GIGAswitch/FDDI System

Special Features Guide

part number: EK–GGGSF–UM .A01

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

Revision/Update Information: This is a new guide.

Digital Equipment Corporation
Maynard, Massachusetts

First Printing, April 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

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	ARP Server
Chapter 3	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
- IEEE 802.1d MAC Bridging Standard

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

- 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

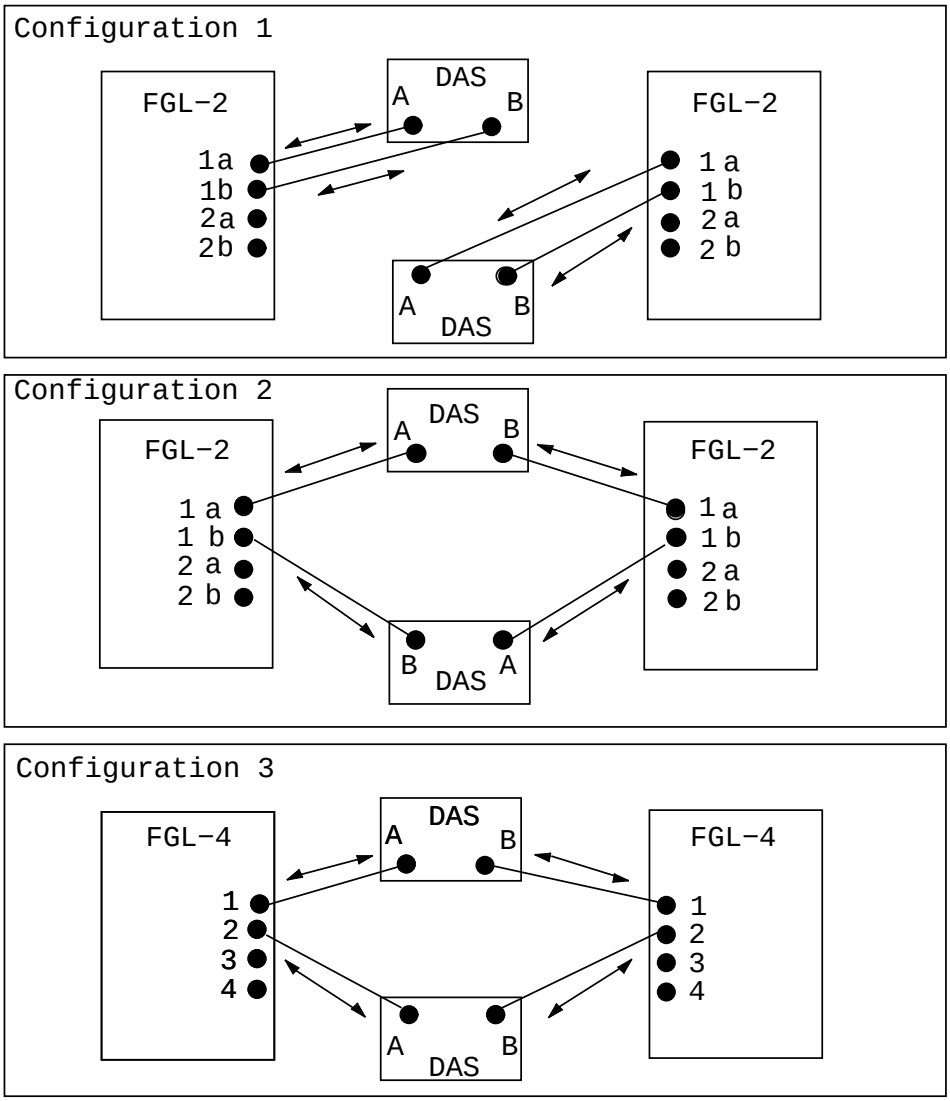

2

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 2–1 provides the default GIGAswitch BL2.1 values for the SNMP objects for the GIGAswitch/FDDI ARP server.

Table 2–1 ARP Server SNMP Object Defaults

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

3

Other Features

This chapter describes two additional features of the GIGAswitch/FDDI System: Rate Limiting and Cut-Through Forwarding.

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 3–1.

Table 3–1 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**.

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.