode Fiber. Used in FDDI networks to support network station connections up to 40K.

SNAP

Subnetwork Access Protocol. Used in protocol ID PID filtering. *See also PID.*

SNMP

Simple network management protocol. A protocol for monitoring and controlling hosts, bridges, routers, and terminal servers on TCP/IP networks with network management applications, such as DECmcc.

source address

See SA.

spanning tree

The logical arrangement created by bridges in an extended LAN in which all LANs are connected and there are no loops. *See also LAN.*

Subnet mask

Address mask. A bit mask used to select bits from an Internet address for subnet addressing. The mask is 32 bits long and selects the network portion of the IP address and one or more bits of the local portion.

TFTP

Trivial file transfer protocol. An Internet facility for transferring electronic files in a TCP/IP environment. TFTP allows authorized users to transfer files over the network.

transparent bridging

The IEEE 802.1d bridging scheme used to interconnect LANs based upon a spanning tree algorithm. The bridge provides all necessary functionality, including address learning and address filtering. Transparent bridging is protocol independent, performs automatic learning and forwarding, and ensures a loop-free topology in a large network, as only one active bridge connects any two LANs. *See also LAN, learning, spanning tree.*

trap

An unsolicited SNMP message sent by an SNMP manageable device to one or more network management stations (NMS). *See also NMS, SNMP.*

Trivial File Transfer Protocol

See TFTP.

WAN

Wide area network. Two or more standard or extended LANs that are joined by routers, gateways, or packet-switched interface (PST) software.

Wide Area Network

See WAN.

UTP

Unshielded twisted pair. Used in networks to support network station connections up to 100 meters.