VNswitch 900 Series

Technical Overview

Part Number: AA-R2LCB-TE

Contents

Preface

1 Switching Overview

2 VNswitch 900 Series Switches

3 Features and Functionality

4 Ethernet to Fast Ethernet Switching

5 Ethernet-to-FDDI Switching

6 Ethernet-to-ATM Switching

7 Switching Between Multiple Technologies

A Accessing MIBs and RFCs

B Spanning Tree

Glossary

Index

Figures

Tables

Preface

Overview

Purpose of This Manual

This manual provides a technical overview of the VNswitch 900 family of high-density switching products. It also includes product specifications and ordering information.

Intended Audience

This manual is intended for use by anyone who needs to become familiar with the VNswitch 900 series products.

Organization

This manual is organized as follows:

Section	Description
Chapter 1	Provides introductory information about switching, including basic switch operation, and introduces the VNswitch family.
Chapter 2	Describes each module in the VNswitch family, and provides product specifications and ordering information.
Chapter 3	Discusses the technical features and functions of the VNswitch modules.
Chapter 4	Provides a brief overview of Fast Ethernet technology and illustrates various configurations of the VNswitch 900EX module in an Ethernet to Fast Ethernet switching environment.
Chapter 5	Provides a brief overview of Fiber Distributed Data Interface (FDDI), discusses the FDDI configuration capabilities of the VNswitch 900EF module, and illustrates various configurations of the module in an Ethernet-to-FDDI switching environment.
Chapter 6	Provides a brief overview of Asynchronous Transfer Mode (ATM) technology and illustrates various configurations of the VNswitch 900EA module in an Ethernet-to-ATM switching environment.
Chapter 7	Describes several possible configurations of the VNswitch modules in multiple-technology environments.
Appendix A	Lists the Management Information Bases (MIBs) supported by the VNswitch family and provides instructions for accessing the MIBs, firmware updates, and release notes.
Appendix B	Provides reference information about the spanning tree algorithm, a process implemented by all VNswitch modules.
Glossary	Provides definitions of terms associated with the VNswitch 900 series.

Conventions

This manual uses the following conventions:

Convention	Description
Special Type	This special type in examples indicates system output.
Boldface	Boldface type in examples indicates user input.
DEChub 900	This term refers to the DEChub 900 MultiSwitch.

Associated Documents

The following documents provide information relating to the module. To order any of the following documents, refer to the directions in How to Order Additional Documentation.

Title and Order Number	Description
VNswitch 900 Series Switch Management AA-R2LD*-TE	Describes how to configure, monitor, and manage a VNswitch 900 series module.
VNswitch 900EE Installation and Configuration EK-DVNEE-IN	Describes the VNswitch 900EE module, including features, installation, and configuration information.
VNswitch 900EX Installation and Configuration EK-DVNEX-IN	Describes the VNswitch 900EX module, including features, installation, and configuration information.
VNswitch 900EF Installation and Configuration EK-DVNEF-IN	Describes the VNswitch 900EF module, including features, installation, and configuration information.
VNswitch 900EA Installation and Configuration EK-DVNEA-IN	Describes the VNswitch 900EA module, including features, installation, and configuration information.
ATM Modular PHY Cards Installation	Provides installation and operating guidelines for installing, verifying, and removing modular PHY cards. Describes cabling and LED information.
DEChub 900 MultiSwitch Owner's Manual EK-DH2MS-OM	Provides installation, use, security, and troubleshooting information for the DEChub 900 MultiSwitch.
DEChub ONE Installation EK-DEHU2-IN	Provides installation and operation guidelines for standalone module configuration, including mounting options and cabling.

Title and Order Number	Description
DEChub ONE-MX Installation EK-DEF1H-IN	Provides installation and operation guidelines for standalone module configuration, including mounting options and cabling.
clearVISN Installation	Provides pre- and post-installation information, as well as actual installation procedures for each application.
clearVISN Overview	Provides an overview of clearVISN, an explanation of each application, and descriptions of all concepts necessary to understand and use the application efficiently.
clearVISN User's Guide	Provides information for starting each application, configuring them, and general use information.
OPEN DECconnect Applications Guide	Provides information to help plan and install networking systems based on Digital's OPEN DECconnect System and networking products.
Event Logging System Messages Guide	Describes messages logged by the Event Logging System.
Bridge and Extended LAN Reference	Describes how bridges are used to create extended local area networks (LANs). The descriptions include the use of bridges in extended LAN configurations, information on LAN interconnections, overall bridge operation, spanning tree, bridge management, and solving bridge-related problems in a network.

Note: The * symbol represents the revision of the manual.

Correspondence

Documentation Comments

If you have comments or suggestions about this document, send them to the Network Products Business Organization.

Attn:	Documentation Project Manager
FAX:	(508) 486-8041
E-MAIL:	doc_quality@lkg.mts.dec.com

Online Services

To locate product-specific information, refer to the following online services:

BBS   To read the Bulletin Board System, set your modem to 8 bits, no parity, 1 stop bit and dial 508-486-5777 (U.S.).

WWW   The Digital Equipment Corporation Network Products Business Home Page on the World Wide Web is located at the following addresses:

North America:	http://www.networks.digital.com
Europe:	http://www.networks.europe.digital.com
Australia:	http://www.digital.com.au/networks

How to Order Additional Documentation

To order additional documentation, use the following information:

To Order:	Contact:
By Telephone	Continental USA: 1-800-DIGITAL (1-800-344-4825)
	Canada: 1-800-267-6215
	Alaska, New Hampshire, and Hawaii: 1-603-884-6660
Electronically (USA only)	Dial 1-800-DEC-DEMO (For assistance, call 1-800-DIGITAL)
By Mail (USA and Puerto Rico)	DIGITAL EQUIPMENT CORPORATION P.O. Box CS2008 Nashua, New Hampshire 03061 (Place prepaid orders from Puerto Rico with the local Digital subsidiary: 809-754-7575)
By Mail (Canada)	DIGITAL EQUIPMENT CORPORATION LTD. 940 Belfast Road Ottawa, Ontario, Canada K1G 4C2 Attn: A&SG Business Manager
Internationally	DIGITAL EQUIPMENT CORPORATION A&SG Business Manager c/o local Digital subsidiary or approved distributor
Internally	U.S. Software Supply Business (SSB) DIGITAL EQUIPMENT CORPORATION 8 Cotton Road Nashua, New Hampshire 03063

Chapter 1

Switching Overview

Overview

Introduction

Bandwidth—the capacity of network media to handle data traffic—has always been a bottleneck in the growth of any network. As a network grows or is bridged with other networks, the problem of reduced bandwidth can become an enterprise-wide problem, slowing the flow of critical information shared between various departments or organizations. The technology of logically organizing a network into different segments and switching data traffic between segments has evolved as a way to achieve maximum available bandwidth—without the necessity of making costly wiring or equipment configuration changes.

This chapter provides an overview of switching technology and explains how the VNswitch family of high-density switches meets the demands of today's enterprise networks.

NOTE

This chapter provides only general information about switching and introduces the VNswitch 900 family of switches. For product-specific information, go to Chapter 2. For detailed information about the VNswitch family features and functionality, go to Chapter 3.

In This Chapter

The following topics are covered in this chapter:

What Is a Switch?

Introduction

In the simplest terms, a switch is a network device that switches data traffic between two or more networks or network segments. Although there are switches for both local area networks (LANs) and wide area networks (WANs), the discussion in this book is focused on LAN switches.

Figure 1-1 illustrates a simple switch connecting two local area networks. Figure 1-2 illustrates a switch being used to segment a local area network into several segments (or individual LANs).

Figure 1-1: Simple Switch Connecting Two LANs

LKG-10180-96F
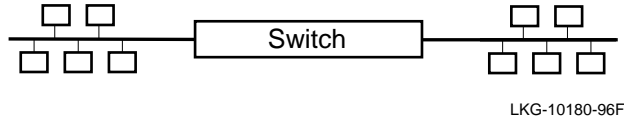
What Is a Switch?

Figure 1-2: Using a Switch to Segment a LAN

BEFORE
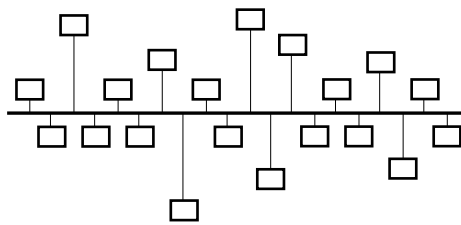
AFTER
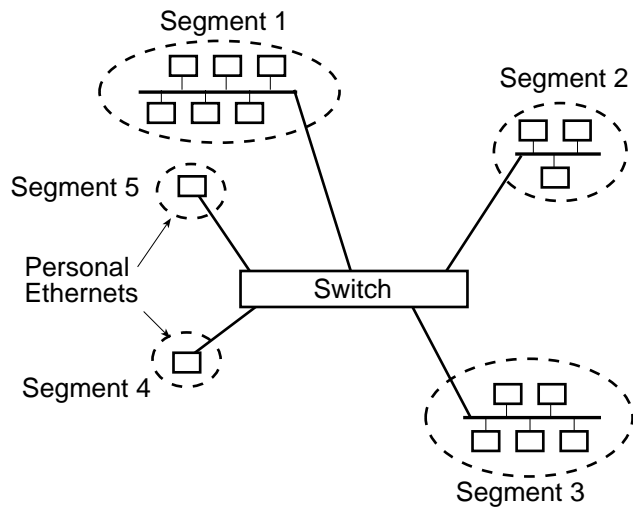
LKG-10182-96F

A switch is sometimes referred to as a multiport bridge, and, in fact, the switch in Figure 1-2 could also be called a multiport bridge. Like a bridge, a switch maintains a table of Layer 2 Media Access Control (MAC) addresses associated with each of its ports. The switch examines the address table to determine how to switch data traffic between the LANs attached to its ports. For example, as shown in Figure 1-3, when a switch (or bridge) receives a data packet from the network attached to port A, it processes the packet differently depending on the packet's destination address:

- If the address is known to exist on port A, the switch filters the packet and keeps the packet local (that is, the packet never leaves the LAN attached to port A), as shown in Example A.

- If the address is known to exist on the LAN attached to port B, the packet is sent out of port B to the node associated with the address, as shown in Example B.

- If the address is an unknown destination address (for example, it does not exist in the address table), the switch floods the packet to all its ports, as shown in Example C.

What Is a Switch?

Figure 1-3: How a Switch Filters Packets

Example A

Example B

Example C

⊞⊞ Data Packet
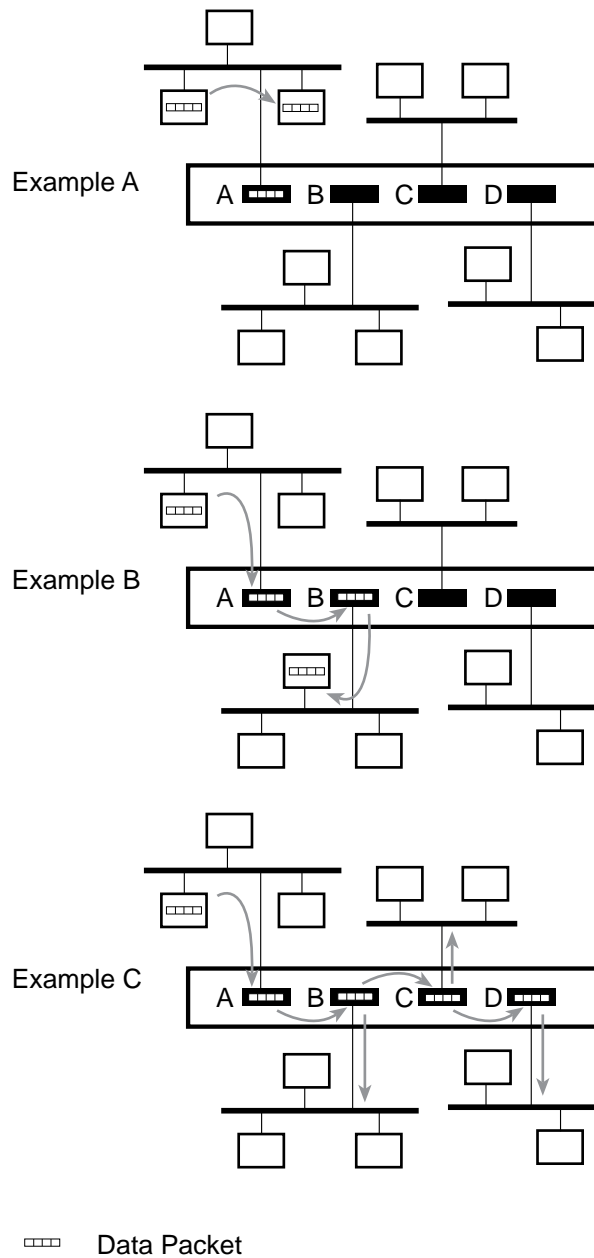
LKG-10293-96F

Types of Switches

A switch is typically classified as either a *frame switch* or a *cell switch*. A frame switch processes and forwards variable-length LAN packets (or frames). A cell switch processes and forwards fixed-length cells.

All the modules in the VNswitch 900 family are classified as frame switches. The VNswitch 900EA high-density Ethernet-to-ATM switch processes *both* LAN frames and ATM cells.

Chapter 3 provides detailed information about the frame formats processed by the VNswitch family of switches.

Features and Benefits of LAN Switches

A switch overcomes the shortcomings of conventional bridges and routers by providing full media bandwidth for each LAN segment. Also, LAN switches can allow the process of segmentation to be extended to its ultimate limit of an individual LAN segment per end station—a dedicated or *personal Ethernet* LAN connection (with full media bandwidth) for each user and server.

The following summarizes the features and benefits of LAN switches:

- Provide dedicated bandwidth for each port

- Increase throughput

- Offer low latency

- Minimize traffic loads by filtering frames that are not destined for another segment

- Implement the spanning tree algorithm, which provides support for multiple paths between segments

- Permit network reconfiguration through software rather than by modifying the network infrastructure

- Provide an economical alternative to expensive routers

Methods of Switching

Introduction

A switch provides the physical connections between users and independent LANs, allowing data packets to be switched from one LAN (or LAN segment) to another. Switching is implemented using either *store-and-forward* or *cut-through* technology.

Store-and-Forward Switching

When a store-and-forward switch processes the data packets passing through it, it validates that the packets have been received without error. As shown in Figure 1-4, a store-and-forward switch waits to receive an entire packet from the network link, verifies packet integrity by examining the frame check sequence (CRC) field, processes its destination address (and optionally, other fields) and, only then, forwards the packet to the appropriate network.

Cut-Through Switching

As shown in Figure 1-5, a cut-through switch, unlike a store-and-forward switch, does not attempt to validate the data in the packets that pass through it. A cut-through switch begins retransmitting a packet as soon as it reads its destination address (if transmit and receive speeds are the same).

Figure 1-4: Store-and-Forward Technology

Stage 1 - Receive
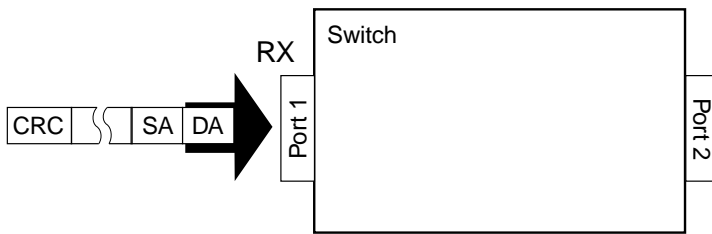
Stage 2 - Receive Complete Frame, Verify Frame Integrity
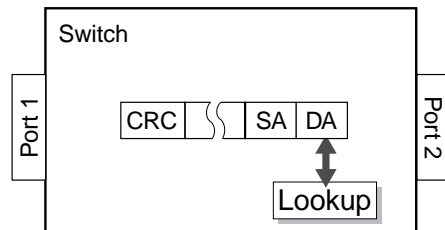
Stage 3 - Forward Frame
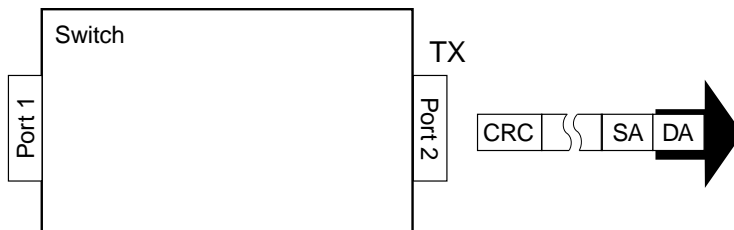
LKG-10105-95I

Figure 1-5: Cut-Through Technology

Stage 1 - Receive
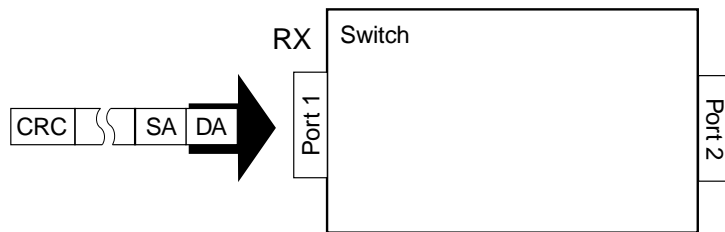
Stage 2 - Receive DA and Lookup
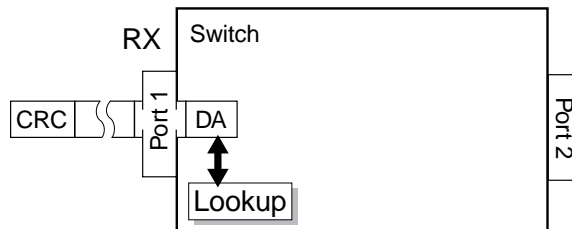
Stage 3 - Forward Packet
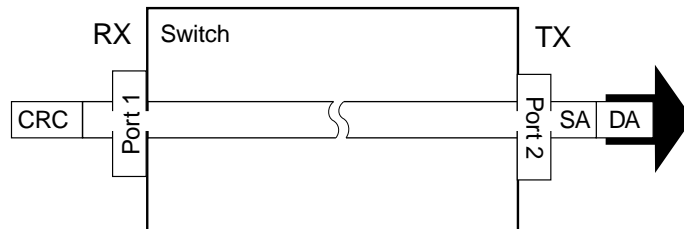
LKG-10106-95

Which Switching Method Is More Effective?

The effectiveness of store-and-forward and cut-through switching methods depends on the network technologies on which they are deployed. This section compares the effectiveness of both switching methods when switches are used to interconnect the following network types:

- Ethernet

- FDDI

- Dissimilar networks, such as Ethernet/FDDI, Ethernet/Fast Ethernet, Ethernet/ATM, or any combination of these

Ethernet

Ethernet (using CSMA/CD) is a collision-oriented approach to transmitting data on a network. If two stations on an Ethernet network attempt to transmit data at the same time, a collision occurs. Collisions can result in packet fragments (or "runt" packets) that do not contain any useful or complete user data.

Runt packets are filtered by store-and-forward switches but forwarded by cut-through switches. This is because cut-through switches only examine the destination address and then forward all packets immediately. When cut-through switches are used to switch Ethernet traffic, the network manager must consider both the integrity of data packets as well as the performance capabilities of the switch.

FDDI

FDDI networks are based on a timed token-passing scheme. Collisions do not occur with this type of network technology, therefore, there are no runt packets. Cut-through switching can be used effectively to improve network performance when forwarding to other FDDI local area networks but are impractical when switching between an FDDI network and a dissimilar network, as described below.

Dissimilar Networks

When switching between different speeds or technologies is involved, cut-through switches are not very useful. For example, differences in speed between FDDI (100 Mb/s) and Ethernet (10 Mb/s), or between *fast* Ethernet (100 Mb/s) or ATM (155 Mb/s) and Ethernet (10 Mb/s), make buffering unavoidable. In addition, switching between FDDI and Ethernet involves translation of frame formats, fragmentation of large packets (if fragmentation is supported), and recalculation of checksums. This typically requires packets to be fully received and stored before forwarding them.

The use of a store-and-forward switch guarantees the integrity of the packets that pass through the switch by detecting bad or incomplete fragments and preventing them from propagating through the extended LAN.

Store-and-forward switches are the only effective solution when switching between dissimilar network technologies. All switches in the VNswitch family implement store-and-forward switching.

Redundancy Feature

IEEE 802.1d-compliant store-and-forward switches (such as the VNswitch family of switches) provide support for redundant paths to ensure network availability.

For example, as shown in Figure 1-6, there are two separate paths between LAN A and LAN C:

- Path 1 through switch 1

- Path 2 through switch 2 to LAN B, then through switch 3

The second path (through switches 2 and 3) is the redundant (backup) path that can be used if switch 1 fails.

Path 2 is not available unless Path 1 becomes unavailable.

Figure 1-6: Redundant Path for Network Availability
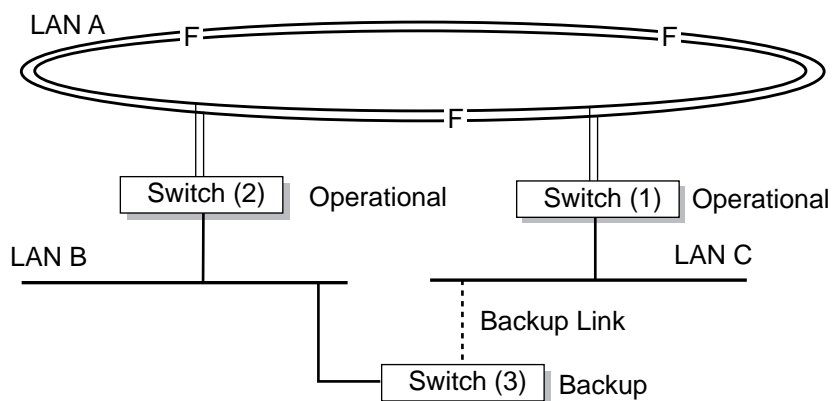
LKG-10118-96F

Network Switching Technologies

Introduction

Switching is not restricted to a particular LAN technology. The VNswitch family includes switches that can be used to connect both similar and dissimilar LANs based on different technologies, including:

- Standard 10 Mb/s Ethernet

- Fast (100 Mb/s) Ethernet

- Fiber Distributed Data Interface (FDDI)

- Asynchronous Transfer Mode (ATM)

Ethernet-to-Ethernet Switching

Standard 10 Mb/s Ethernet switching technology is used to connect standard Ethernet LANs, or to segment a single Ethernet LAN. This technology is employed in small workgroup settings or wherever a 10 Mb/s bandwidth capacity is sufficient to handle network traffic.

Using this technology, a network manager can divide an Ethernet LAN into several segments (or independent LANs). For example, a 4-port switch could interconnect up to four Ethernet LANs, forming one extended LAN. To the users and applications connected to the individual LANs, the four LANs appear as one. However, the switch treats each LAN as an individual LAN segment and ensures that the full 10 Mb/s bandwidth is available to each segment.

Figure 1-7 illustrates a simple 4-port switch used in an Ethernet-to-Ethernet, half-duplex switching environment.

Figure 1-7: Ethernet-to-Ethernet Switching

4 Port 10 Mb/s
Ethernet Switch
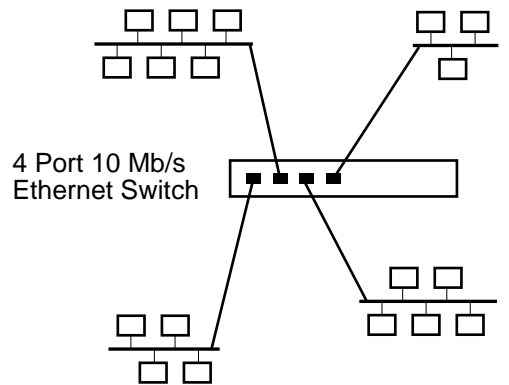
LKG-10181-96F

The VNswitch family includes a 24-port Ethernet switch called the VNswitch 900EE. The VNswitch 900EE is described in detail in the following chapter.

Unlike ordinary Ethernet switches, all switches in the VNswitch family can operate in either half- or full-duplex mode when used in a personal Ethernet configuration. For example, as shown in Figure 1-8, a VNswitch (the VNswitch 900EX Fast Ethernet switch, described later) is used as a backbone connecting 12 personal Ethernet LANs to each other and to a high-speed server. In this configuration, all connections to the switch can operate in full-duplex mode. By allowing bidirectional data transmissions, full-duplex switches can effectively double the throughput of ordinary half-duplex switches.

Figure 1-8: Full-Duplex Switching

Personal Ethernet

Server

100 Mb/s Fast
Ethernet Link

Full Duplex
10 Mb/s Links
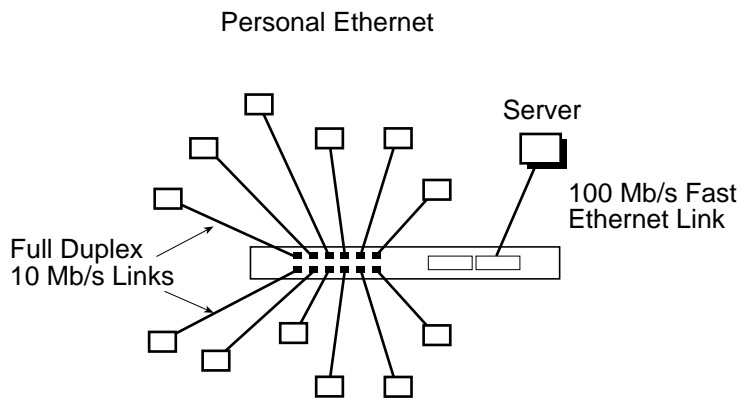
LKG-10295-96F

Ethernet to Fast Ethernet Switching

The 10 Mb/s bandwidth provided by standard Ethernet can become a bottleneck when switching traffic between LANs populated with processing-intensive workstations, servers, and commercial multimedia applications. This can require a medium to large scale network to use a high-speed backbone to interconnect its network segments.

Fast Ethernet technology uses the ISO/IEC 8802-3 CSMA/CD MAC protocol, but implements the protocol at 100 Mb/s. MAC frame formats and sizes are the same as those used by 10 Mb/s Ethernet. Therefore, there is no translation of frames between 10 Mb/s Ethernet and Fast (100 Mb/s) Ethernet.

Figure 1-9 illustrates a switch used in a Fast Ethernet switching environment. In the scenario depicted in Figure 1-9, a Fast Ethernet-capable switch functions as a 100 Mb/s backbone for 12 local area networks. Each LAN (or LAN segment) is provided with the full 10 Mb/s bandwidth offered by the standard Ethernet wire connected to the switch's front port. However, all transfers between the 10 Mb/s LAN segments and the server connected to the Fast Ethernet backbone occur at 100 Mb/s.

Figure 1-9: Simple Fast Ethernet Backbone Switching Environment

LKG-10183-96F
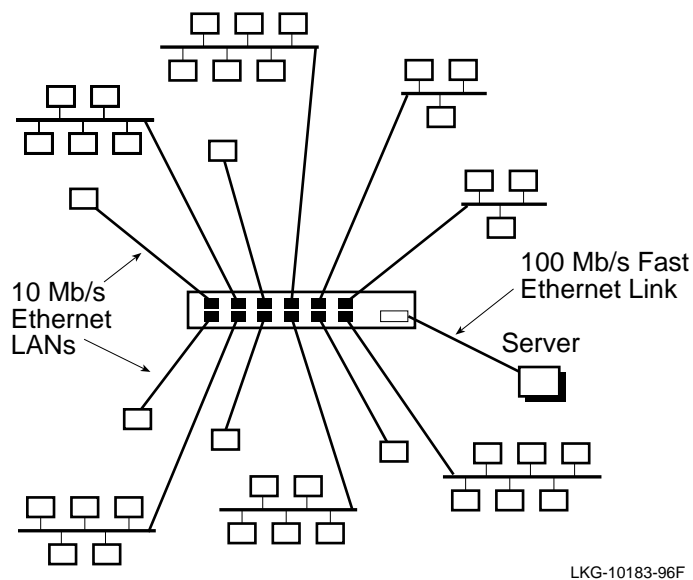
Figure 1-10 illustrates another use of a Fast Ethernet switch. In this illustration, a pair of Fast Ethernet switches connect two distinct networks over a 100 Mb/s Fast Ethernet backbone cable.

Figure 1-10: Two Fast Ethernet Switches Connecting Multiple LANs

100 Mb/s
Fast Ethernet
Link

LKG-10188-96F
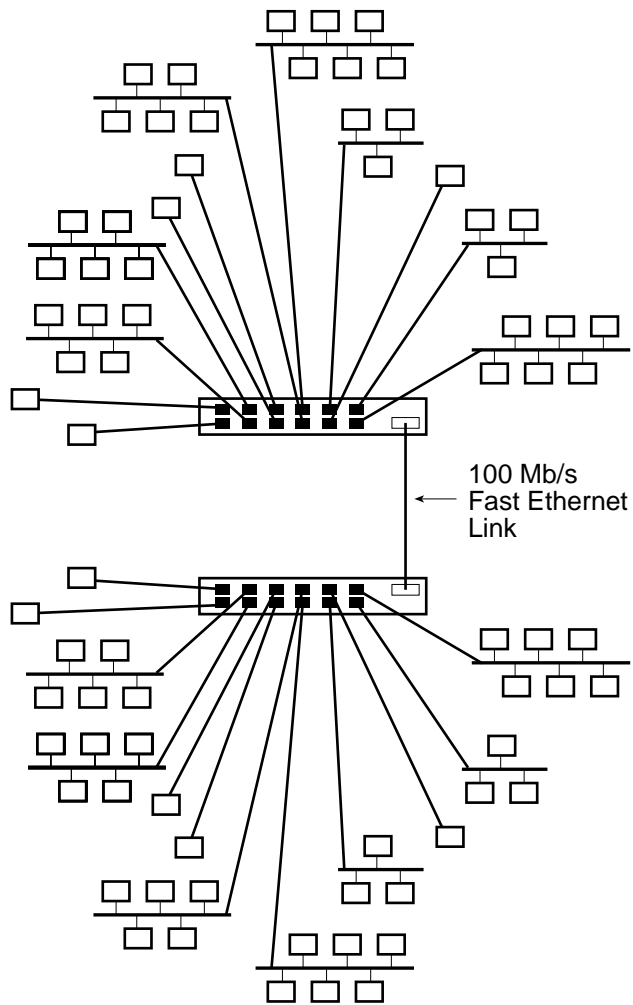
The VNswitch family includes a Fast Ethernet switch called the VNswitch 900EX. The VNswitch 900EX is an Ethernet to Fast Ethernet switch that can connect multiple LANs to two 100 Mb/s Ethernet backbones. The VNswitch 900EX features two Fast Ethernet ports and twelve 10BaseT ports. The VNswitch 900EX is described in detail in Chapter 2. For more information about Fast Ethernet switching, including an overview of Fast Ethernet technology, refer to Chapter 4.

Ethernet-to-FDDI Switching

Fiber Distributed Data Interface (FDDI) is another LAN technology offering high-speed data communications. By collapsing multiple Ethernet segments onto a high-speed FDDI backbone, users can take advantage of the high-speed services offered by FDDI—such as application and file services. Figure 1-11 illustrates an Ethernet-to-FDDI switch used as a high-speed backbone connecting several Ethernet LANs to an FDDI network.

Figure 1-11: Example of Ethernet-to-FDDI Switching
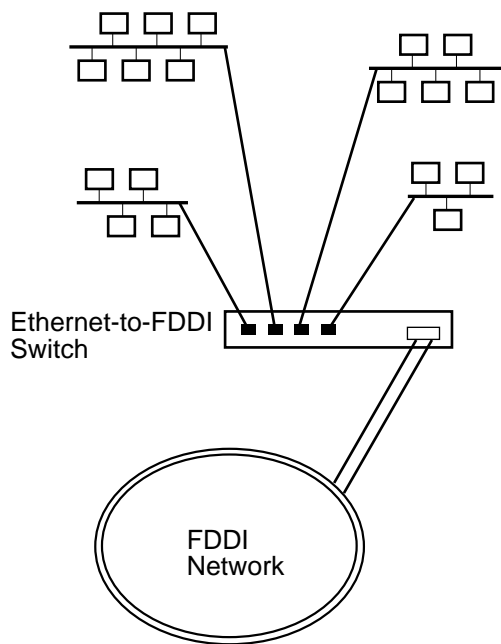
LKG-10187-96F

The VNswitch family includes an Ethernet-to-FDDI switch called the VNswitch 900EF. The VNswitch 900EF is an Ethernet-to-FDDI switch that can connect multiple LANs to an FDDI backbone or network. The VNswitch 900EF features a single-port pair (DAS or SAS) FDDI interface and twelve 10BaseT ports. The VNswitch 900EF is described in detail in Chapter 2. For more information about Ethernet-to-FDDI switching, including an overview of FDDI technology, refer to Chapter 5.

Ethernet-to-ATM Switching

Asynchronous Transfer Mode (ATM) is a high-speed networking technology that switches traffic between nodes on dedicated point-to-point network paths. This switching approach enables ATM to provide link bandwidth capacities ranging in the hundreds of megabits per second, and even gigabits per second.

A significant difference between ATM and other LAN technologies is that ATM requires the setting up of a virtual connection before the end systems can communicate. Traditional LANs (also called *legacy LANs*), such as Ethernet and FDDI, are *connectionless* because all the stations on an Ethernet or FDDI LAN use a shared medium, allowing data to be transmitted without the necessity of first establishing a connection between the end points.

Data is transferred through an ATM network in fixed-length cells, as opposed to the variable length packets used in traditional LANs. Thus, Ethernet-to-ATM switching involves the translating of Ethernet frames to ATM cells, and vice versa.

Figure 1-12 illustrates an Ethernet-to-ATM switch used to connect several LANs across an ATM network.

Figure 1-12: Ethernet-to-ATM Switching Environment

10 Mb/s Ethernet
LANs

Ethernet-to-ATM
Switch

ATM Cloud

Ethernet-to-ATM
Switch

10 Mb/s Ethernet
LANs

LKG-10190-96F
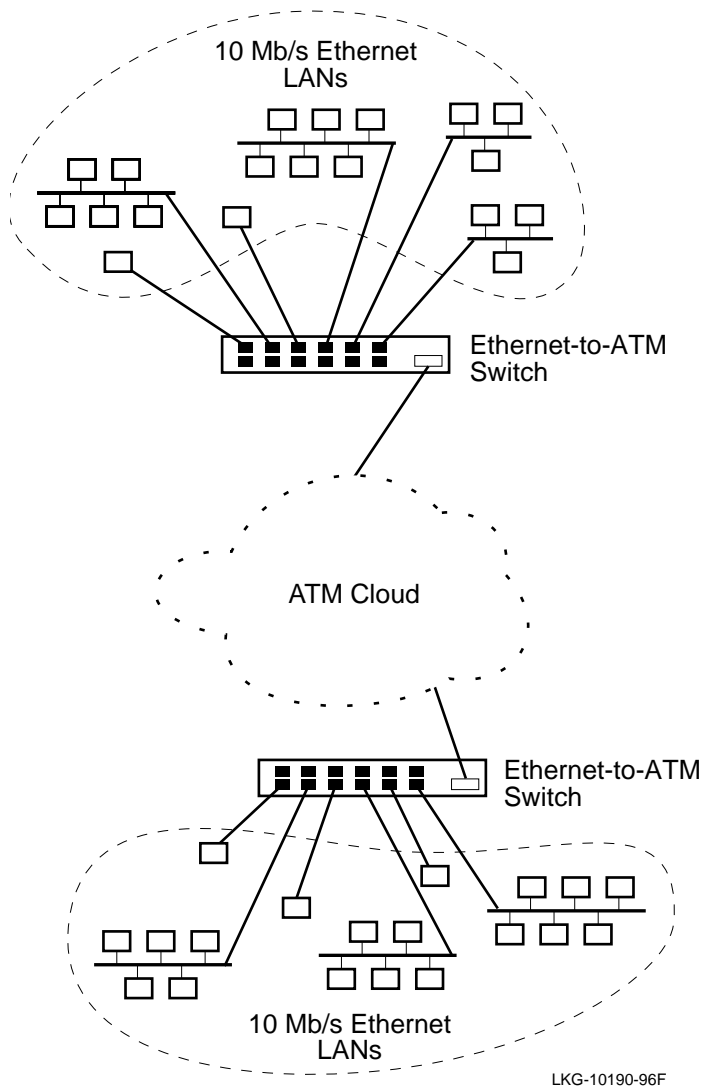
The VNswitch family includes an Ethernet-to-ATM switch called the VNswitch
900EA.The VNswitch 900EA is an Ethernet-to-ATM switch that can connect multiple
LANs to an ATM network. The VNswitch 900EA features a single ATM port and
twelve 10BaseT ports. The VNswitch 900EA is described in detail in Chapter 2. For
more information about Ethernet-to-ATM switching, including an overview of ATM
technology, refer to Chapter 6.

LAN Emulation

LAN Emulation, or LANE (also known as ATM Emulated LANs, or ELANs), is an ATM service described in the ATM Forum's LAN Emulation Specification. This specification defines how LAN applications can run unaltered in an ATM environment.

The VNswitch 900EA supports LAN Emulation (LANE), allowing you to define as many as 16 emulated 802.3/Ethernet LANs that can span an ATM network. To the ATM network, each emulated LAN (ELAN) appears as a LANE client. To the stations connected to the Ethernet ports of the VNswitch, each ELAN is viewed as an Ethernet LAN that spans the ATM cloud and operates at up to 155 Mb/s. Note that an ELAN can also include ATM hosts connected to the ATM network.

Figure 1-13 depicts an emulated LAN on two VNswitch 900EA switches that spans an ATM network.

Figure 1-13: Example of an ELAN Spanning an ATM Network
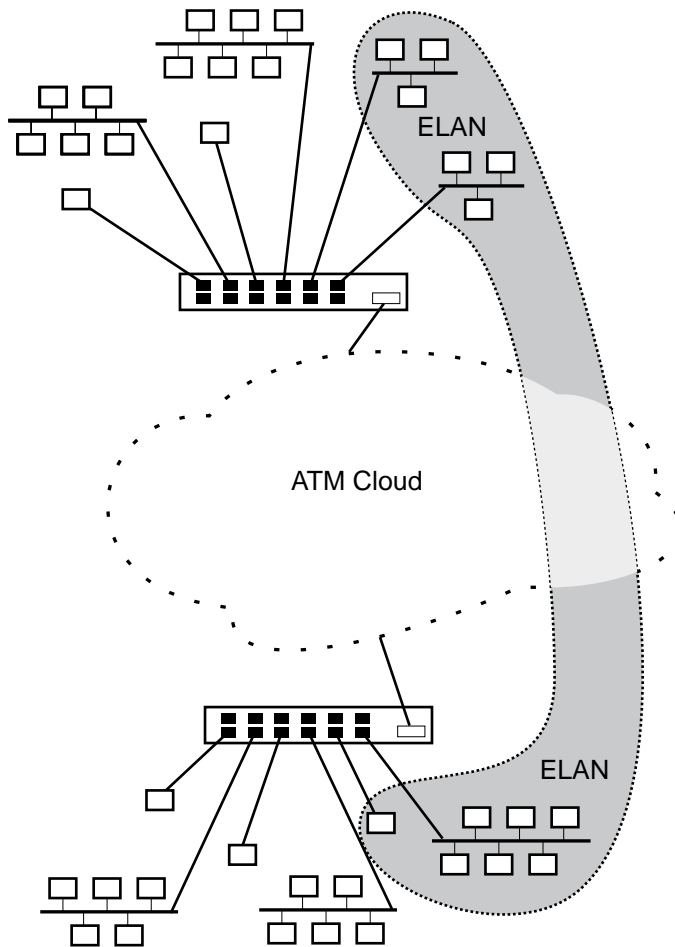
LKG-10191-96F

More detailed information about LAN Emulation, as well as other features of the VNswitch 900EA (such as bridge *tunneling*) is provided in Chapter 3. Also refer to Chapter 6.

VLAN Overview

Introduction

By means of software, LANs attached to ports on one or more switches can be combined into virtual LANs by grouping their associated ports.

NOTE

This section introduces the concept of port-based virtual LANs (VLANs) as implemented in Version 1.5 of the VNswitch 900 series switches. Other types of VLANs are also possible, including those based on MAC address, protocol type, and IP address. The VNswitch 900 series currently supports only port-based VLANs. Future releases of firmware will add support for additional types of VLANs.

What Is a Virtual LAN?

The simplest definition of a virtual LAN (VLAN) is a group of LAN segments that are logically connected through software and isolated from other LAN segments. One or more segments connected to a switch can be logically defined as a single virtual LAN by grouping selected ports into a port group. For example, using the VNswitch 900EX, the twelve 10 Mb/s ports and two Fast Ethernet ports can be divided into various port groups to define several virtual LANs, as illustrated in Figure 1-14.

Figure 1-14: Grouping Ports to Create Virtual LANs

LKG-10186-96F
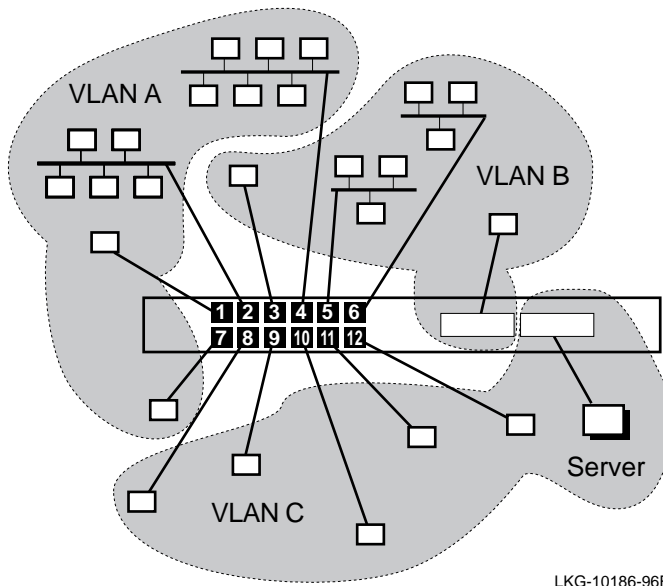
As shown in Figure 1-14, the LAN attached to port 2 has been grouped with the LAN attached to port 4, along with the personal Ethernet LANs attached to ports 1 and 7. These four LANs form a single VLAN (VLAN A). LANs attached to ports 3, 5, and 6 are grouped with one of the Fast Ethernet ports to form another VLAN (VLAN B). The personal Ethernet LANs attached to ports 8 through 12, along with the other Fast Ethernet port, are grouped into VLAN C.

When two or more VNswitch 900 series modules are used in a single DEChub 900 MultiSwitch, ports from the different switches can be combined (through software) to form multiple VLANs, irrespective of the switch to which the ports belong. For example, as shown in Figure 1-15, the upper six ports from switches 1 and 2 are combined to form a single VLAN (VLAN A). The remaining ports from the two switches are grouped to form a second VLAN (VLAN B).

NOTE

To enable VLANs to span multiple modules in the hub, the VNswitch 900 series switches implement a virtual network bus called the *VNbus*. The VNbus is discussed in the following section.

Figure 1-15: Virtual LANs Across Multiple Switches

DEChub 900 MultiSwitch

LKG-10194-96F
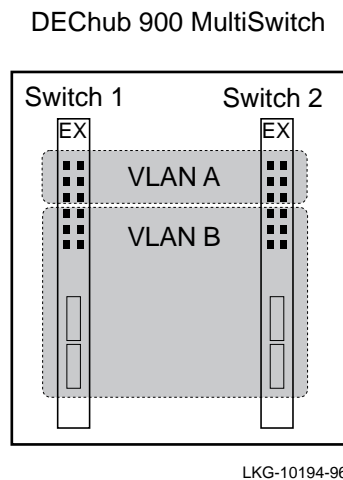
Virtual LANs are a powerful feature of today's switches. They enable a network manager to logically group physically diverse LANs. For example, a LAN on the first floor of an office building can be logically grouped with a LAN on the third floor to form a single virtual LAN. The virtual LAN can then be managed as a unique entity within the network even though the physical components of the VLAN (such as workstations, servers, and so forth) are wired to physically diverse LAN segments.

All members of the VNswitch 900 series support VLAN functionality. The specific VLAN features supported by the VNswitch 900 series switches are described in detail in Chapter 3.

Integrating Multiple Technologies

Introduction

A key feature of the modules in the VNswitch 900 family of switches is their ability to be used together in a single DEChub 900 MultiSwitch to integrate multiple technologies. When configured together in the DEChub 900, the modules are connected by the *VNbus*, a 400 Mb/s bus implemented by the VNswitch 900 series switches on the hub backplane.

Function of the VNbus

The VNbus provides a channel for virtual LANs (VLANs) to span the modules in the hub and also allows bridging between all four technologies (Ethernet, Fast Ethernet, FDDI, and ATM).

Through the VNbus, LANs attached to one module seamlessly connect to LANs attached to another module, as shown in Figure 1-16.

Figure 1-16: Two VNswitch Modules Connected Over a VNbus

LKG-10210-96F
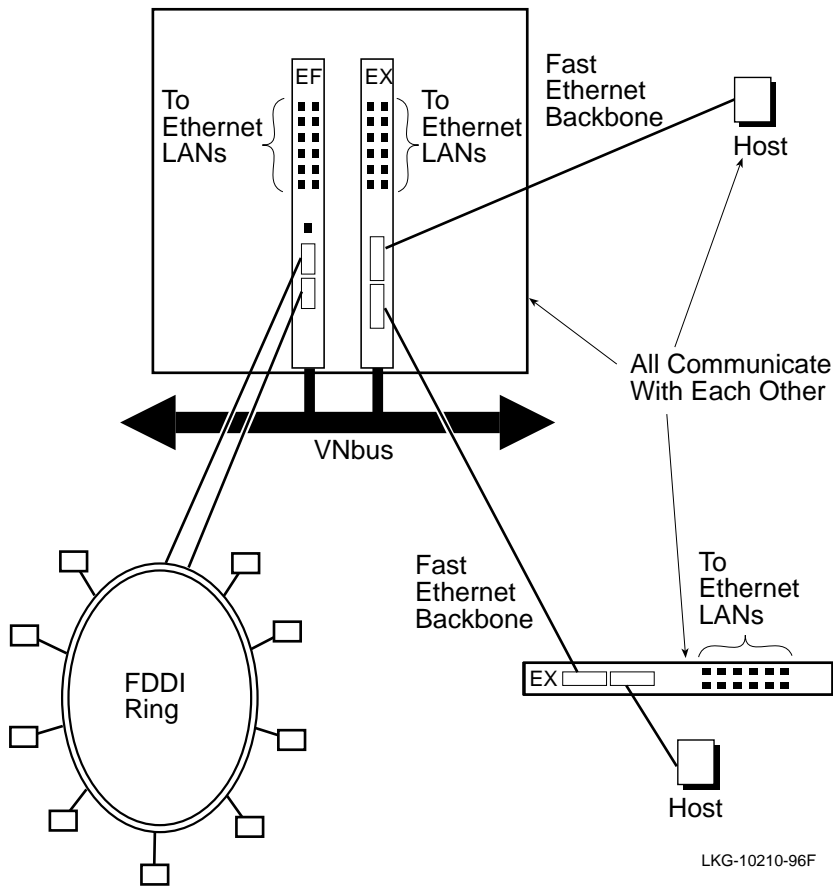
In Figure 1-16, the FDDI LAN attached to the VNswitch 900EF can communicate across the VNbus to the VNswitch 900EX, then over the Fast Ethernet backbone to the Ethernet LANs attached to a remote VNswitch 900EX.

Although not shown in this figure, the VNbus also allows virtual LANs (VLANs) to span the VNswitch modules plugged into the backplane. For example, VLANs on the VNswitch 900EF can be grouped with VLANs on the VNswitch 900EX module to form a single VLAN, as shown earlier in Figure 1-15. Another illustration of multiple VLANs across several VNswitch modules is illustrated in Figure 1-17.

Figure 1-17: Example of Multiple VLANs Across the VNbus
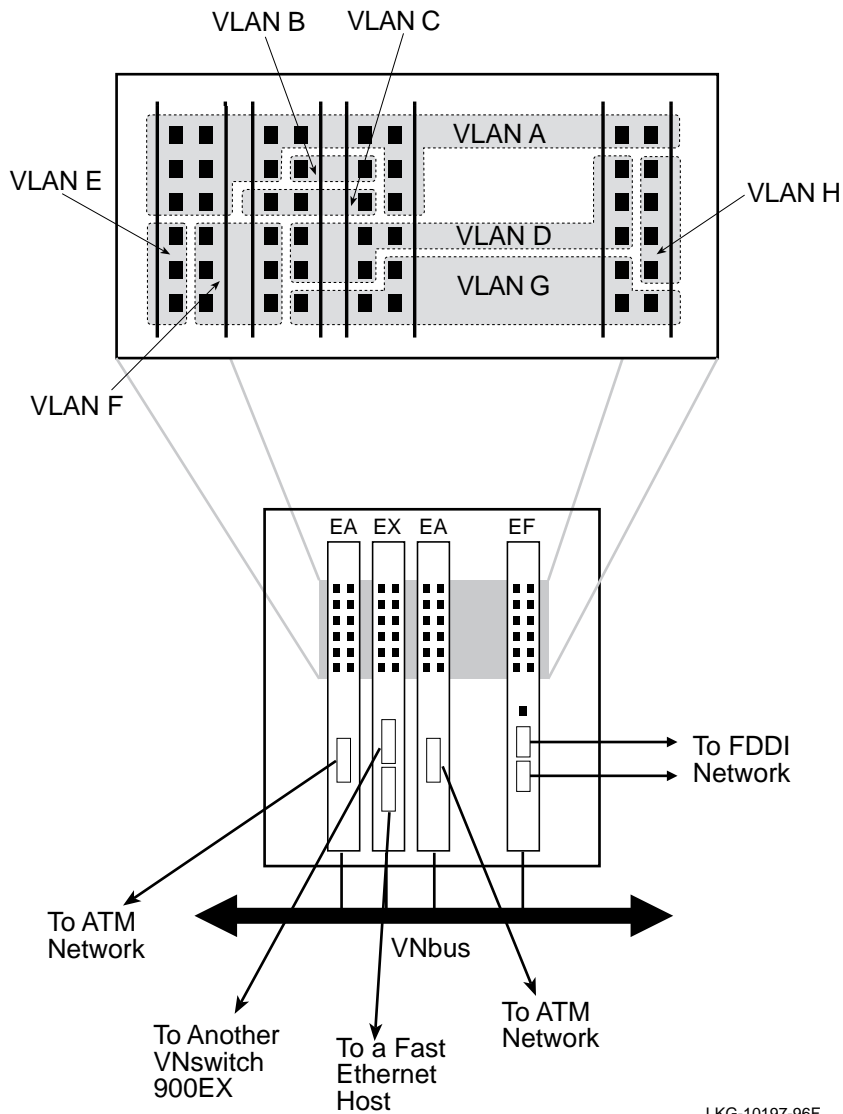
LKG-10197-96F

NOTE

When the VNswitch modules are first plugged into the DEChub 900, all of the attached LANs (including those attached to the high-speed ports) are automatically configured into a single virtual LAN that spans a single 400 Mb/s bus (VNbus) on the backplane. Each VNswitch module can participate in as many as 32 VLANs, and a single VNbus supports up to 63 port-based VLANs. For more information about VLANs and the VNbus, refer to Chapter 3.

Standards-Based VLANs over ATM

Using two or more VNswitch 900EA modules attached to an ATM network, you can define a virtual LAN (VLAN) comprised of LAN segments from multiple switches, even though the switches (and LANs) are located across the ATM network. Digital refers to this feature as *standards-based VLANs over ATM* because it is an implementation of VLANs coupled with the ATM Forum's LAN Emulation standard (LANE V1.0).

Figure 1-18 illustrates a single extended VLAN comprised of ports from two VNswitch 900EA modules. In this illustration, VLAN A from switch 1 is joined with VLAN A from switch 2 to form a single ELAN, effectively creating a single extended VLAN (VLAN A) that spans the ATM network.

Figure 1-18: Virtual LANs Across an ATM Network

LKG-10192-96F
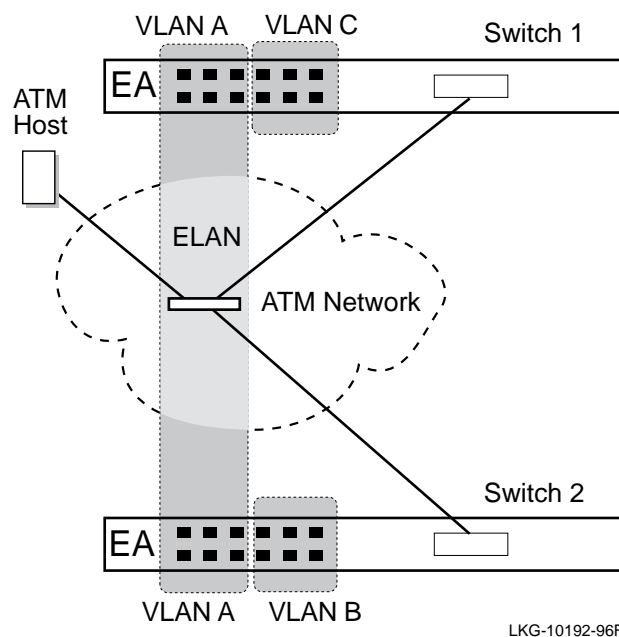
As many as 16 emulated LANs (ELANs) are supported on each VNswitch 900EA module. Each ELAN may consist of VLANs that span multiple modules in the hub, allowing LANs from the VNswitch 900EE, EX, or EF to be joined with LANs on other VNswitch modules residing in a separate hub on the other side of an ATM network. Refer to Figure 1-19.

Figure 1-19: Spanning Multiple Technologies with the VNswitch 900 Series Switches
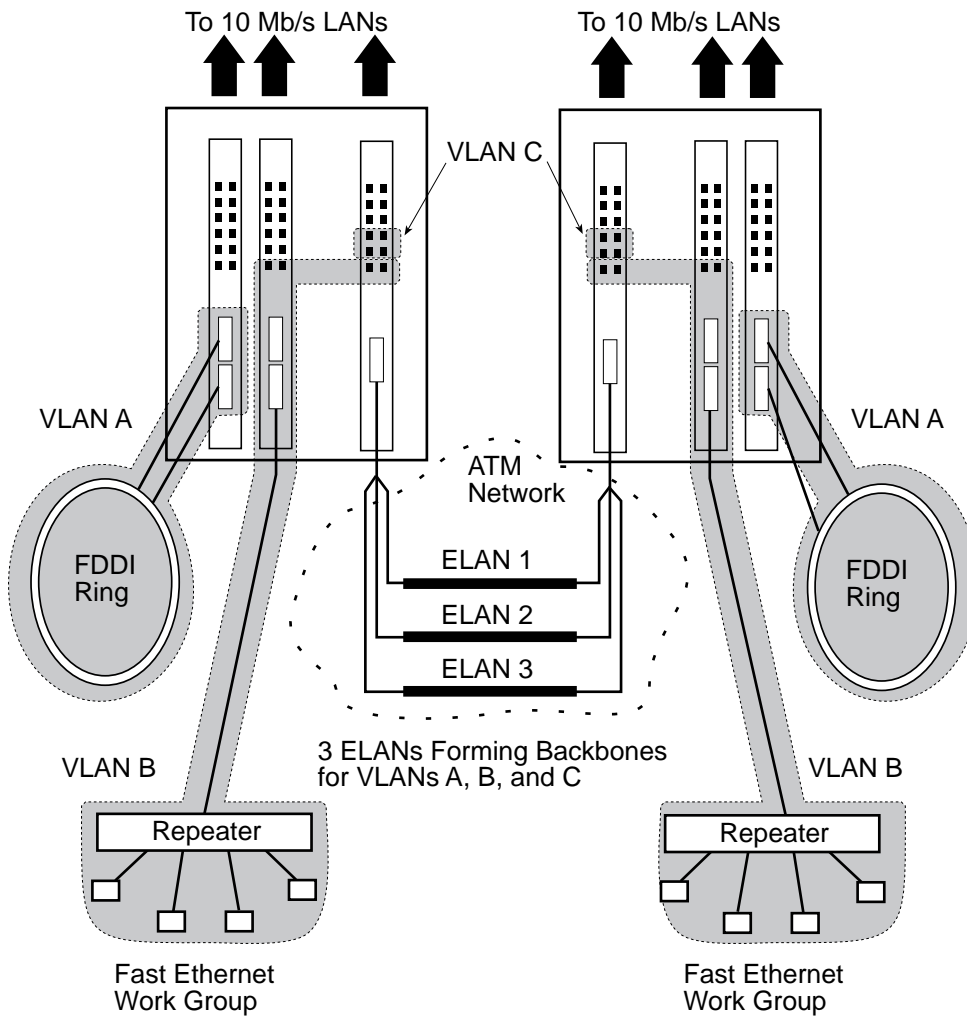
LKG-10215-96F

Figure 1-19 illustrates a more complex example of standards-based VLANs over ATM, and also shows how the VNswitch modules enable seamless switching between LANs of different network technologies, even across an ATM network. In this illustration, VLANs A, B, and C are connected to their counterparts over ELANs 1, 2, and 3 respectively, allowing each VLAN to span the ATM network.

Summary

Effectively, any combination of Ethernet, Fast Ethernet, FDDI, and ATM technologies is possible using the different modules in the VNswitch 900 series. Chapter 7 discusses some of the many combinations and illustrates several possible configurations.

Chapter 2

VNswitch 900 Series Switches

Overview

Introduction

This chapter describes each module in the VNswitch family. It includes a description of each module, its features, and its specifications. It also provides information about how to order each module.

In This Chapter

The following topics are covered in this chapter:

VNswitch 900 Family of Products

Introduction

The VNswitch 900 series switch products are high-performance virtual network switches. The term *virtual network switch* refers to the capability of the VNswitch to connect virtual local area networks (VLANs) that can span multiple technologies. The VNswitch 900 series products bring the benefits of high performance and IEEE 802.1d-compliant switching to a wide range of network applications, from *personal Ethernets* to enterprise-wide networking.

Description

Currently, the VNswitch family consists of the following modules:

Product Name	Part Number	Description
VNswitch 900EE	DVNEE-MA	24-port Ethernet-to-Ethernet virtual network switch.
VNswitch 900EX	DVNEX-MX	12-port Ethernet to Fast Ethernet (two Fast Ethernet ports) virtual network switch.
VNswitch 900EF	DVNEF-MM DVNEF-MX	12-port Ethernet-to-FDDI virtual network switch. There are two versions of this product: • The -MM version is equipped with factory installed ANSI MIC connectors for FDDI over Multimode fiber (MMF). • The -MX version uses an optional Modular Media Interface (MMI). The MMI options are described later in this chapter.
VNswitch 900EA	DVNEA-MX	12-port Ethernet-to-ATM virtual network switch.

These powerful and flexible 802.1d-compliant switches combine data integrity with high performance. They are designed to connect multiple local area networks (LANs) into a *virtual switched network*. A virtual switched network is any network that employs switching between multiple LANs configured by means of software into logical network segments (called virtual LANs). A small virtual network could consist

of a single VNswitch used as a hub to segment a network into several virtual LANs (consisting of both personal Ethernets and workgroup Ethernet segments), as shown in Figure 2-1. A large, enterprise-wide virtual network might consist of numerous LANs and VLANs connected across multiple technologies such as Ethernet, FDDI, and ATM as shown in Figure 2-2.

Figure 2-1: Small Switched Virtual Network

VNswitch 900EE
Installed in a
DEChub ONE
Docking Station

= Virtual LAN
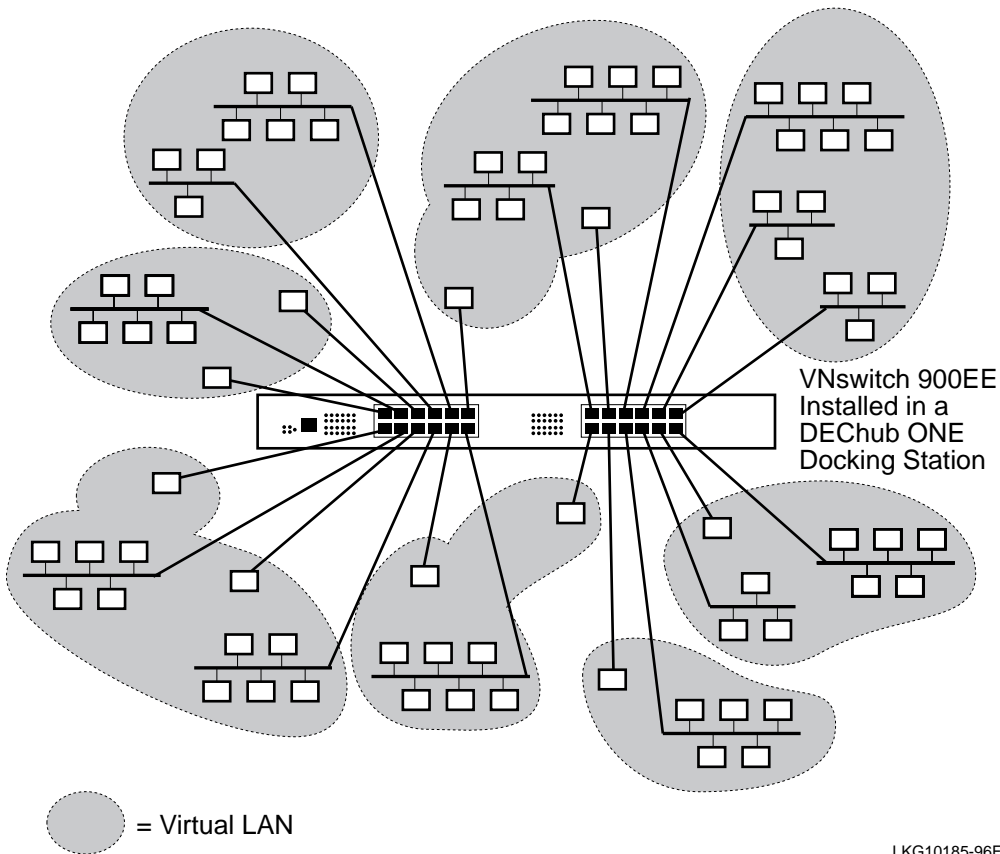
LKG10185-96F

Figure 2-2: Large Switched Virtual Network

Single VLAN
Comprised of 24
10 Mb/s LANs
and FDDI Ring

EF

EA

EA

VLAN with
6 LANs

VLAN with
6 LANs and the
High-speed
ATM Port

FDDI
Ring

ATM
Network

ATM
Network

Single VLAN
Comprised of 72
10 Mb/s LANs

To 10 Mb/s
Ethernet LANs

Fast
Ethernet
Host

EX EF EE EA

To Fast
Ethernet
Repeater

To 10 Mb/s
Ethernet LANs
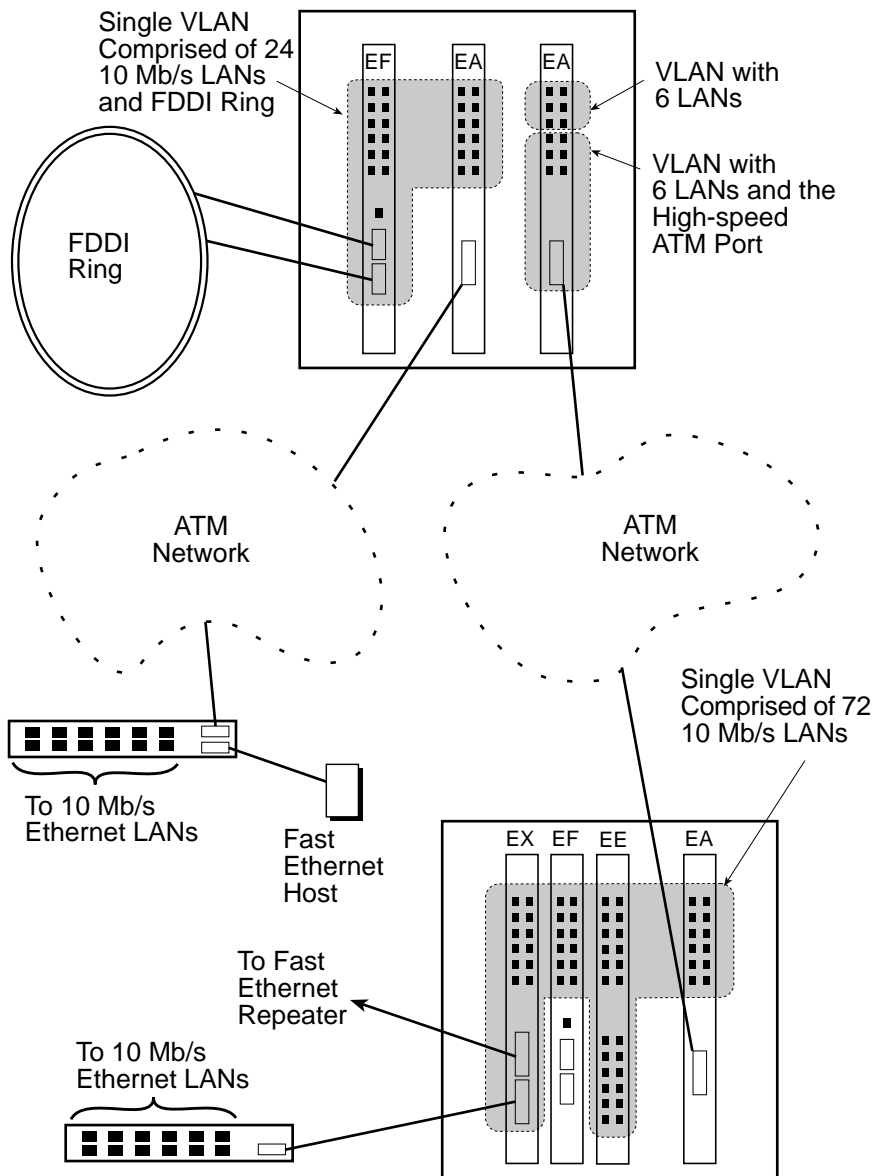
LKG-10199-96F

NOTE

By default, all ports (including the high-speed ports) on the VNswitch modules are grouped into a single VLAN, called the *default VLAN*. Through software commands, the ports can be arbitrarily divided into port groups (as shown in Figure 2-2). Any ports not included in a port group (or VLAN) remain in the default VLAN. Thus, the ports in this illustration that are not shaded are all members of the default VLAN. More information about VLANs is provided in Chapter 3, Features and Functionality.

Each VNswitch 900 series module can be installed standalone in a DEChub ONE docking station (as shown in Figure 2-1) or configured with other modules in a DEChub 900 MultiSwitch (as shown in Figure 2-2). When installed in a DEChub 900, the modules interconnect by means of a 400 Mb/s VNbus. When fully configured in a DEChub 900, the VNswitch 900 series modules can provide as many as 192 switched 10BaseT ports in a single hub.

Each VNswitch 900 series module performs layer 2 line-rate switching of LAN frames on all ports and between different technologies, including Ethernet, Fast Ethernet, FDDI, and ATM. A software routing option is available that enables each module to perform *multilayer switching*, that is, both layer 2 *and* layer 3 switching. Information about the optional routing software is provided in Chapter 3.

VNbus

The VNbus is an integral component of the DEChub 900 architecture that provides interconnectivity between attached VNswitch 900 series modules. The VNswitch 900 series modules support three separate VNbuses, providing a total capacity of 1.2 Gb/s (gigabits per second). Each VNbus provides the following functionality:

- Allows VLANs to span across VNswitch 900 series modules in the DEChub 900

- Allows a VNswitch 900EE module to be used as an expansion module for any other module in the VNswitch 900 series

- Allows bridging between all four technologies (Ethernet, Fast Ethernet, FDDI, and ATM)

- Allows routing between all four technologies when equipped with the optional routing software

For more information about the VNbus, refer to Chapter 3.

Highlights

All modules in the VNswitch 900 series have the following features:

- Auto-configuring at powerup (Plug-n-Play functionality) using factory-set defaults

- Fully IEEE 802.1d compliant

- Support multiple spanning trees

- Allow up to 12 front panel Ethernet ports to be redirected to the backplane

- Support both half- and full-duplex operation on all Ethernet ports

- Full line-rate bridge forwarding—up to 750,000 packets per second (pps) per module

- User-settable rate limiting

- Offer multiple filtering options

- Support port-based VLANs:

  — Each module supports as many as 32 VLANs

  — Each VNbus supports as many as 63 VLANs

- Software-upgradable to routing option via downline load

- Installable as either a standalone unit in a DEChub ONE docking station or configured with other modules in a DEChub 900 MultiSwitch

- Support up to 8192 network addresses:

  — 7936 MAC addresses

  — 256 internal addresses for management and other purposes

- Manageable via SNMP and/or Command Line Interface (CLI) management

- Fully interoperable with other modules installed in the DEChub 900 (including both DEChub 900 and DEChub 90 modules)

VNswitch 900EE Module

Introduction

The VNswitch 900EE module (Figure 2-3) is a 24-port SNMP manageable Ethernet switch, providing full-speed switching between each of its twenty-four 10BaseT 802.3/Ethernet ports. It supports various filtering capabilities (source address, destination address, protocol type, and VLAN) as well as providing a large address table (8192 entries). The module is fully IEEE 802.1d standards compliant.

Figure 2-3: VNswitch 900EE Network Module
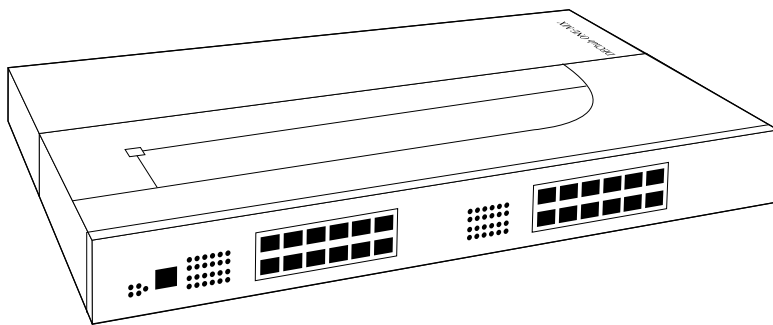
LKG-10246-96F

All of the module's 24 front panel 10BaseT 802.3/Ethernet ports are auto-configured to connect to the 400 Mb/s VNbus on the DEChub 900 backplane. Additionally, as many as 12 of the ports can be individually redirected to a DEChub 900 backplane channel (LAN segment) to form a backplane LAN. Also, port 12 can be redirected to the DEChub 900 backplane ThinWire LAN segment.

When connected to a DEChub ONE docking station (DEF1H or DEHUA), the module's (front panel) port 11 can be optionally assigned to the docking station's AUI port.

When installed into a DEChub 900 and connected to the VNbus, the unit can function as an expansion module for any of the other VNswitch 900 series modules, as illustrated in Figure 2-4.

Figure 2-4: VNswitch 900EE Used as an Expansion Module

VLAN Composed of
36 10 Mb/s LANs and
two 100 Mb/s Connections

EX

EE

To Other
VNswitch 900EX
Modules
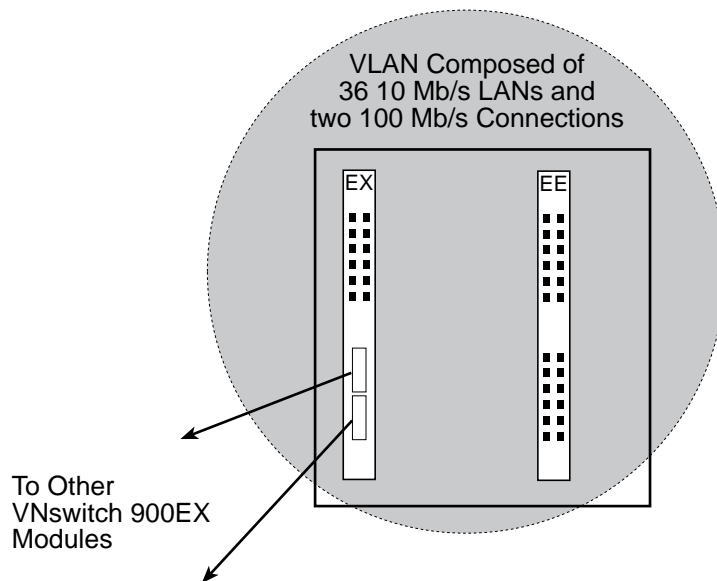
LKG-10195-96F

The VNswitch 900EE switch includes nonvolatile flash memory for software upgrades of the device firmware using the Trivial File Transfer Protocol (TFTP) load protocol (available on a wide variety of platforms including DOS, ULTRIX, UNIX, and OpenVMS operating systems). This eliminates the need to replace or upgrade hardware in the future, and avoids the associated costs to network users.

Like all VNswitch 900 series modules, the VNswitch 900EE also features integrated SNMP for easy, comprehensive management.

Highlights

Following are the highlights of the VNswitch 900EE switch:

- Provides 24 low-cost, high-performance, switched 10BaseT Ethernet ports.

- Supports half- or full-duplex operation over all ports.

- Functions as an expansion unit for the other VNswitch 900 series modules.

- Offers standalone (DEHUA or DEF1H) or DEChub 900 MultiSwitch configuration options.

- Improves Ethernet LAN configuration flexibility and performance by dynamically switching multiple Ethernet LANs.

- Improves efficiency and utilization of current Ethernet networks.

- Supports VLANs across the DEChub 900 backplane via the VNbus.

- Features upgradable device firmware (in nonvolatile flash memory) using Trivial File Transfer Protocol (TFTP) with HUBloader or through the OBM port with any TFTP server.

- Offers user-configurable rate limiting for broadcast and multicast packets by address and specific protocol.

- Supports 8192 network addresses.

- Features a Command Line Interface (CLI) that is accessible from a directly attached terminal or remotely through a Telnet connection.

- Features a built-in SNMP management agent. It supports the comprehensive graphical user interface (GUI), via clearVISN, that is identical for both in-band and out-of-band management.

- Supports a complete set of standard and private MIBS, including: MIB-II (RFC 1213), Bridge MIB (RFC 1493), Ethernet MIB (RFC 1398), SNMP management (RFC 1157), and Digital enterprise MIBs for full management and configuration capability via SNMP. (Refer to Appendix A for a complete list of MIBs supported by the VNswitch family of products.)

- Features out-of-band management using SNMP over SLIP and PPP through the OBM connector as an alternative to normal in-band management.

- Supports full-performance filtering between all of its twenty-four 10BaseT ports at full Ethernet speeds (14,880 packets/second/Ethernet).

- Provides 802.1d spanning tree loop detection per port.

- Provides the ability to turn off the spanning tree algorithm on a per-port basis.

- Supports backplane LANs by allowing up to 12 ports to be individually directed (through management software) to the DEChub 900 backplane.

Department Backbone with the VNswitch 900EE

Figure 2-5 shows the VNswitch 900EE module used as a high-performance Ethernet switch for a department backbone.

Figure 2-5: VNswitch 900EE Departmental Backbone

LAN Work Groups

Standalone
in a
Docking Station

OR

Hub-Based
(DEChub 900)

Personal Ethernets

Server
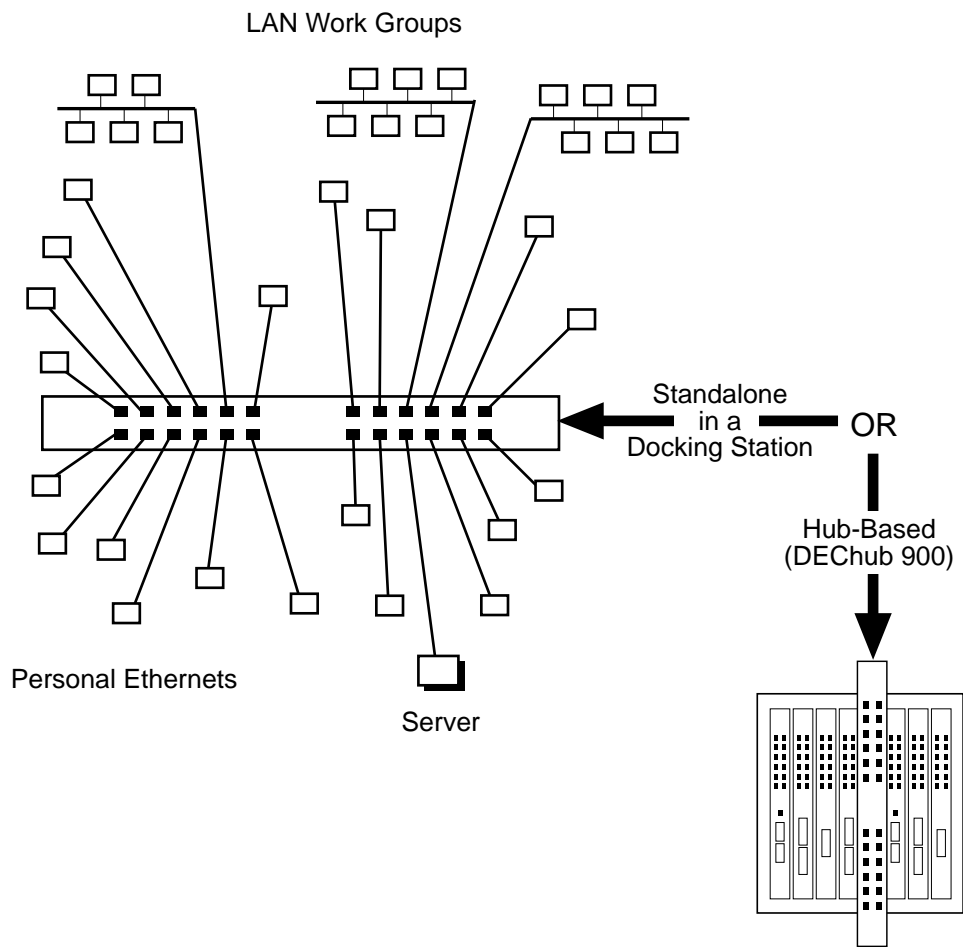
LKG10184-96F

VNswitch 900EE Ordering Information

Description	Order Number
VNswitch 900EE network module.	DVNEE-MA
DEChub 900 MultiSwitch 8-slot intelligent hub chassis.	DMHUB-*xx*
DEChub ONE docking station for standalone use of the VNswitch 900EE module.	DEHUA-*xx* or DEF1H-*xx*

Note: *xx* = Country kit code. Check with your Digital sales representative for the appropriate code for your location.

VNswitch 900EE Product Specifications

Parameter	Specification
Environment	
Operating temperature[1]	5°C to 50°C (41 °F to 122° F)
Relative humidity	10% to 95% noncondensing
Altitude	
• Operating	Sea level to 2400 m (8,000 ft)
• Non-operating	Sea level to 4900 m (16,000 ft)
Power	50.8 W, total power
	8.6 A, 5 Vdc
	0.12 A, 12 Vdc
	0.52 A, 15 Vdc
Physical	
Height	44.45 cm (17.5 in)
Width	4.45 cm (1.75 in)
Depth	15.25 cm (6.0 in); 25.40 cm (10.0 in) with a DEChub ONE docking station
Weight	1.9 kg (4.3 lb)[2]
Connectors	24 shielded 10BaseT (MJ-8 RJ45 type), supporting both UTP and ScTP cabling
Shock (Class A/B for products weighing under 100 lb)	10 G / 10 ms half sine pulse in three orthogonal axes
Vibration (Class C)	5 to 200 Hz sine sweep @ 0.25 G limited by 0.5mm (0.02 in) displacement DA* 200 to 500 Hz sine sweep @ 0.10 G
Certification	CE, CSA, FCC, TÜV, UL, VCCI
Standards Compliancy	IEEE 802.1d, IEEE 802.3, Ethernet V1 & V2, 802.1h, 802.1i, 802.2

[1.] For sites above 2400 m (8,000 ft), decrease the operating temperature specification by 1.8°C for each 1000 m or 3.2°F for each 3200 ft.

[2.] Actual module weight (no shipping container). Include an additional 1.59 kg (3.5 lb) when attached to a DEChub ONE docking station.

Acoustical Specifications

Product	Sound Power Level L_{WAd}, B	Sound Pressure Level L_{pAm}, dBA (bystander positions)
	Idle/Operate	Idle/Operate
DVNEE	5.0	37
DVNEE + DEHUA	5.3	40
DVNEE + DEF1H	5.4	39

Note: Current values for specific configurations are available from Digital Equipment Corporation representatives. 1B = 10 dBA.

For More Information

Topic	Where to Get More Information
Additional information about the functions and features of the VNswitch family	Refer to Chapter 3.
Examples of network configurations using the VNswitch 900EE module	Refer to Chapters 4 through 7 for various examples of network configurations using the VNswitch 900EE module.

VNswitch 900EX Module

Introduction

The VNswitch 900EX module (Figure 2-6) is a 12-port (10BaseT ports) to 2-port Fast Ethernet switch that provides automatic dual-speed switching between standard 10 Mb/s Ethernet and 100 Mb/s Ethernet. The switch is SNMP manageable and supports various filtering capabilities (source address, destination address, protocol type, and VLAN) as well as providing a large address table (8192 entries). The module is fully IEEE 802.1d standards compliant.

Figure 2-6: VNswitch 900EX Network Module
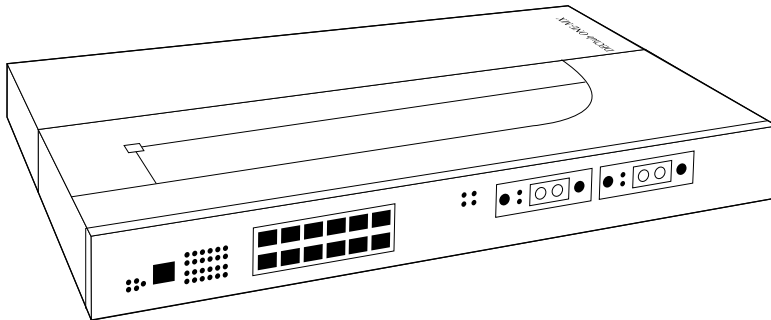
LKG-10245-95F

The unit is auto-configured to connect to the 400 Mb/s VNbus, which provides interoperability across the backplane with other VNswitch modules while maintaining front panel connectivity. Additionally, each of the module's 12 front-panel 10BaseT 802.3/Ethernet ports can be individually redirected to a DEChub 900 backplane LAN segment to create a backplane LAN. Also, port 12 can be redirected to the DEChub 900 backplane ThinWire LAN segment.

When connected to a DEChub ONE docking station (DEF1H or DEHUA), the module's (front panel) port 11 can be optionally assigned to the docking station's AUI port.

The VNswitch 900EX module includes nonvolatile flash memory for software upgrades of the device firmware using the Trivial File Transfer Protocol (TFTP) load protocol (available on a wide variety of platforms including DOS, ULTRIX, UNIX, and OpenVMS operating systems). This eliminates the need to replace or upgrade hardware in the future, and avoids the associated costs to network users.

Like all VNswitch 900 series modules, the VNswitch 900EX also features integrated SNMP for easy, comprehensive management.

Highlights

Following are the highlights of the VNswitch 900EX switch:

- Provides switching between two high-speed 100 Mb/s Fast Ethernet ports and twelve 10BaseT Ethernet ports.

- Supports both 100BaseTX and 100BaseFX Fast Ethernet media standards.

- Improves Ethernet LAN configuration flexibility and performance by dynamically switching between 10 Mb/s Ethernet and 100 Mb/s Ethernet.

- Supports front-insertable, Modular Media Interfaces (MMIs):

  — TX (UTP-5 with RJ-45 connectors)

  — FX (Multimode fiber with SC connectors)

- Allows the TX port to be configured to operate at either 10 Mb/s or 100 Mb/s.

- Supports half- and full-duplex operation over all ports, including the Fast Ethernet ports.

- Supports 8192 network addresses.

- Supports auto-negotiation as specified by IEEE 802.3u.

- Offers standalone (DEHUA or DEF1H) or DEChub 900 MultiSwitch configuration options.

- Supports VLANs across the DEChub 900 backplane via the VNbus.

- Features upgradable device firmware (in nonvolatile flash memory) using Trivial File Transfer Protocol (TFTP) with HUBloader or through the setup port with any TFTP server.

- Offers user-configurable rate limiting for broadcast and multicast packets by address and specific protocol.

- Features a Command Line Interface (CLI) that is accessible from a directly attached terminal or remotely through a Telnet connection.

- Features a built-in SNMP management agent. It supports a comprehensive graphical user interface (GUI), via clearVISN, that is identical for both in-band and out-of-band management.

VNswitch 900EX Module

- Supports a complete set of standard and private MIBs for full management and configuration capability via SNMP. Supported MIBs include MIB-II (RFC 1213), Bridge MIB (RFC 1493), Ethernet MIB (RFC 1398), SNMP management (RFC 1157), and Digital enterprise MIBs. (Refer to Appendix A for a complete list of MIBs supported by the VNswitch family of products.)

- Features out-of-band management using SNMP over SLIP and PPP through the OBM connector as an alternative to normal in-band management.

- Supports full-performance filtering between all of its twelve 10BaseT ports at full Ethernet speeds (14,880 packets/second/Ethernet).

- Provides 802.1d spanning tree loop detection per port.

- Provides the ability to turn off the spanning tree algorithm on a per-port basis.

- Supports backplane LANs by allowing the 10BaseT ports to be individually directed (through management software) to the DEChub 900 backplane.

Enterprise Switching with the VNswitch 900EX

When installed into a DEChub 900, the VNswitch 900EX module can function as a high-speed 100 Mb/s Fast Ethernet backbone for LANs connected to other modules in the hub.

Figure 2-7 shows the VNswitch 900EX module connecting 36 local area networks to a 100 Mb/s Fast Ethernet enterprise backbone. In this illustration, the VNswitch 900EE is used as an expansion module, adding 24 more addressable ports to the VNswitch 900EX.

Figure 2-7: VNswitch 900EX Used as a High-Speed Backbone

VLAN Comprised of
18 10 Mb/s LANs

100 Mb/s
Backbone

To Another
VNswitch 900EX

100 Mb/s
Backbone

To a Fast Ethernet
Work Group

VLAN Comprised of 12
10 Mb/s LANs
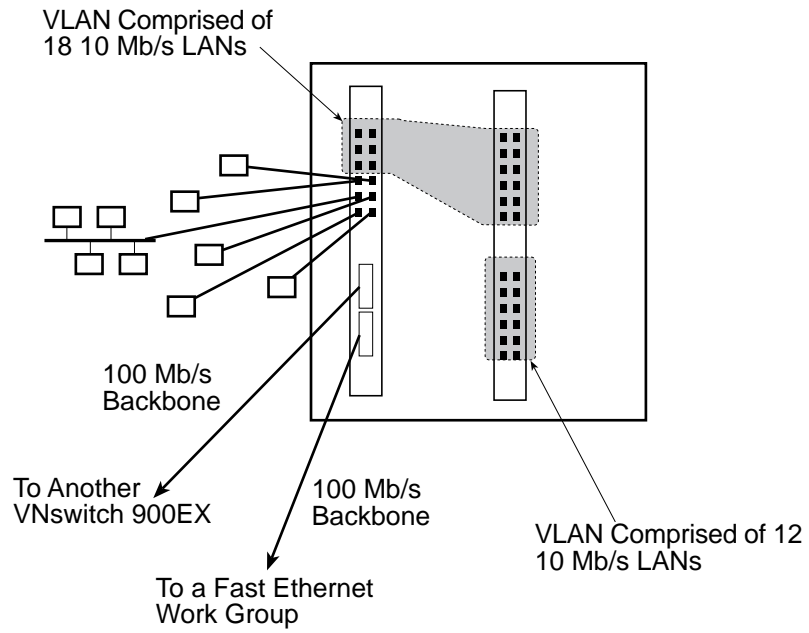
LKG-10196-96F

VNswitch 900EX Ordering Information

Description	Order Number
VNswitch 900EX network module.	DVNEX-MX
DEChub 900 MultiSwitch 8-slot intelligent hub chassis.	DMHUB-*xx*
DEChub ONE docking station for standalone use of the VNswitch 900EX module.	DEHUA-*xx* or DEF1H-*xx*

Note: *xx* = Country kit code. Check with your Digital sales representative for the appropriate code for your location.

VNswitch 900EX Product Specifications

Parameter	Specification
Environment	
Operating temperature[1]	5°C to 50°C (41°F to 122°F)
Relative humidity	10% to 95% noncondensing
Altitude	
• Operating	Sea level to 2400 m (8,000 ft)
• Non-operating	Sea level to 4900 m (16,000 ft)
Power	50.8 W, total power
	8.6 A, 5 Vdc
	0.12 A, 12 Vdc
	0.52 A, 15 Vdc
Physical	
Height	44.45 cm (17.5 in)
Width	4.45 cm (1.75 in)
Depth	15.25 cm (6.0 in); 25.40 cm (10.0 in) with a DEChub ONE docking station
Weight	1.9 kg (4.2 lb)[2]
Connectors	12 shielded 10BaseT (MJ-8 RJ45 type) supporting both UTP and ScTP cabling
	2 MMI slots supporting both UTP-5 with RJ-45 connectors (TX) and multimode fiber with SC connectors (FX)
Shock (Class A/B for products weighing under 100 lb)	10 G / 10 ms half sine pulse in three orthogonal axes
Vibration (Class C)	5 to 200 Hz sine sweep @ 0.25 G limited by 0.5mm (0.02 in) displacement DA* 200 to 500 Hz sine sweep @ 0.10 G
Certification	CE, CSA, FCC, TÜV, UL, VCCI
Standards Compliancy	IEEE 802.1d, IEEE 802.3, Ethernet V1 & V2, 803.3aa, 802.1h, 802.1i, 802.2, 802.3u

[1.] For sites above 2400 m (8,000 ft), decrease the operating temperature specification by 1.8°C for each 1000 m or 3.2°F for

VNswitch 900EX Module

each 3,200 ft·

2. Actual module weight (no shipping container). Include an additional 1.59 kg (3.5 lb) when attached to a DEChub ONE docking station.

Acoustical Specifications

Product	Sound Power Level L_{WAd}, B	Sound Pressure Level L_{pAm}, dBA (bystander positions)
	Idle/Operate	Idle/Operate
DVNEX	5.0	37
DVNEX + DEHUA	5.3	40
DVNEX + DEF1H	5.4	39

Note: Current values for specific configurations are available from Digital Equipment Corporation representatives. 1B = 10 dBA.

For More Information

Topic	Where to Get More Information
Additional information about the functions and features of the VNswitch family	Refer to Chapter 3.
Examples of network configurations using the VNswitch 900EX module	Refer to Chapter 4 and Chapter 7.

VNswitch 900EF Module

Introduction

The VNswitch 900EF module (Figure 2-8) is a high-throughput, SNMP manageable, multiport network switch for linking Ethernet LANs with an FDDI backbone or local area network. The module improves the configuration flexibility of your network by easily integrating Ethernet and FDDI LANs within a single standalone or hub-based architecture. It provides switching for up to 12 Ethernet LANs and 1 high-speed FDDI network.

Figure 2-8: VNswitch 900EF Network Module
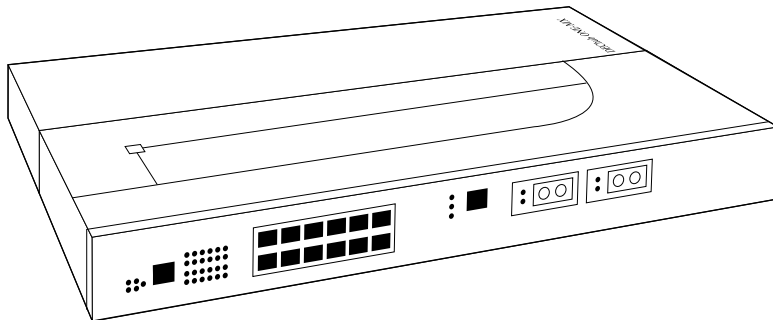
LKG-10244-95F

The VNswitch 900EF module has twelve 10BaseT Ethernet switched ports on the front panel and one FDDI port pair (DAS or SAS).

The unit is auto-configured to connect to the 400 Mb/s VNbus, which provides interoperability across the backplane with other VNswitch modules while maintaining front panel connectivity. Additionally, each of the module's 12 front panel 10BaseT 802.3/Ethernet ports can be individually redirected to a DEChub 900 backplane LAN segment to create backplane LANs.

With clearVISN management, the FDDI PHY ports (1A/M and 1B/S) can also be individually switched to a DEChub 900 backplane LAN. Also, port 12 can be redirected to the DEChub 900 backplane ThinWire LAN segment or to other backplane LANs.

When connected to a DEChub ONE docking station (DEF1H or DEHUA), the module's (front panel) port 11 can be optionally assigned to the docking station's AUI port.

VNswitch 900EF Module

The VNswitch 900EF network module is a true backbone network switch that supports complete filtering capabilities (source address, destination address, protocol type). It includes a large address table (8192 entries) and is fully IEEE 802.1d standards compliant, ensuring the high performance and packet integrity required in large networks.

The VNswitch 900EF network module also offers nonvolatile flash memory for software upgrades of the device firmware using Trivial File Transfer Protocol (TFTP) load protocol (available on a wide variety of platforms including DOS, ULTRIX, UNIX, and OpenVMS operating systems). This eliminates the need to replace or upgrade hardware in the future, and avoids the associated costs to network users.

Like all VNswitch 900 series switches, the VNswitch 900EF network module features integrated SNMP for easy, comprehensive management.

Highlights

Following are the highlights of the VNswitch 900EF switch:

- Provides 12 Ethernet ports (UTP-5) plus 1 DAS or SAS FDDI port pair.

- MX-version provides front-insertable Modular Media Interfaces (MMIs):

  — UTP-5 with RJ-45 connectors

  — Multimode fiber with SC connectors

  — Single-mode fiber with SC connectors

- MM-version features factory-installed ANSI MIC connectors for multimode fiber.

- FDDI ports can be directed to the DEChub backplane.

- Features a Command Line Interface (CLI) that is accessible from a directly attached terminal or remotely through a Telnet connection.

- Features a built-in SNMP agent for in-band and out-of-band management, whether configured in a single-slot standalone DEChub ONE chassis or installed in a DEChub 900 MultiSwitch enterprise hub.

- Features upgradable device firmware (in nonvolatile flash memory) using Trivial File Transfer Protocol (TFTP) with HUBloader or through the setup port with any TFTP server.

- Offers multiple filtering options, including source and destination address and protocol.

- Supports high filtering rates of 14,880 pps on each Ethernet port, and a filtering rate of more than 440,000 pps on the FDDI port.

- Supports full-duplex FDDI at 200 Mb/s.

- Supports 8192 network addresses.

- Handles IP fragmentation of large packets from FDDI per RFC 791 (Internet Protocol) and RFC 1191 (Path MTU Discovery).

- Provides translation between FDDI and 802.3/Ethernet frame formats for direct transparent connections, translation of AppleTalk V1 and V2 AARP packets, and handling of raw 802.3 Novell IPX packets.

- Provides Appletalk NTP support.

- Supports a complete set of standard and private MIBs for full management and configuration capability via SNMP. Supported MIBs include MIB-II (RFC 1213), Bridge MIB (RFC 1493), Ethernet MIB (RFC 1398), FDDI MIB (RFC 1512), SNMP management (RFC 1157), and Digital enterprise MIBs. (Refer to Appendix A for a complete list of MIBs supported by the VNswitch family of products.)

- Features out-of-band management using SNMP over SLIP and PPP through the OBM connector as an alternative to normal in-band management.

- Provides user-configured rate limiting for broadcast and multicast packets by address and specific protocol.

- Supports operation in either an FDDI tree or dual ring configuration. FDDI port 1A/M can be switched to emulate an M port of a concentrator. This action causes FDDI port 1B/S to automatically switch to emulate an S port.

- Provides 802.1d spanning tree loop detection.

- Provides the ability to turn off the spanning tree algorithm on a per-port basis.

- Supports backplane LANs by allowing the 10BaseT ports to be individually directed (through management software) to the DEChub 900 backplane.

- Supports backplane FDDI dual rings or trees for attachment to concentrators by allowing both FDDI ports to be directed to the DEChub 900 backplane.

Virtual Network Switching with the VNswitch 900EF

Figure 2-9 shows the VNswitch 900EF module used as a high-performance switch connecting multiple LANs and VLANs in an enterprise-wide network. In this illustration, the FDDI ports on the EF module adjacent to the FDDI concentrator (the FH module) are directed to the backplane to join an FDDI backplane LAN with the concentrator.

Figure 2-9: VNswitch 900EF Virtual Network Switch

All LANs Are Connected
and Have Access
to FDDI Rings

84 LANs   12 LANs

EF EE EE EE   FH EF

FDDI Ports
Directed to the
Backplane

FDDI
Ring

FDDI Stations
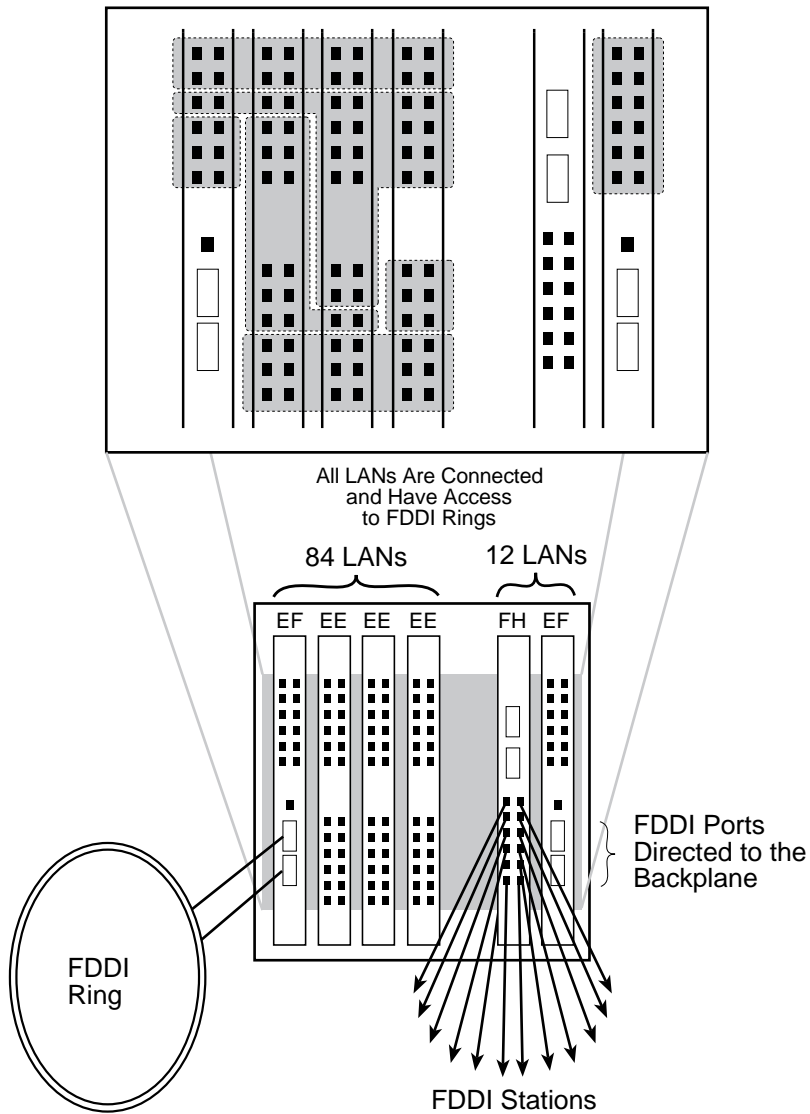
LKG-10198-96F

VNswitch 900EF Ordering Information

Description	Order Number
VNswitch 900EF network module with factory-installed ANSI MIC connectors for FDDI over Multimode fiber (MMF). This model includes FDDI Optical Bypass Relay (OBR) support.	DVNEF-MM
VNswitch 900EF network module with op-tional Modular Media Interface (MMI). Refer to the following section for a list of available MMIs.	DVNEF-MX
DEChub 900 MultiSwitch 8-slot intelligent hub chassis.	DMHUB-*xx*
DEChub ONE docking station for standalone use of the VNswitch 900EF module.	DEHUA-*xx* or DEF1H-*xx*

Note: *xx* = Country kit code. Check with your Digital sales representative for the appropriate code for your location.

Optional Modular Media Interface (MMI)

The following optional MMIs are available with the MX version of the VNswitch 900EF:

Media	Order Number
Multimode optics (SC connector)	DEFYM-AA
Single-mode optics (SC connector)	DEFYS-AA
Unshielded twisted-pair (RJ45 connector)	DEFYU-AA

VNswitch 900EF Product Specifications

Parameter	Specification
Environment	
Operating temperature[1]	5°C to 50°C (41°F to 122°F)
Relative humidity	10% to 95% noncondensing
Altitude	
• Operating	Sea level to 2400 m (8,000 ft)
• Non-operating	Sea level to 4900 m (16,000 ft)
Power	46.8 W, total power
	7.8 A, 5 Vdc
	0.12 A, 12 Vdc
	0.52 A, 15 Vdc
Physical	
Height	44.45 cm (17.5 in)
Width	4.45 cm (1.75 in)
Depth	15.25 cm (6 in); 25.40 cm (10.0 in) with a DEChub ONE docking station
Weight	1.94 kg (4.3 lb)[2]
Connectors	12 shielded 10BaseT (MJ-8 RJ45 type) supporting both UTP and ScTP cabling
	2 MMI or ModPMD slots and an OBR port, depending on product version:
	• DVNEF-MM version includes factory-installed ANSI MIC connectors and an OBR port.
	• DVNEF-MX version supports the MMI connectors listed in the previous section.
Shock (Class A/B for products weighing under 100 lb)	10 G / 10 ms half sine pulse in three orthogonal axes
Vibration (Class C)	5 to 200 Hz sine sweep @ 0.25 G limited by 0.5 mm (0.02 in) displacement DA* 200 to 500 Hz sine sweep @ 0.10 G
Certification	CE, CSA, FCC, TÜV, UL, VCCI

Parameter	Specification
Standards Compliancy	IEEE 802.1d, IEEE 802.1h, IEEE 802.1i, IEEE 802.3, Ethernet V1 & V2, ANSI X3/ISO 9314

[1] For sites above 2400 m (8,000 ft), decrease the operating temperature specification by 1.8°C for each 1000 m or 3.2°F for each 3,200 ft.

[2] Actual module weight (no shipping container). Include an additional 1.59 kg (3.5 lb) when attached to a DEChub ONE docking station.

Acoustical Specifications

Product	Sound Power Level L_{WAd}, B	Sound Pressure Level L_{pAm}, dBA (bystander positions)
	Idle/Operate	Idle/Operate
DVNEF	5.0	37
DVNEF + DEHUA	5.3	40
DVNEF + DEF1H	5.4	39

Note: Current values for specific configurations are available from Digital Equipment Corporation representatives. 1B = 10 dBA.

For More Information

Topic	Where to Get More Information
Additional information about the VNswitch's functions and features.	Refer to Chapter 3.
Examples of network configurations using the VNswitch 900EF module.	Refer to Chapter 5 and Chapter 7.

VNswitch 900EA Module

Introduction

The VNswitch 900EA module (Figure 2-10) is a high-throughput, SNMP manageable, multiport network switch for linking LANs with an ATM backbone. The module enables Ethernet LANs to access other Ethernet LANs or ATM hosts across an ATM network. It provides switching for up to 12 Ethernet LANs and 1 high-speed ATM backbone, and transparently converts MAC frames into ATM cells and vice versa.

Figure 2-10: VNswitch 900EA Network Module
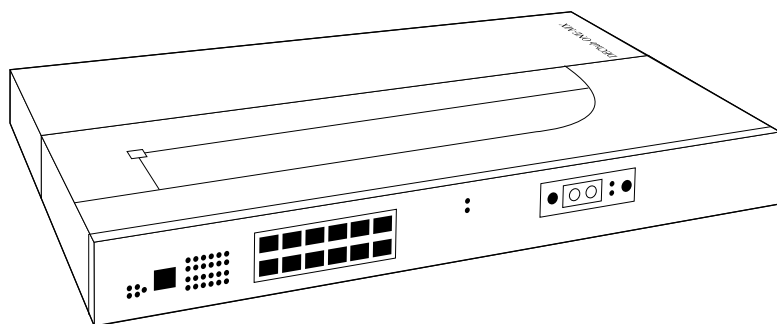
LKG-10247-96F

The VNswitch 900EA module has twelve 10BaseT Ethernet switched ports on the front panel and a single high-speed ATM port. The ATM port incorporates as many as 16 logical ports (as described in the section on ATM Features in Chapter 3).

The unit is auto-configured to connect to the 400 Mb/s VNbus, which provides interoperability across the backplane with other VNswitch modules while maintaining front panel connectivity. Additionally, each of the module's ports (including the ATM port) can be individually redirected to a DEChub 900 backplane LAN segment to create a backplane LAN. Also, port 12 can be redirected to the DEChub 900 backplane ThinWire LAN segment.

When connected to a DEChub ONE docking station (DEF1H or DEHUA), the module's (front panel) port 11 can be optionally assigned to the docking station's AUI port.

The VNswitch 900EA network module is a true backbone network switch that supports complete filtering capabilities (including source address, destination address, protocol type, and VLAN). It includes a large address table (8192 entries) and is fully IEEE 802.1d standards compliant, ensuring the high performance and packet integrity required in large switched networks.

VNswitch 900EA Module

The VNswitch 900EA network module also offers nonvolatile flash memory for software upgrades of the device firmware using the Trivial File Transfer Protocol (TFTP) load protocol (available on a wide variety of platforms including DOS, ULTRIX, UNIX, and OpenVMS operating systems). This eliminates the need to replace or upgrade hardware in the future, and avoids the associated costs to network users.

Like all VNswitch 900 series switches, the VNswitch 900EA network module features integrated SNMP for easy, comprehensive management.

Highlights

Following are the highlights of the VNswitch 900EA module:

- Provides a flexible ATM port using a variety of front-insertable ModPHYs, including:

    — 155 Mb/s (multimode, single-mode, or UTP-5)

    — T3 or E3 coaxial (45 Mb/s or 34 Mb/s)

- ATM 155 Mb/s port can be directed to hub backplane.

- Supports up to 16 ATM Forum LAN Emulation (LANE) clients (also called emulated LANs or ELANs).

- Enables standards-based VLANs over an ATM network.

- Supports up to 16 virtual ports over the single ATM physical port for various combinations of LANE and bridge tunnels. (The default configuration has one FDDI bridge tunnel, one Ethernet bridge tunnel, and one LANE client.)

- Supports flow control on the ATM link (using FLOWmaster).

- Supports up to 1022 virtual circuits (PVCs and SVCs) on the physical ATM port (UNI 3.0 & 3.1).

- Supports four ATM virtual paths (VPI=0, plus three non-zero VPs) for connection to public ATM networks.

- Features a Command Line Interface (CLI) that is accessible from a directly attached terminal or remotely through a Telnet connection.

- Features a built-in SNMP agent for in-band and out-of-band management, whether configured in a single-slot standalone DEChub ONE chassis or installed in a DEChub 900 MultiSwitch enterprise hub.

- Features out-of-band management using SNMP over SLIP and PPP through the OBM connector as an alternative to normal in-band management.

- Features upgradable device firmware (in nonvolatile flash memory) using Trivial File Transfer Protocol (TFTP) with HUBloader or through the setup port with any TFTP server.

- Offers multiple filtering options, including source and destination address and protocol.

- Supports line rate filtering of 14,880 pps on each Ethernet port.

- Supports 8192 network addresses.

- Supports a complete set of standard and private MIBs for full management and configuration capability via SNMP. Supported MIBs include MIB-II (RFC 1213), Bridge MIB (RFC 1493), Ethernet MIB (RFC 1398), AToM MIB (RFC 1695), SONET MIB (RFC 1495), SNMP management (RFC 1157), and Digital enterprise MIBs. (Refer to Appendix A for a complete list of MIBs supported by the VNswitch family of products.)

- Features out-of-band management using SNMP over SLIP through the OBM connector as an alternative to normal in-band management.

- Provides user-configured rate limiting for broadcast and multicast packets by address and specific protocol.

- Provides 802.1d spanning tree loop detection.

- Provides the ability to turn off the spanning tree algorithm on a per-port basis.

- Supports backplane LANs by allowing the 10BaseT ports and the ATM port to be individually directed (through management software) to the DEChub 900 backplane.

Virtual Network Switching with the VNswitch 900EA

Figure 2-11 shows the VNswitch 900EA module used as a high-performance switch connecting multiple LANs and VLANs in an enterprise-wide network. The shaded areas in this figure represent different VLANs. One of the VLANs in this figure spans the ATM cloud and includes an ATM host. This is accomplished using an ATM emulated LAN (ELAN) and is referred to as *standards-based VLANs over ATM*. More information about VLANs and ELANs is provided in Chapter 3.

Figure 2-11: VNswitch 900EA Virtual Network Switch
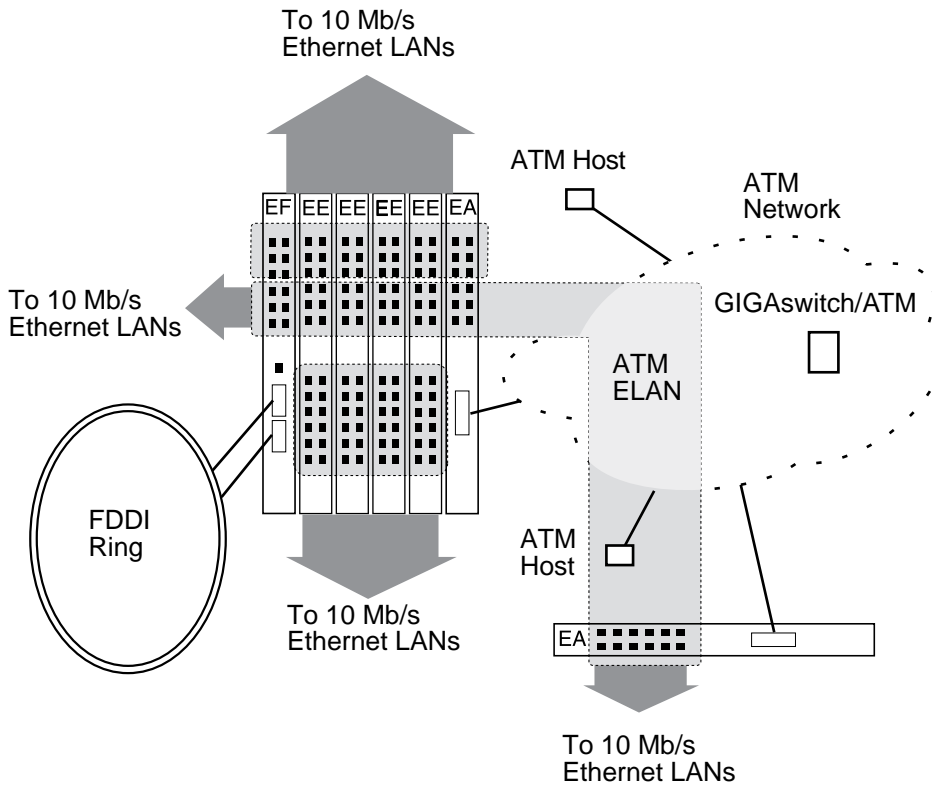
LKG-10253-96F

VNswitch 900EA Ordering Information

Description	Order Number
VNswitch 900EA network module.	DVNEA-MX
DEChub 900 MultiSwitch 8-slot intelligent hub chassis.	DMHUB-*xx*
DEChub ONE docking station for standalone use of the VNswitch 900EA module.	DEHUA-*xx* or DEF1H-*xx*

Note: *xx* = Country kit code. Check with your Digital sales representative for the appropriate code for your location.

VNswitch 900EA Product Specifications

Parameter	Specification
Environment	
Operating temperature[1]	5°C to 50°C (41°F to 122°F)
Relative humidity	10% to 95% noncondensing
Altitude	
• Operating	Sea level to 2400 m (8,000 ft)
• Non-operating	Sea level to 4900 m (16,000 ft)
Power	48.8 W, total power
	8.2 A, 5 Vdc
	0.12 A, 12 Vdc
	0.52 A, 15 Vdc
Physical	
Height	44.45 cm (17.5 in)
Width	4.45 cm (1.75 in)
Depth	15.25 cm (6 in); 25.40 cm (10.0 in) with a DEChub ONE docking station
Weight	1.86 kg (4.1 lb)[2]
Connectors	12 shielded 10BaseT (MJ-8 RJ45 type) supporting both UTP and ScTP cabling
	Single ATM modPHY slot supporting a variety of modPHYs
Shock (Class A/B for products weighing under 100 lb)	10 G / 10 ms half sine pulse in three orthogonal axes
Vibration (Class C)	5 to 200 Hz sine sweep @ 0.25 G limited by 0.5 mm (0.02 in) displacement DA* 200 to 500 Hz sine sweep @ 0.10 G
Certification	CE, CSA, FCC, TÜV, UL, VCCI
Standards Compliancy	IEEE 802.1d, 802.1h, 802.1i, 802.2, IEEE 802.3, Ethernet V1 and V2, LANE 1.0, UNI 3.1

[1] For sites above 2400 m (8,000 ft), decrease the operating temperature specification by 1.8°C for each 1000 m or 3.2°F for each 3,200 ft.

[2] Actual module weight (no shipping container). Include an additional 1.59 kg (3.5 lb) when attached to a DEChub ONE docking station.

Acoustical Specifications

Product	Sound Power Level L_{WAd}, B	Sound Pressure Level L_{pAm}, dBA (bystander positions)
	Idle/Operate	Idle/Operate
DVNEA	5.0	37
DVNEA + DEHUA	5.3	40
DVNEA + DEF1H	5.4	39

Note: Current values for specific configurations are available from Digital Equipment Corporation representatives. 1B = 10 dBA.

For More Information

Topic	Where to Get More Information
Additional information about the VNswitch's functions and features	Refer to Chapter 3.
Examples of network configurations using the VNswitch 900EA module	Refer to Chapter 6 and Chapter 7.

Chapter 3

Features and Functionality

Overview

Introduction

This chapter discusses the technical features and functions of the network modules in the VNswitch 900 series. It also expands on the switching overview given in Chapter 1 and describes the bridging method implemented by the VNswitch. An overview of the VNswitch architecture is also provided.

In This Chapter

The following topics are covered in this chapter:

802.3 LAN Switching Overview

Introduction

All VNswitch 900 series modules operate as 802.3 multiport bridges and implement the *transparent bridging* algorithm. This section provides an overview of 802.3 LAN switching and discusses the bridging method implemented by the VNswitch family. Keep in mind that the term *bridge,* as used in this chapter, can also mean *switch*.

What Is an 802.3 LAN Switch?

All VNswitch 900 series modules are classified as 802.3 LAN switches. An 802.3 LAN switch functions as a multiport bridge that links multiple local area networks (LANs) or LAN segments. Like a bridge, an 802.3 LAN switch accepts data frames from each connected network segment and then decides whether to forward each frame based on the MAC header contained in the frame.

The following summarizes some of the features and functions of 802.3 LAN switches:

- Transparently extend the connectivity between systems attached to LAN segments by forming a single extended LAN.

- Form an extended LAN from segments of different media types.

- Increase the effective throughput capacity of the extended LAN by automatically restricting traffic to only those segments that need to see it.

- Enable partitioning of the extended LAN into separate collision domains for each segment, effectively increasing the throughput capacity. (A collision domain is a shared physical medium with which systems must contend for access.)

- Provide interconnection between systems using nonroutable protocols (such as the Local Area Transport protocol) when those systems are attached to different LAN segments.

Switches and Bridges Compared to Routers

Internetworking devices such as switches, bridges, and routers connect network segments. However, these devices use different methods to establish and maintain their LAN-to-LAN connections. Routers connect LANs at Layer 3 (Network Layer) of the OSI model while switches and bridges connect LANs at Layer 2 (Link Layer).

NOTE

When equipped with the optional routing software, the VNswitch provides switching at *both* layers. This is referred to as *multilayer switching*, and is one of the unique features of the VNswitch family of switches.

Router Connections

Network layer protocols efficiently move information in large and diverse network configurations. Connecting at the Network Layer (Layer 3) with a router allows connectivity and path selection between end stations at large geographical distances. Routing protocols are then used to select the best path for connecting distant and diverse LANs.

Switch and Bridge Connections

Connecting a switch or bridge at Layer 2 provides connectivity across a physical link. This connection is essentially transparent to the host connected on the network.

The Link Layer maintains physical addressing schemes, line discipline, topology reporting, error notification, flow control, and ordered delivery of data frames. Isolation from upper layer protocols is one of the advantages of switching. Because 802.3 bridging switches function at the Link Layer, they do not examine the protocol information that occurs at the upper layers. This results in lower processing overhead and faster communication of network layer protocol traffic. Because switches are not concerned with Layer 3 information, they may also forward different types of protocol traffic (such as IP, DECnet, and IPX) between two or more networks.

Advantages of Bridge Filtering

Switches and bridges filter frames based on Layer 2 fields. This means that the VNswitch may be configured to accept and forward frames of a certain type or ones that originate from a particular LAN segment. This ability to configure filters is very useful for maintaining effective traffic flow.

Multilayer Switching

All VNswitch modules are designed for *multilayer switching*, that is, switching at *both* Layer 2 *and* Layer 3. Multilayer switching is enabled by ordering the optional routing software which can be downloaded to the VNswitch. The optional routing software is introduced later in this chapter.

Basic Bridge or Switch Operation

According to the IEEE 802.3 LAN standard, all station addresses are specified at the Media Access Control (MAC) level. Therefore, the MAC level is the level at which a bridge or 802.3 LAN switch functions. The 802.3 bridges and LAN switches interconnect separate LANs by relaying data frames between each LAN's MAC entity.

VNswitch 900 series switches perform additional functions such as modifying frames, and recognizing and transmitting frames from different LAN technologies (for example, FDDI and Fast Ethernet).

Transparent Bridging

All modules in the VNswitch 900 series implement the Spanning Tree Protocol. A device that implements the Spanning Tree Protocol is commonly called a *transparent bridge* or a Spanning Tree Bridge (STB).

The term *transparent* refers to the fact that the bridge silently forwards non-local traffic to attached LANs in a way that is transparent to or unseen by the originating and receiving systems. End station applications do not know about the presence of the bridge.

NOTE

All modules in the VNswitch 900 series implement transparent bridging conforming to the IEEE 802.1d standard. All Ethernet bridges and switches within the same network must be 802.1d spanning tree bridges. This version of the Spanning Tree Protocol is not compatible with bridges that implement Digital's LAN Bridge 100 spanning tree protocol.

How Does Transparent Bridging Work?

A bridge that implements the Spanning Tree Protocol learns about the presence of end stations by listening to traffic. From this listening process, it builds a database of addresses of end stations attached to its LANs.

For each frame it receives, the bridge checks the frame's destination address against the ones in its database. If the frame's destination is an end station on the same LAN, the frame is not forwarded. If the destination is on another LAN, the frame is forwarded. If the destination address is not present in the database, the frame is forwarded to all the LANs that are connected to the bridge (except the LAN from which it originated).

Function of Transparent Bridging or the Spanning Tree Protocol

> In a mesh topology where more than one bridge or switch is connected between two LANs, looping can occur. In such cases, data packets bounce back and forth between the two LANs' parallel bridges. This creates a redundancy in data traffic and produces the phenomenon known as looping.
>
> To prevent looping, all VNswitch modules use the Spanning Tree Protocol and algorithm. The spanning tree algorithm produces and maintains a loop-free topology in a bridged network that contains loops in its physical design. With spanning tree, a self-configuring algorithm allows a bridge to be added anywhere in the LAN without creating loops. Upon adding the new bridge, the spanning tree transparently reconfigures all bridges on the LAN into a single loop-free spanning tree.

Multiple Spanning Trees

> Unlike ordinary 802.3 LAN switches, the modules in the VNswitch family implement *multiple spanning trees*. Each port-based virtual LAN (VLAN) on one or more VNswitch modules has its *own* separate spanning tree process, logically isolating it from other ports or port groups, and thus ensuring the highest degree of stability and security between VLANs. As many as 192 spanning trees can be implemented within a DEChub 900 configuration of 8 VNswitch 900EE modules.

Where to Get More Information

> For more information about transparent bridging and the spanning tree protocol, refer to Appendix B. Appendix B describes transparent bridging operation and lists all of the spanning tree parameters, including those that can be set using management software.

VNswitch Architecture

Introduction

The VNswitch 900 series products are designed with a modular architecture, consisting of a main board and technology-specific secondary boards. This section provides an overview of the architecture of the main board and each secondary board. It also discusses the features of the DEChub 900 backplane that provide connectivity for the VNswitch modules.

Main Board

The main board contains the following three processing engines, each optimized to deliver high performance for a specific set of networking or management tasks:

- Application Processor (AP)

- Fastpath Processor (FP)

- Data Moving Engine (DME)

Each processing engine has its own dedicated memory subsystem designed for the speed of the processor and its unique processing requirements.

Figure 3-1 illustrates the internal architecture of the main board. Each of the components shown in the figure is described in the following paragraphs.

Figure 3-1: VNswitch 900 Series Main Board Architecture
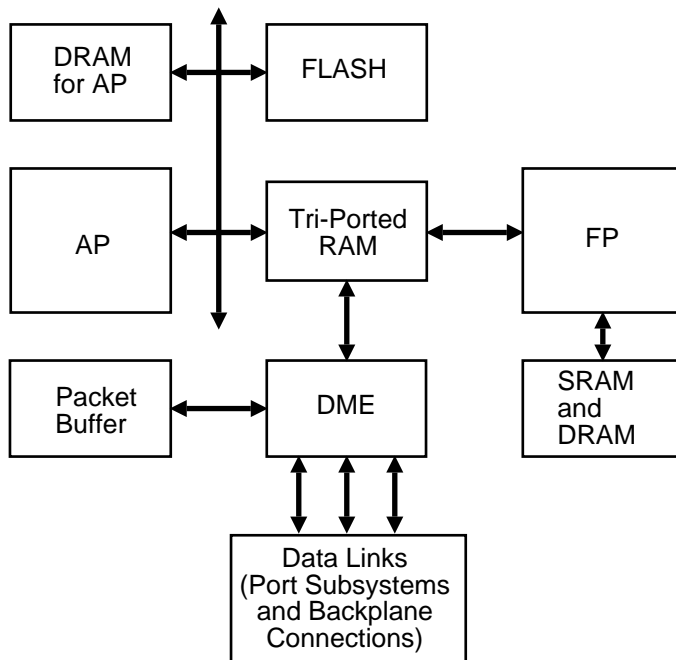
LKG-10248-96F

Application Processor (AP)

The AP is a high-performance 32-bit CISC processor. It is used for management interfaces (SNMP and Telnet), 802.1d spanning tree operation, and other high-level functions in the switch.

For modules equipped with the optional routing software upgrade, the AP also performs RIP and OSPF route calculations.

Fastpath Processor (FP)

The FP is a high-performance 32-bit RISC processor. It is used for multicast bridging, assisting with advanced VLAN features, FDDI-to-Ethernet translation bridging, and IP fragmentation for FDDI.

For modules equipped with the optional routing software upgrade, the FP also provides the high-performance packet forwarding engine for Layer 3 routing.

Data Moving Engine (DME)

The DME is a sophisticated custom application-specific integrated circuit (ASIC) designed by Digital to perform unicast transparent bridge forwarding (that is, switching), advanced bridge filtering, and core VLAN functions at very high speeds. It includes three datalink ports, two of which attach to the various datalink subsystems implemented on the secondary boards. The third datalink port is dedicated to the VNbus interface.

Memory Subsystems

Each main processing engine has its own memory subsystem, which is described in this section.

AP Memory Subsystem

The AP (Application Processor) is equipped with 8 Mbytes of DRAM in the base configuration (upgradable to 32 Mbytes) plus 5 Mbytes of nonvolatile FLASH memory.

The DRAM is used for data structures required by the bridging and routing application code. The DRAM is also used for storing performance statistics, alarms, and system event data.

The FLASH memory is used to store the module's configuration, and to store and run the operational code.

FP Memory Subsystem

The FP (Fastpath Processor) is equipped with 1 Mbyte of DRAM and 0.5 Mbytes of high-speed SRAM. These are used for running the code needed to conduct Layer 3 packet forwarding and filtering operations, and operations related to advanced Layer 2 functions.

DME Memory Subsystem

The DME (Data Moving Engine) has a dedicated 4 Mbyte packet buffer shared dynamically across all ports on the switch. This is used for speed matching (between different links) and congestion control.

NOTE

Unlike the VNswitch 900 series switches, ordinary switches are typically designed with per-port buffers (sometimes as low as 192 Kbytes per port), and are unsuitable for high-performance applications that require switching between 10 Mb/s Ethernet and a high-speed link such as 100 Mb/s Ethernet, FDDI, or ATM.

Common Memory System

All three major processing engines share a common high-speed, triple-ported RAM used for communications between all processor and memory subsystems.

Secondary Boards

The secondary boards comprise the various datalink subsystems implemented by the VNswitch 900 series products. These port subsystems include the following:

- 10 Mb/s Ethernet port subsystem

- 100 Mb/s Ethernet (Fast Ethernet) port subsystem

- FDDI port subsystem

- ATM port subsystem

All VNswitch modules are comprised of two secondary boards. With the exception of the VNswitch 900EE module (which has two 10 Mb/s Ethernet port systems), each module includes a 10 Mb/s port system and one of the high-speed datalink port systems (100 Mb/s Ethernet, FDDI, or ATM). Each of these subsystems is described in the following sections.

10 Mb/s Ethernet Port Subsystem

The Ethernet port subsystem is illustrated in Figure 3-2. As shown in Figure 3-2, this subsystem consists of the MACMUX and a bank of 12 MACEs. Each MACE is a slave MAC device that is polled by the MACMUX.

The MACEs interface to external cables through 12 MJ-8 connectors. MACE 12 is also used to connect to the backplane ThinWire segment. MACE 11 is used to attach to the docking station AUI port.

The MACEs attach to the 12 backplane Ethernet LAN segments by means of the following ASICs:

- A single IMBI ASIC is used for attaching to backplane LAN segments 1–6.

- Four SIMBI ASICs are used for attaching to backplane LAN segments 7–12.

Figure 3-2: Ethernet Port Subsystem

LKG-10249-96F
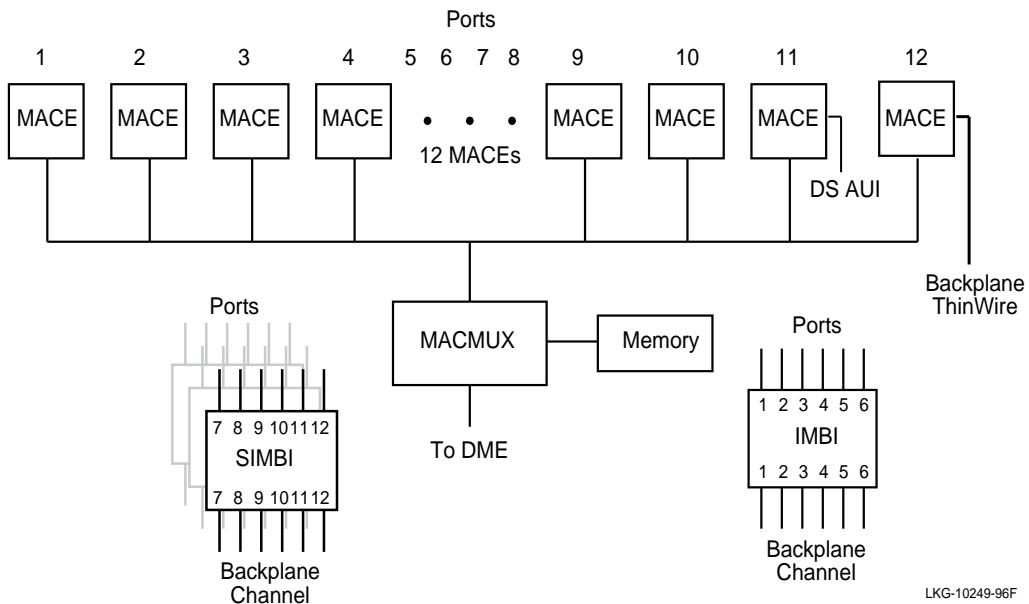
Fast Ethernet Port Subsystem

The Fast Ethernet port subsystem, shown in Figure 3-3, includes two SEEQ 10/100 EtherMACs tied to external FIFOs and controlled by a Xilinix FPGA and PALS. Each SEEQ MAC attaches to an MMI connector through its MII interface. The MMI connector supports a set of modules that provide fiber and UTP connections.

Figure 3-3: Fast Ethernet Port Subsystem
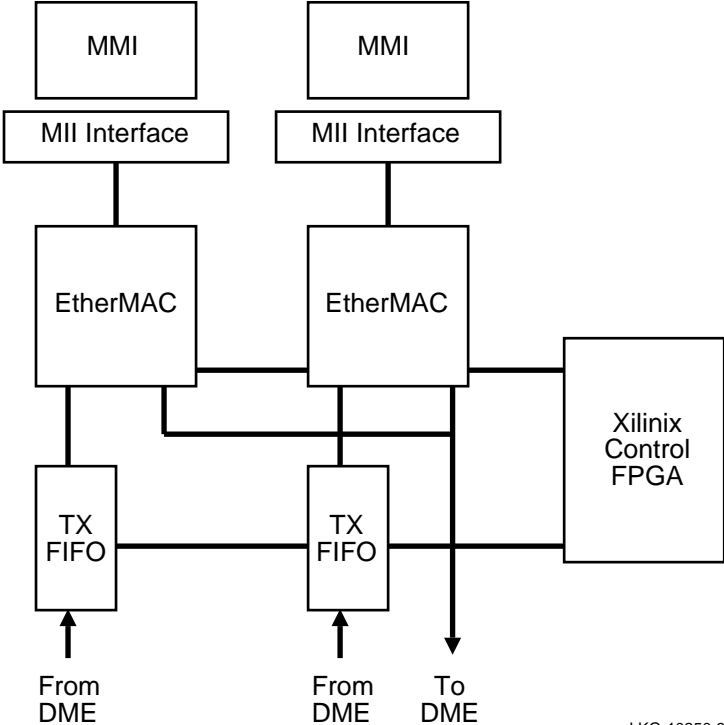
LKG-10250-96F

FDDI Port Subsystem

The FDDI port subsystem, shown in Figure 3-4, consists of the MELMAC and ELM ASICs, along with the DMSS backplane interface chip. The dual Mod-PMD or MMI connectors allow the switch to be configured with two external DAS ports.

Figure 3-4: FDDI Port Subsystem
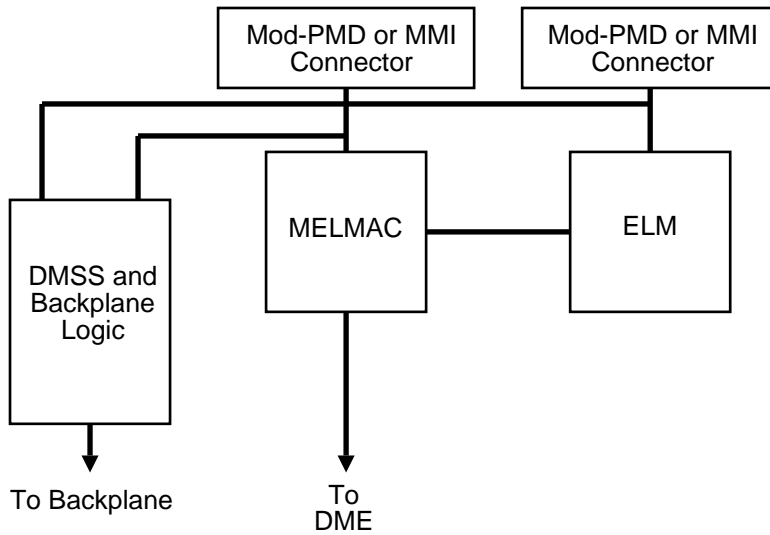
LKG-10251-96F

ATM Port Subsystem

The ATM subsystem, shown in Figure 3-5, consists of an ATOM3+ Segmentation and Reassembly (SAR) ASIC with its cell memory. The chip attaches to either a Mod-PHY over a connector or to a local SUNI chip for attachment to the backplane. The ATOM3+ chip can operate up to 155 Mb/s full duplex.

Figure 3-5: ATM Port Subsystem
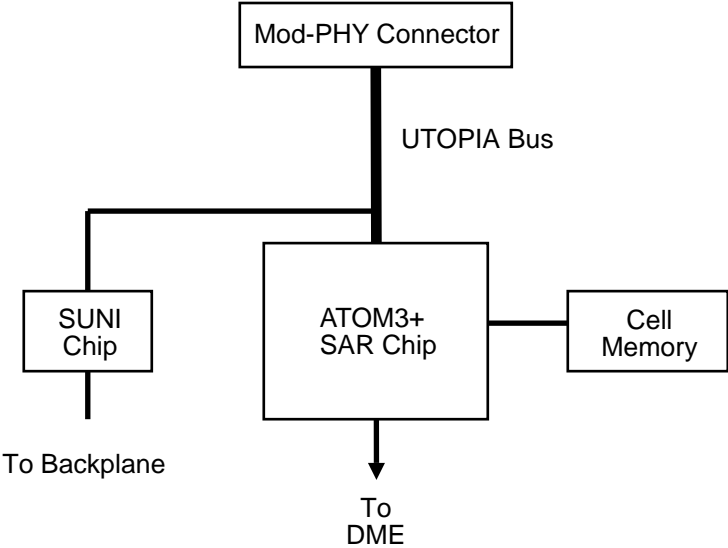
LKG-10252-96F

DEChub 900 Backplane Connections

The VNswitch modules connect to the DEChub 900 backplane by means of two connectors: a 48-pin backplane connector and a 160-pin matrix backplane connector. Together, these connectors provide the following features and functionality:

- One dedicated Ethernet ThinWire segment — The dedicated ThinWire segment provides backward compatibility for DEChub 90 modules.

- Multiple flexible channels (or LAN segments) — The flexible channels (also called *flex channels*) are technology-independent high bandwidth channels that support the interconnect needs of the VNswitch modules and other DEChub 900 modules.

- Four channels for ATM on the backplane — Four channels are provided for connecting the VNswitch 900EA module to an ATMswitch 900 on the backplane.

- Management Agent Module (MAM) — The MAM provides the management interface to the modules installed in the backplane. The modules are linked to the MAM by a full-duplex, point-to-point serial link. The functions of the MAM are accessed by means of either a setup port or OBM port on the DEChub 900.

The DEChub 900 backplane connections are illustrated in Figure 3-6.

Figure 3-6: Backplane Connections
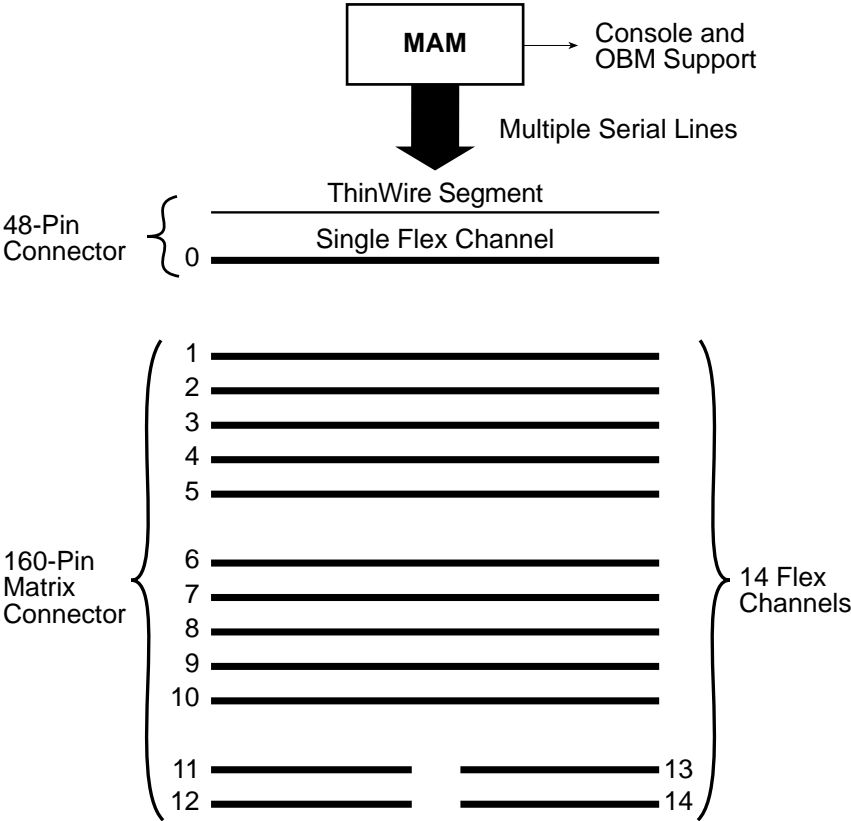
LKG-10202-96F

Channel Assignment

As shown in Figure 3-6, a total of 15 flexible channels (0–14) can be used by the VNswitch modules. The DEChub 900 automatically assigns data links to the flexible channels to obtain maximum bandwidth. No intervention is necessary on the part of the network manager. However, certain flexible channels are currently used by specific technologies, as illustrated in Figure 3-7.

Figure 3-7: Flex Channel Assignment
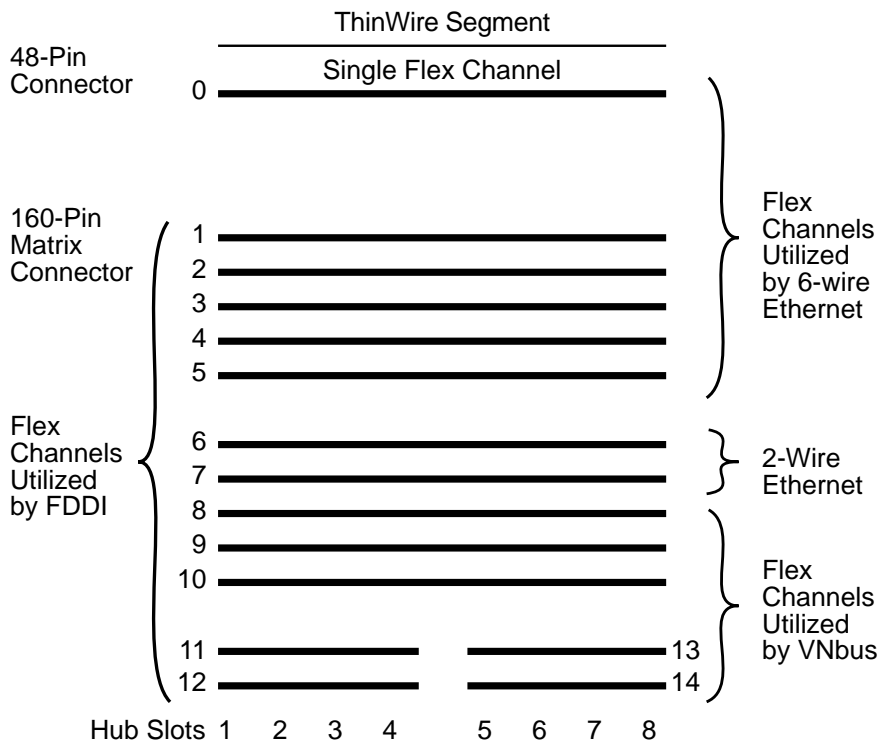
LKG-10201-96F

NOTE

Four channels are provided for ATM on the backplane. These channels are not illustrated in the above figure.

As shown in Figure 3-7, flexible channels 0–5 are used for the 6-wire Ethernet protocol (IMB-6). Channels 6 and 7 are reserved for the 2-wire Ethernet protocol (IMB-2), and support three links each, adding an additional six Ethernet LAN segments to the backplane. Thus, flexible channels 0–7 provide a total of 12 separate Ethernet LAN connections.

FDDI is supported on channels 1–14; however, a maximum of four FDDI rings are allowed.

ATM is supported on four independent channels (not illustrated), and the VNbus (described in the following section) is implemented on channels 8, 9, and 11–14.

Figure 3-8 illustrates how the ports on the VNswitch modules are mapped to the various flexible channels.   The number of data links allowed for each flexible channel is listed in Table 3-1.

Figure 3-8: Mapping of VNswitch Ports to Flexible Channels

LKG-10203-96F
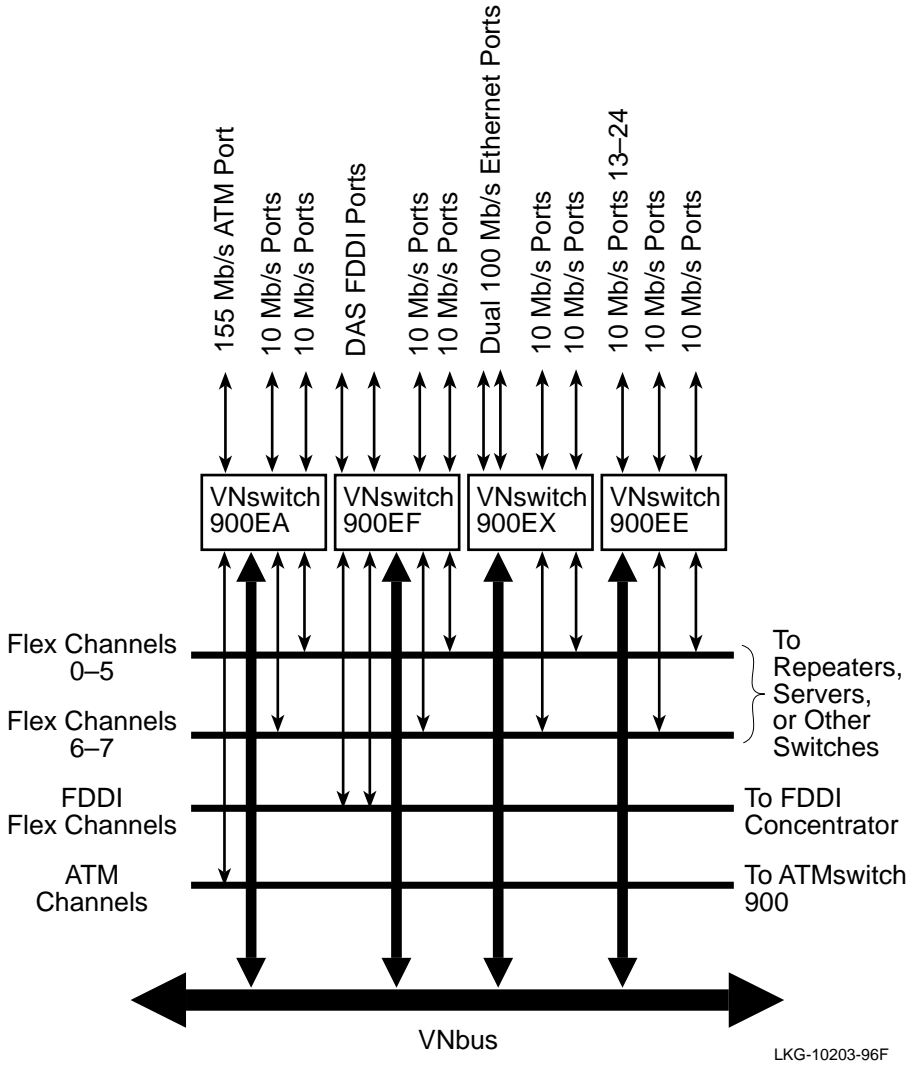
NOTE

The ATM channels are not included in the following table because they differ
from the flexible channels used by the other modules.

Table 3-1: Data Links Allowed Per Flexible Channel

Flex Channel No.	Number of Data Links Allowed		
	VNbus	Ethernet	FDDI
0	0	1	0
1	0	1	1
2	0	1	1
3	0	1	1
4	0	1	1
5	0	1	1
6	0	3	1
7	0	3	1
8	1	0	1
9		0	1
10	0	0	1
11	1	0	1
12		0	1
13	1	0	1
14		0	1
Total Data Links:	3*	12	14**

* Two flex channels per VNbus. **Maximum of 4 FDDI rings.

Multiple technologies cannot be used at the same time on a single flexible channel. For example, if the VNbus is implemented on channels 8 and 9, these two channels cannot be used by FDDI. Similarly, if channels 6 and 7 are used by FDDI, neither of these channels can be used at the same time by an Ethernet LAN.

Virtual Network Bus (VNbus)

Introduction

All VNswitch modules can be joined by a high-speed bus called the *VNbus*, which is implemented on the DEChub 900 backplane. Three VNbuses are supported, each providing 400 Mb/s of bandwidth, for a total capacity of 1.2 Gb/s.

What Is the VNbus?

The VNbus is a high-speed, *virtual network bus* (VNbus) that can be used to connect VNswitch modules on the DEChub 900 backplane. Refer to Figure 3-9.

Figure 3-9: VNswitch Modules Connected by the VNbus
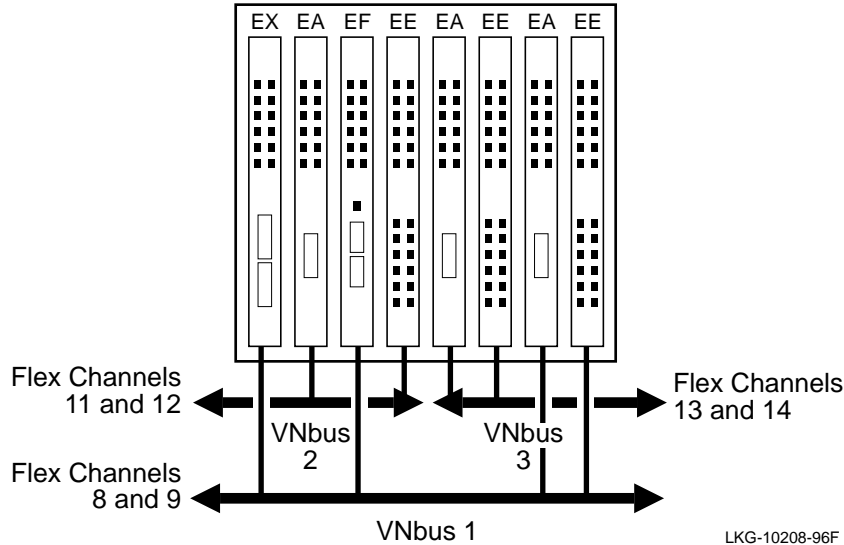
As shown in Figure 3-9, three distinct VNbuses are supported. The VNbuses are physically isolated from each other and each forms a separate network.

NOTE

The separate networks attached to the different VNbuses can be connected by means of a high-speed link (FDDI, ATM, or Fast Ethernet). This is accomplished by bridging modules together across the front panel.

VNbus 1 (implemented on flexible channels 8 and 9) spans the entire hub and can connect modules from any of the eight slots in the hub. The other two VNbuses span half the width of the hub: VNbus 2 (implemented on channels 11 and 12) spans across slots 1–4, and VNbus 3 (implemented on channels 13 and 14) spans across slots 5–8.

The default configuration is for all VNswitch modules to connect to VNbus 1. Thus, when first inserted into a DEChub 900 slot, the modules have the configuration illustrated in Figure 3-10.

Figure 3-10: VNswitch Default Configuration

400 Mb/s VNbus
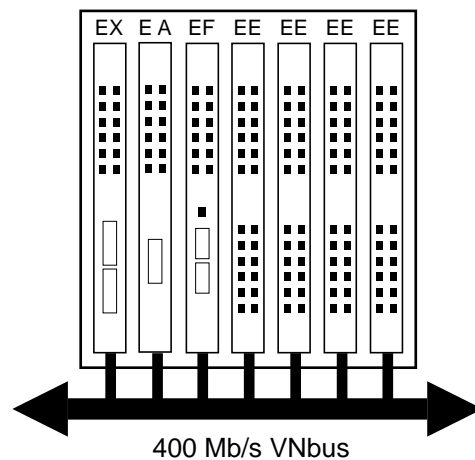
LKG-10209-96F

Through software commands, a module can be disconnected from VNbus 1 and connected to another VNbus, or it can be disconnected from all three VNbuses.

What Is the Function of the VNbus?

The VNbus provides a means for creating a high-speed LAN that connects the VNswitch modules, without having to redirect any front panel ports to the backplane. Thus, all front panel ports remain active.

NOTE

Individual ports can still be redirected to the backplane to form backplane LANs, as described later in this chapter. For more information, see the section Backplane LANs.

An important function of the VNbus is to allow virtual LANs (VLANs) to span across the modules in the hub as illustrated in Figure 3-11. Each VNbus supports 63 VLANs.

Figure 3-11: VLANs Spanning Across Modules Using the VNbus
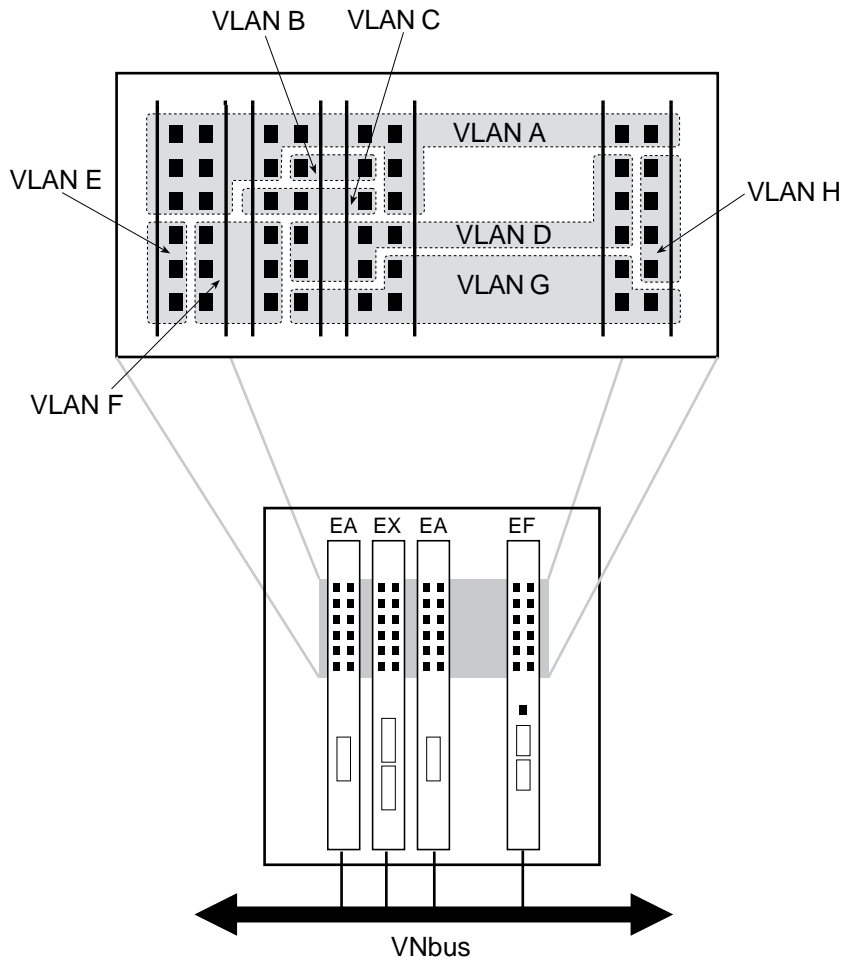
LKG-10237-96F

Virtual Network Bus (VNbus)

Another important function of the VNbus is to allow a VNswitch module equipped
with the optional routing software to serve as a router for one or more other VNswitch
modules in the hub, as shown in Figure 3-12.

Figure 3-12: Single Router Serving All Modules Using the VNbus

R = Optional Routing Software
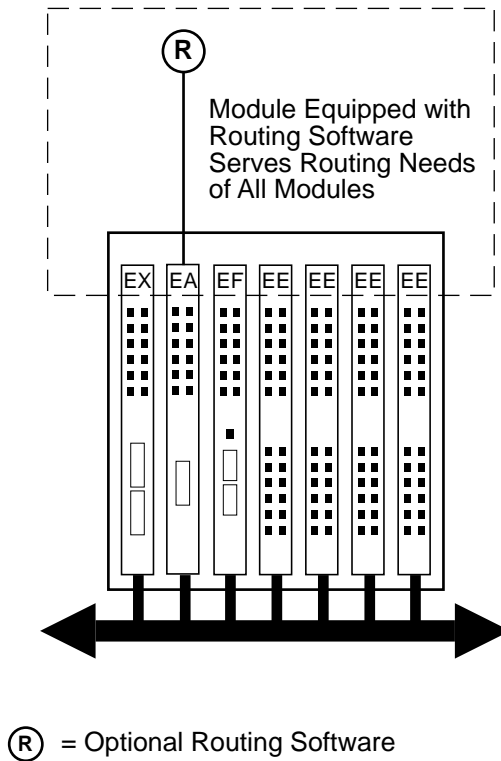
LKG-10207-96F

Summary of VNbus Features

The VNbus features a low latency of less than 3 microseconds, and includes fairness guarantees for all modules on the bus. For example, to ensure that small frames do not have to wait for transmission when larger frames are sent over the VNbus, large frames are broken into fragments before they are transmitted.

Also, if several VNswitch modules attached to the same VNbus transmit at the same time, automatic flow control mechanisms are implemented to ensure that no frames are lost.

The following summarizes the features of the VNbus:

- Provides connectivity between all VNswitch modules

- Allows bridging or routing between all four technologies (Ethernet, Fast Ethernet, FDDI, and ATM)

- Provides up to 1.2 Gb/s of bandwidth (using three VNbuses) on the backplane for exclusive use by VNswitch modules

- Allows VLANs to span across modules in the hub

- Fully compatible with existing the DEChub 900 chassis—no hardware upgrade to the chassis is required

- Features low latency of less than 3 microseconds

- Includes automatic flow control mechanisms

- Ensures low latency by breaking up large frames into smaller fragments

- Supports simultaneously both Ethernet and FDDI frame formats without performing any translation

Frame Formats

Introduction

This section describes the formats of the various frames that are processed by the VNswitch 900 series switches. The following formats are described:

- Ethernet frame formats

- FDDI frame format

- Raw IPX frame format

- VNswitch 900EA frame formats, including:

  — ATM cell format

  — AAL5 format

  — LAN Emulation (LANE) frame format

  — Bridge tunnel formats

Ethernet Frame Formats

The 10 Mb/s Ethernet and Fast Ethernet process the same two types of Ethernet frames: 802.3 (CSMA/CD) frames and the original Ethernet V2.0 frames.

802.3 Ethernet Frame Format

Figure 3-13 illustrates the format of 802.3 Ethernet frames processed by all VNswitch 900 modules.

Figure 3-13: 802.3 Ethernet Frame Format

802.3 Frame Format

6 octets	6 octets	2 octets	1 octet	1 octet	1 octet	2 octets	variable	4 octets
DA	SA	Length	DSAP	SSAP	CTL	PID	Data	FCS
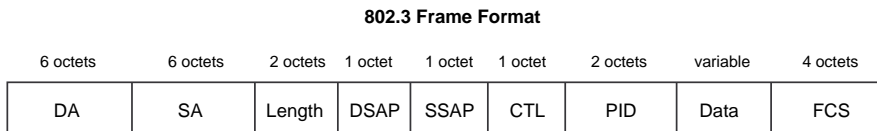
LKG-10235-96F

Table 3-2 defines the 802.3 Ethernet frame fields.

Table 3-2: 802.3 Ethernet Frame Field Descriptions

Field	Description	Size (Octets)
DA	Destination Address	6
SA	Source Address	6
Length	Length of contents that follow, up to (but not including) the FCS field	2
DSAP	Destination Service Access Point	1
SSAP	Source Service Access Point	1
CTL	802.2 control field	1
PID	Protocol Identifier	2
Data	User data	46–1500
FCS	Frame Check Sequence	4

Ethernet V2.0 Frame Format

Figure 3-14 illustrates the format of Ethernet V2.0 frames processed by all VNswitch 900 modules.

Figure 3-14: Ethernet V2.0 Frame Format

Ethernet V2 Frame Formats

6 octets	6 octets	2 octets	variable	4 octets
DA	SA	PT	Data	FCS
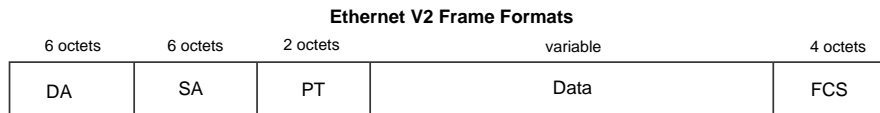
LKG-10236-96F

Table 3-3 defines the Ethernet V2.0 frame fields.

Table 3-3: Ethernet V2.0 Frame Field Descriptions

Field	Description	Size (Octets)
DA	Destination Address	6
SA	Source Address	6
PT	Protocol Type	2
Data	User data	41–1495
FCS	Frame Check Sequence	4

FDDI Frame Format

The format of FDDI frames processed by the VNswitch 900EF module is shown in Figure 3-15.

Figure 3-15: FDDI Frame Format

FDDI Frame Format

1 octet	6 octets	6 octets	1 octet	1 octet	1 octet	2 octets	variable	4 octets
FC	DA	SA	DSAP	SSAP	CTL	PID	Data	FCS

LKG-10234-96F

Table 3-4 describes the FDDI frame fields.

Table 3-4: FDDI Frame Field Descriptions

Field	Description	Size (Octets)
FC	Frame Control	1
DA	Destination Address	6
SA	Source Address	6
DSAP	Destination Service Access Point	1
SSAP	Source Service Access Point	1
CTL	802.2 control field	1
PID	Protocol Identifier	2
Data	User data	0–4473
FCS	Frame Check Sequence	4

Raw 802.3 IPX Frame Format

Novell stations source two types of IPX frames on Ethernet LANs:

- The normal Ethernet V2.0 format with a Protocol Type (PT) of 81–37

- Raw 802.3 frame

NOTE

Two other formats (not discussed here) are rarely used but are also sourced by some IPX stations: 802 SAP and 802 SNAP.

The Raw 802.3 format violates IEEE specifications because it does not contain the standard 802.2 LLC required header. This format is known to cause problems in extended LANs with FDDI components.

Although the VNswitch supports both the Raw 802.3 format and the Ethernet V2.0 format, it is recommended, whenever possible, to configure IPX stations to use the Ethernet V2.0 format (see Figure 3-16).

Figure 3-16: Novell IPX Ethernet V2.0 Frame Format

6 octets	6 octets	2 octets	variable	4 octets
DA	SA	01 - 32	Data	FCS

PT
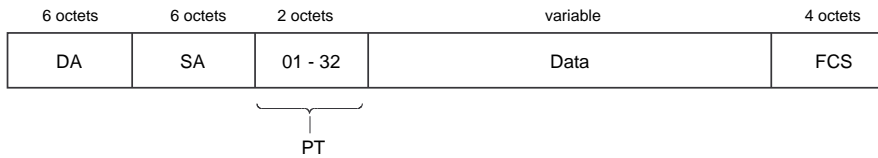
LKG-10132-96F

As shown in Figure 3-17, the Raw 802.3 frame does not have the 802.2 LLC header. Instead, after the length field on the 802.3 header, it has the IPX (network layer) header.

Figure 3-17: Novell IPX Raw 802.3 Frame Format

6 octets	6 octets	2 octets	1 octet	1 octet	variable	4 octets
DA	SA	Length	FF	FF		FCS

IPX Data
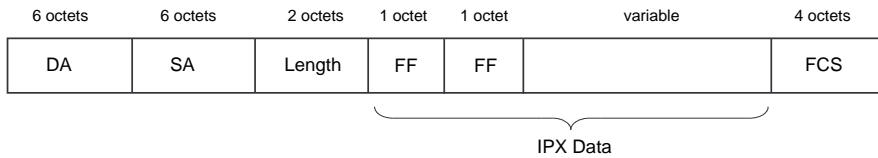
LKG-10133-96F

The first two octets of the IPX header can be used to identify a Raw 802.3 IPX frame. The first two octets of the IPX header (checksum field) are always FF-FF (checksum is disabled in IPX).

VNswitch 900EA Frame Formats

The VNswitch 900EA translates ATM cells into LAN frames, and LAN frames into ATM cells. This is accomplished by using the ATM interface and service protocol known as the ATM Adaption Layer (AAL). Specifically, the VNswitch 900EA uses an implementation of AAL known as AAL5.

This section describes the format of an ATM cell and the AAL5 trailer that is used when translating cells into LAN frames. It also describes the formats of LAN frames when transported over the ATM network using LAN Emulation (LANE) and bridge tunnels.

ATM Cell Format

Figure 3-18 illustrates the structure of an ATM cell. As shown in the figure, each cell contains 53 bytes of both control and user information.

Figure 3-18: ATM Cell Format


```
┌──────────────┬──────────────┐ ╲
│ Flow Control │     VPI      │  ╲
├──────────────┼──────────────┤   ╲
│     VPI      │     VCI      │    ╲
├──────────────┴──────────────┤     ╲  Header
│            VCI              │      ╲  5 bytes
├──────────┬─────────┬────────┤      ╱
│   VCI    │ Payload │  CLP   │     ╱
│          │ Type    │        │    ╱
├──────────┴─────────┴────────┤   ╱
│    Header Error Control      │  ╱
├─────────────────────────────┤ ╱
│      User Information       │ ╲ Payload
│   (voice, data, video)      │ ╱ 48 bytes
└─────────────────────────────┘
```

LKG-10259-86F
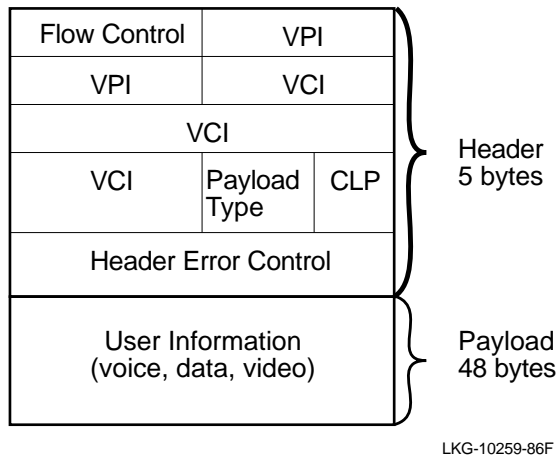
An ATM cell is broken into two main sections, the header and the payload.

- Header — 5 bytes of addressing information

- Payload — 48 bytes of user information (voice, data, or video)

Header

Each ATM cell header includes two numbers that identify the connection and destination addresses. These identifiers are called the Virtual Path Identifier (VPI) and the Virtual Circuit Identifier (VCI). The VPI/VCI permits cells to be addressed to several destinations by establishing multiaddress virtual circuits.

Payload

The payload portion of the ATM cell contains the actual user data.

AAL5 Trailer Format

Figure 3-19 shows the AAL5 trailer that is used when transporting LAN frames over the ATM network by means of LAN Emulation (LANE) and bridge tunnels. The AAL5 trailer is appended to LANE and bridge tunnel frames as illustrated in Figures 3-20 through 3-22.

Figure 3-19: AAL5 Trailer Format

Data	variable
PAD	
Reserved	2 octets
Length	2 octets
CRC	4 octets

LKG-10301-86F

LAN Emulation Frame Format

The VNswitch 900EA supports the LANE format specified in the ATM Forum's *LAN Emulation Over ATM V1.0* document. The LAN Emulation frame format is shown in Figure 3-20.

Figure 3-20: LAN Emulation Frame Format

LAN Emulation Header (LEC ID)	2 octets
DA (Destination Address)	6 octets
SA (Source Address)	6 octets
Type/Length	2 octets
Data	variable
PAD	
Reserved	2 octets
Length	2 octets
CRC	4 octets

LKG-10260-86F

Bridge Tunnel Formats

The VNswitch 900EA supports bridge tunnels, as specified in RFC 1483. The format of an FDDI frame transported over the ATM network through a bridge tunnel is shown in Figure 3-21. Figure 3-22 shows the format of an Ethernet frame transported over the ATM network through a bridge tunnel.

Figure 3-21: FDDI Bridge Tunnel Format

Null	2 octets
FC	1 octet
DA	6 octets
SA	6 octets
DSAP	1 octet
SSAP	1 octet
CTL	1 octet
PID	2 octets
Data	variable
PAD	
Reserved	2 octets
Length	2 octets
CRC	4 octets

LKG-10296-86F

Figure 3-22: Ethernet Bridge Tunnel Format

Field	Size
Null	2 octets
DA	6 octets
SA	6 octets
PT	
Data	variable
PAD	
Reserved	2 octets
Length	2 octets
CRC	4 octets

LKG-10297-86F

Translation

Introduction

Unlike proprietary encapsulating switches (that require a decapsulating device at the receiving end), all VNswitch 900 series switches guarantee interoperability and transparency to upper level protocols by creating standardized frames and performing translation of frames between all interconnected LANs.

Transparent Translation

The VNswitch 900 series switches translate all frames that they forward between dissimilar interconnected LANs. Frames destined to be forwarded are first converted into the native frame format of the destination LAN. This allows all stations on an FDDI LAN to communicate transparently with all stations on an 802.3/Ethernet LAN. Similarly, stations on an 802.3/Ethernet LAN can communicate across an ATM network with stations on another 802.3/Ethernet LAN.

Ethernet-to-FDDI Translation

Translation from the IEEE 802.3 format to the FDDI frame format is fairly straightforward, since both FDDI and 802.3 frames contain an 802.2 Logical Link Control (LLC) header. However, when translating a frame from the Ethernet V2.0 format to the FDDI format, an 802.2 LLC header must be generated that conforms to Internet standard RFC 1390. Figure 3-23 provides an example.

Figure 3-23: Ethernet-to-FDDI Translation Example
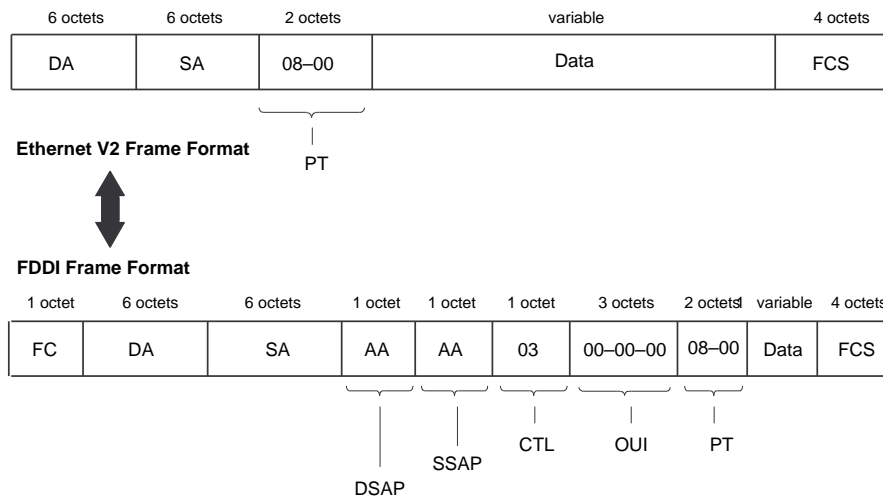
LKG-10130-96F

NOTE

AppleTalk ARP (one of the protocols related to AppleTalk) presents a special case that deviates from the descriptions provided in this section (refer to the section AppleTalk Translation).

A standard 802.2 LLC header is created for Ethernet V2.0 format frames to be sent on FDDI as follows:

- Both the Destination Service Access Point (DSAP) and Source Service Access Point (SSAP) values are set to Sub-Network Access Protocol (SNAP) AA (hexadecimal).

- The Control field is set to U_I (Unnumbered Information), which has a value of 03.

- The Organizationally Unique Identifier (OUI), the first 3 octets of the Protocol Identifier (PID), is set to 00-00-00 (the Ethernet OUI).

- The remaining 2 octets in the PID are set to the Ethernet Protocol Type (PT).

Translation

When translating a frame from the FDDI format to the Ethernet V2.0 format, the process is reversed to allow the generation of an Ethernet V2.0 format frame from an FDDI frame.

VNswitch 900 series switches ensure that all 802.3/Ethernet frames transferred between pairs of (Digital) switches are retranslated into their original format. A table in each switch contains the PIDs of special protocols that could otherwise be retranslated into the wrong format. The AppleTalk PID is entered into the table by default.

IP Fragmentation

The VNswitch 900EF Ethernet-to-FDDI switch performs IP fragmentation on IP frames received on the FDDI that are to be forwarded on the Ethernet, but are larger than the maximum frame size that Ethernet supports.

The maximum FDDI frame size, including the CRC, is 4495 octets. The maximum Ethernet frame size is 1518 octets. RFC 791, the standard that describes the Internet Protocol, specifies the rules for fragmentation when there is a mismatch in maximum data link size between the source and destination data link.

As shown in the example (Figure 3-24), the VNswitch 900EF switch breaks up (fragments) the received frame into legal Ethernet frames in accordance with this specification. An IP header is generated for each of the fragments that are generated, and the destination node reassembles the fragments when received.

Figure 3-24: IP Fragmentation Example

IP Frame on FDDI with DF Bit Clear | Ethernet

| 4352 data octets | Fragment → | 1500 data octets |
| + |
| 1500 data octets |
| + |
| 1352 data octets |
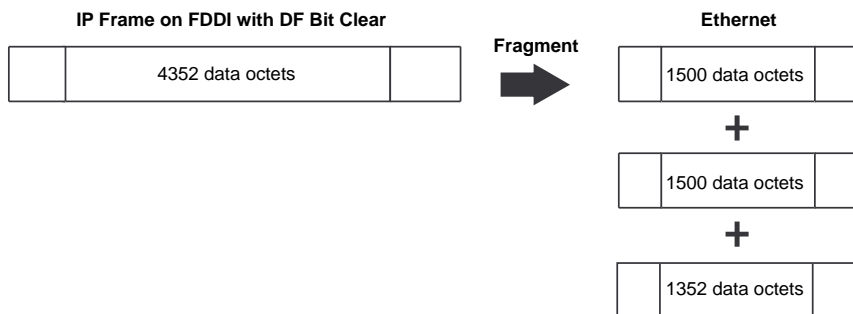
LKG-10131-96F

Note that IP fragments are generated only when the IP fragmentation option is enabled (by the management software), the Don't Fragment (DF) bit in the IP header is set to Clear, and the IP header is error free.

NOTE

The DF bit is set by the station originating the frame.

By default, IP fragmentation is *enabled*. If IP fragmentation is not occurring, the reason for dropping the IP frames can be detected by examining the IP fragment-related counters. See the documentation on your network management software for more information.

IPX Translation

IPX translation (Figure 3-25) involves translating raw 802.3 IPX frames on the Ethernet LAN to SNAP encapsulated frame (RFC 1390) containing a PT of 81-37 on the FDDI side (the opposite is applied from FDDI to Ethernet).

Figure 3-25: Frame Formats When IPX Translation Is Enabled
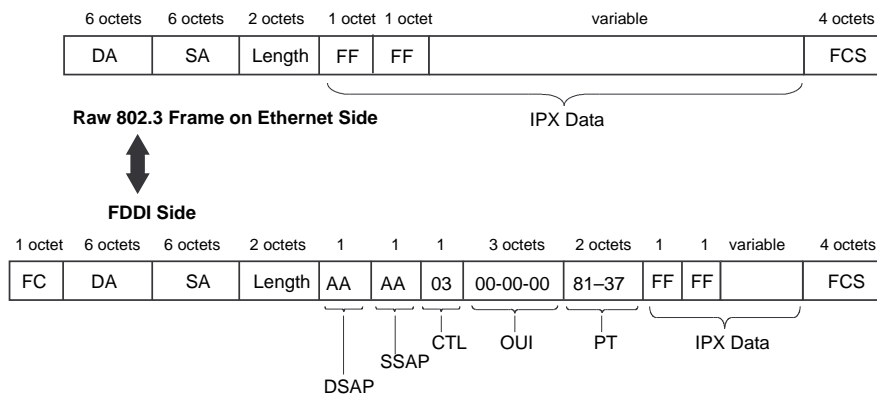
LKG-10134-96F

When the switch receives an 802.3 frame while DSAP filtering, it checks the frame's DSAP and SSAP fields. If the field values are FF-FF (indicating a raw IPX frame), the switch generates a SNAP encapsulated frame on the FDDI with PT of 81-37 (Ethernet protocol type for IPX).

Conversely, when the switch receives a SNAP encapsulated IPX frame from the FDDI that contains a PT of 81-37, the switch converts the frame to a Raw 802.3 IPX frame for transmission to an Ethernet LAN.

IPX translation is disabled by default (Figure 3-26). It can be enabled via the VNswitch Command Line Interface (CLI) or by use of clearVISN management software. For information about the CLI and clearVISN management software, refer to the section Switch Configuration and Management.

Figure 3-26: Frame Formats When IPX Translation Is Disabled

6 octets	6 octets	2 octets	1 octet	1 octet	variable	4 octets
DA	SA	Length	FF	FF		FCS

Raw 802.3 Frame on Ethernet Side

IPX Data

FDDI Side

1 octet	6 octets	6 octets	1	1	variable	4 octets
FC	DA	SA	FF	FF		FCS
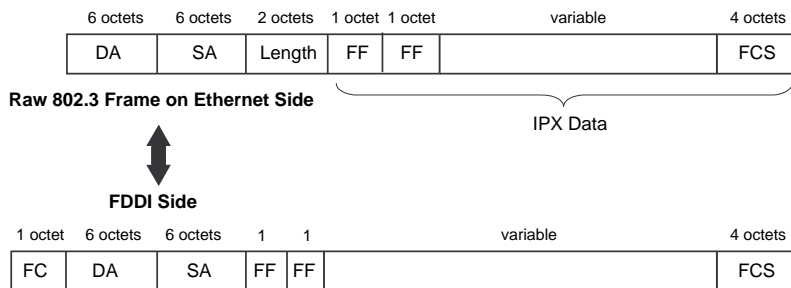
LKG-10135-96F

AppleTalk Translation

On Ethernet, AppleTalk V1 uses Ethernet format frames with a PT value of 80-F3. AppleTalk V2 uses 802 SNAP format frames with a PID value of 00-00-00-80-F3.

As shown in Figure 3-27, if a frame is translated to FDDI format without formal rules, both frames can translate to SNAP frames with a PID value of 00-00-00-80-F3. These frames are then translated back to Ethernet frames at the second switch (or bridge).

Most protocols that use both formats (for example, TCP/IP) treat Ethernet format frames with a PT of *xx-yy* and 802 SNAP format frames with a PID value of 00-00-00-*xx-yy* as equivalents, as specified by RFC 1122. AppleTalk treats them as two separate protocol versions that are not compatible. The solution is to use special rules when dealing with protocols such as AppleTalk.

Using these special rules, an AppleTalk V1 frame translates, on the FDDI side, to a frame with a PID value of 00-00-F8-80-F3. An AppleTalk V2 frame translates to a frame with a PID value of 00-00-00-80-F3.

At the next switch, each frame is translated back into its original form and the version information is preserved.

Figure 3-27: AppleTalk-to-FDDI Translation

AppleTalk V1 Frame on Ethernet Side

6 octets	6 octets	2 octets	variable	4 octets
DA	SA	80-F3	Data	FCS

PT

Standard Translation

Translated Frames on FDDI Side

1 octets	6 octets	6 octets	1	1	1	3 octets	2 octets	variable	4 octets
FC	DA	SA	AA	AA	03	00-00-00	81-37	Data	FCS

DSAP SSAP CTL OUI PT

Indistinguishable After Translation

Standard Translation

AppleTalk V2 Frame on Ethernet Side

6 octets	6 octets	2 octets	1	1	1	3 octets	2 octets	variable	4 octets
DA	SA	Length	AA	AA	03	00-00-00	81-37	Data	FCS

DSAP SSAP CTL OUI PT

LKG-10136-96F

Standard Ethernet-to-FDDI Translation

Table 3-5 describes the Standard Ethernet to FDDI translation rules.

Table 3-5: Standard Ethernet-to-FDDI Translation Rules

Rule	Description
1	The 802 format frames traverse the network with the LLC field unchanged.
2	Ethernet format frames are converted into 802 SNAP frames with a Protocol ID consisting of 3 octets of zeros, followed by the Ethernet Protocol Type.

AppleTalk ARP Ethernet-to-FDDI Translation

The VNswitch 900 series switches are provided with a non-translated protocols (NTP) table that lists protocols which require special handling. Normally, the NTP table contains a single entry: the protocol type for AppleTalk ARP (80-F3).

The VNswitch 900EF switch uses the rules in Tables 3-6 and 3-7 for translating between Ethernet and FDDI formats.

Table 3-6: AppleTalk ARP Ethernet-to-FDDI Translation Rules

Rule	Description
1	The 802.3 format frames traverse the network with the LLC field unchanged (see Figure 3-28).
2	For Ethernet frames, look up the Protocol Type (PT) in the NTP table: If the PT is found in the NTP table, translate to a SNAP frame with Protocol ID (PID) of 00-00-F8 followed by the PT (see Figure 3-29). Otherwise, translate to a SNAP frame with PID of 00-00-00 followed by the PT in the Ethernet frame.

Table 3-7: AppleTalk ARP FDDI-to-Ethernet Translation Rules

Rule	Description
1	If the frame is SNAP with Protocol ID (PID) beginning with 00-00-F8, translate to Ethernet format (see Figure 3-29).
2	If the frame is SNAP with PID beginning with 00-00-00, check to see if the last 2 octets of the PID are in the special handling table. If not, translate to Ethernet format. If it is in the table, leave the frame as an 802 format frame (see Figure 3-28).
3	All other frames traverse the network with the LLC field unchanged.

Figure 3-28: AppleTalk ARP V2 Translation
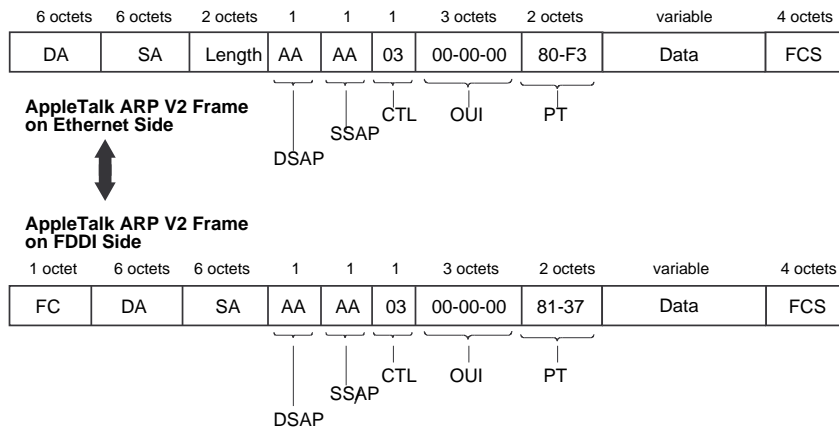
LKG-10137-96F

Figure 3-29: AppleTalk ARP V1 Translation

6 octets	6 octets	2 octets	variable	4 octets
DA		80-F3	Data	FCS

PT

AppleTalk ARP V1 Frame on Ethernet Side

Translated AppleTalk ARP V1 Frame on FDDI Side

1 octet	6 octets	6 octets	1	1	1	3 octets	2 octets	variable	4 octets
FC	DA	SA	AA	AA	03	00-00-00	80-F3	Data	FCS

DSAP SSAP CTL OUI PT
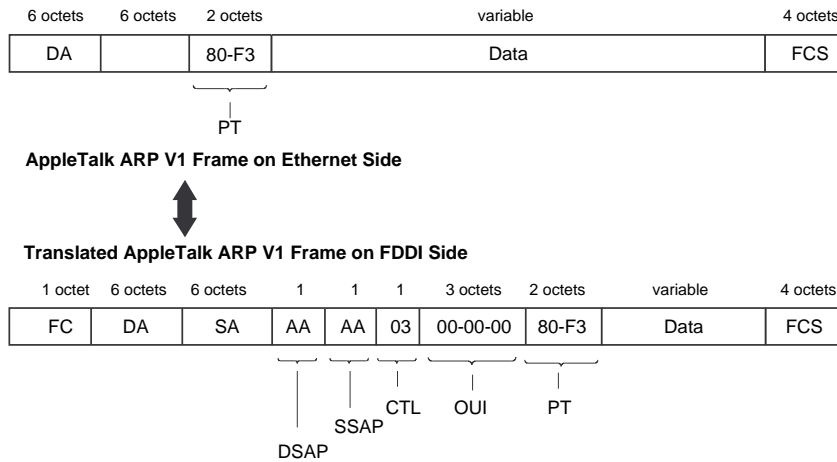
LKG-10138-96F

Ethernet-to-ATM Translation

The VNswitch 900EA translates Ethernet frames (both V2.0 and 802.3 frames) and FDDI frames into ATM fixed-length cells before passing the information onto the ATM network.

Translation is accomplished using the ATM Adaption Layer 5 (AAL5) protocol along with either bridge tunnels (RFC 1483) or LAN Emulation (LANE 1.0). Frames are first translated into the appropriate LANE or bridge tunnel formats and then segmented into ATM cells. For information about ATM, AAL5, LANE, and bridge tunnel formats, refer to the section VNswitch 900EA Frame Formats.

Rate Limiting

Introduction

Uncontrolled rates of broadcast and multicast traffic (broadcast storms) place severe strain on network connectivity. Broadcast storms can use up so much bandwidth that the network effectively becomes unusable—nodes, overcome with excessive traffic, are prevented from communicating information on the network.

While it is rare to see an indefinite broadcast storm, it is fairly common for lengthy peaks (several seconds to minutes) to occur in medium- to large-sized networks.

Curbing Broadcast Storms

The VNswitch 900 series switches include a feature called *rate limiting* for containing broadcast storms. Specified multicast addresses can be included in the rate limit category. Also, specified protocols (as encoded in the Ethernet PT, SAP, or SNAP fields) can be included in the rate limit if they occur in any multicast frame. Finally, the rate limit can be set in packets-per-second (pps) with a granularity of 100 pps.

The VNswitch 900 series switches do not forward frames in excess of the preset rate limit for specified frames. This feature regulates further propagation of this traffic into the extended LAN by eliminating the peaks of this traffic while forwarding, thus isolating it to the LAN on which the storm originates.

For detailed instructions on how to set rate limiting, refer to either the *VNswitch 900 Series Switch Management* guide or the *clearVISN Configuration and Use* guide.

Frame Filtering

Introduction

The VNswitch 900 series switches offer the following types of filtering:

- Destination address filtering

- Source address filtering

- Multicast/broadcast filtering

- Manual mode/secure mode

- Protocol filtering

- Unknown-protocol filtering

- VLAN filtering

Destination Address Filtering

Destination address filtering decreases network traffic by keeping local traffic local and only forwarding frames that are destined for the opposite side of the switch.

Each switch has an address database that contains addresses of stations connected to its LANs. When the switch detects a frame on any of its attached LANs, it checks the destination address of that frame against the addresses stored in its database.

If the frame is addressed to a node that has an address that was learned on the *same port* of the switch, the switch filters the frame (it does not forward it to the other switch ports).

Address information is entered in the database in either of two ways:

1   Listening/learning — The switch listens to network traffic and acquires a working knowledge of the nodes that are connected to its LAN ports.

It acquires this knowledge by reading the source address of each incoming frame and noting on which of the LAN ports the source address is located.

2   Switch management — The address database can also receive entries from switch management.

By specifying a station address and the port number of the switch where that station is located, you can lock that station to that port. This creates a permanent address for that node and causes the switch to ignore any learned information that differs from this permanent address.

Source Address Filtering

Using switch management, you can prevent all frames emanating from a specified node from being forwarded to the other ports of the switch (regardless of their destination address).

Source address filtering can also be used to lock down a specified source node to a single port or group of ports on the switch. This protects against masquerading nodes.

For example, consider a secure station (station A) that is known to be on switch port 2. To prevent stations on other LANs from masquerading as station A, switch management can be used to lock down station A to switch port 2.

Multicast/Broadcast Filtering

Switch management can also be used to channel all broadcast and multicast traffic to a backbone port.

For example, suppose your servers are on the FDDI backbone and you do not want the Ethernet LAN's multicast address to go to any of the other Ethernets. In this case, you could lock down a multicast or broadcast address to the FDDI backbone by defining an address filter that allows this address on the FDDI port only. This filter prevents multicast or broadcast frames from flooding any of the other Ethernet ports.

Manual Mode/Secure Mode

The switch can be set to manual mode (also known as secure mode) on an arbitrary set of ports. When in manual mode, the switch shuts off its learning capability on the specified set of ports and forwards frames based only on the address and protocol filters configured by management.

It is important to note that manual mode filtering can be specified on a per-port basis. If a port has been placed into manual mode and a packet is attempting to *enter* the switch on that port, the packet's source address (SA) is compared to the list of addresses that was manually loaded by the network manager. If the address is found in the table, the packet will be forwarded (if no protocol filter is in effect to block it); if the address is not found, the packet is blocked from entering the switch. On the other hand, if a packet is attempting to *exit* the switch at that port, its destination address (DA) is compared to the list of addresses that was manually loaded, and, if a match is found, the packet will be forwarded; otherwise, the packet is discarded.

Manually loaded address entries are never aged out of the learning tables. They are permanently stored in the table until the network manager explicitly removes them. This is true for all manually entered addresses—whether the port is in manual mode (address learning turned off) or auto-learning mode (standard bridge learning).

Protocol Filtering

Switch management can also be used to prevent specified protocols from being forwarded across the switch.

Protocol filtering can be used in an inclusive or an exclusive fashion. For example, if you specify the Local Area Transport (LAT) protocol for protocol filtering, you can have the switch filter only LAT protocol messages and forward all other protocols (if no other filters match), or you can have the switch forward only LAT protocol messages (if no other filters match) and filter all others.

Switches can be set to restrict a specific Ethernet Protocol type, an 802 DSAP, or an 802 SNAP PID to a set of allowed ports. Frames containing these protocols are only forwarded if both of the following situations are true:

- The source port is in the set of allowed ports for the protocol.

- The destination port is also in the set of allowed ports for the protocol.

Note that filtering always takes precedence over forwarding.

For example, if you specify that the switch always forward messages to a particular destination address (for example, node Z), and you also specify that the switch always filter LAT protocols, then any LAT protocol message destined for node Z will be filtered. The switch will forward a frame only if no filters are set to filter that frame.

Unknown-Protocol Filtering

Using switch management, you can set the switch to filter or forward on any port, any protocol that does not match a configured protocol filter. By default, a switch will forward an unknown protocol.

VLAN Filtering

Establishing VLAN filters is a function of creating your VLANs. For example, when you create a port-based VLAN (or VSD), you effectively create a port filter that restricts delivery of frames to only those ports in the VLAN.

Unlike the other switch filtering mechanisms described in this section (which provide counters), no VLAN counters are currently implemented.

ATM Features

Introduction

The VNswitch 900EA switch functions as an edge device in an ATM network, incorporating 16 virtual ports within a single, flexible, high-speed connector. The 16 virtual ATM ports support any combination of the following communication schemes:

- LAN Emulation (LANE)

- Bridge tunneling

LAN Emulation (LANE)

LAN Emulation (LANE) is a specification that defines how legacy LANs can communicate across an ATM network with other LANs, ATM hosts, or ATM edge-devices. LANE-compliant devices can be linked together in an emulated LAN (ELAN) to share information across the ATM network, regardless of each device's underlying technology (Ethernet, Fast Ethernet, FDDI, or ATM).

The VNswitch 900EA supports as many as 16 ELANs on its ATM port. The ATM port (which is configured into 16 virtual ports) supports up to 1022 virtual circuits (VCs). A separate VC is created for each client in the ELAN, creating a full mesh of virtual circuits, as shown in Figure 3-30.

Figure 3-30: Full Mesh of Virtual Circuits

LKG-10211-96F
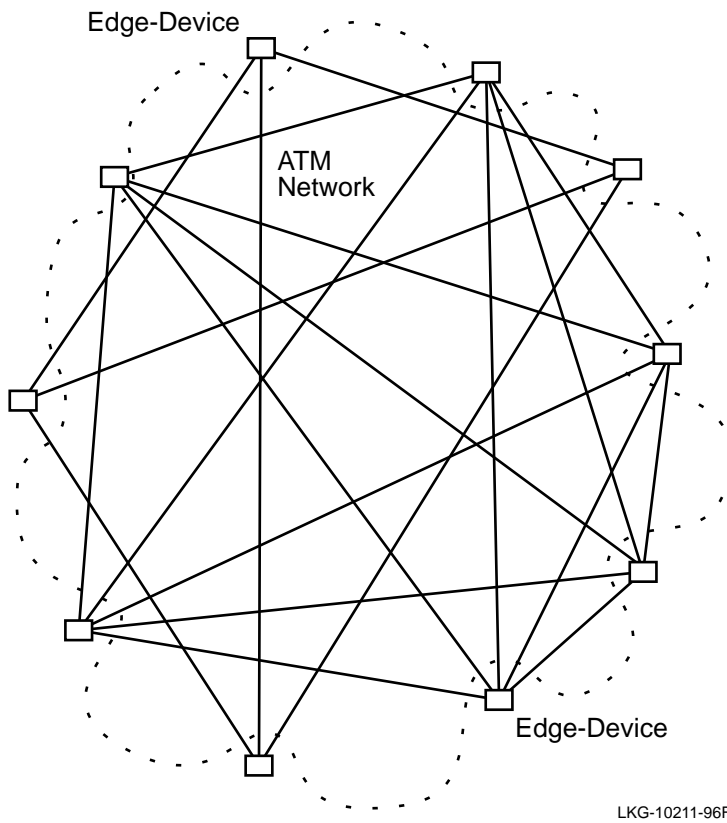
The ability of the VNswitch 900EA to support many virtual circuits provides flexibility in designing the size of your ELANs. For example, you could connect the VNswitch 900EA to 16 different ELANs, each containing as many as 55 edge-devices or ATM hosts. This configuration would require 960 virtual circuits (55 x 16 = 880, plus 5 server VCs for each ELAN for a total of 80 more; thus, 880 + 80 = 960). Several of the edge-devices could be other VNswitch 900EA modules, with each module supporting hundreds or thousands of attached stations.

Instead of using all 16 virtual ports on the VNswitch 900EA, you could configure the VNswitch in a single ELAN comprised of over 1000 members (ATM edge-devices and hosts). Figure 3-31 illustrates a single ELAN with only six members. In this configuration, only 20 VCs are used (note that the server VCs are not illustrated).

Figure 3-31: Example of a Small ELAN
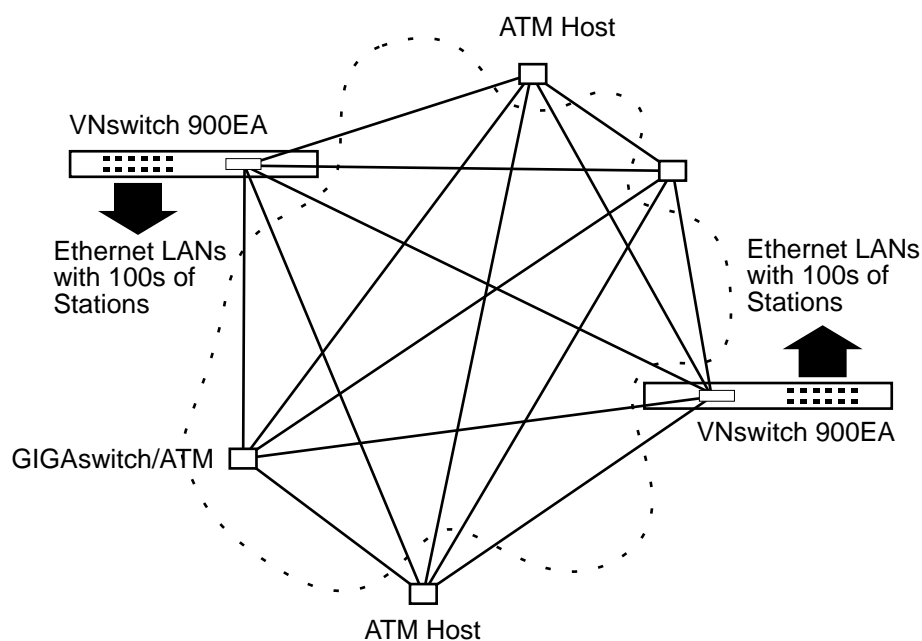
LKG-10212-96F

LAN Emulation involves the interaction of various LANE components to successfully transport data between members of the ELAN. The VNswitch 900EA module functions as a LANE client (supporting up to 16 LANE clients—one for each virtual ATM port) and interfaces with other LANE components including the LAN Emulation Server (LES), LAN Emulation Configuration Server (LECS), and Broadcast/ Unknown Server (BUS). These other LANE components reside in other devices (such as the GIGAswitch/ATM) attached to the ATM network.

For more information about the various components used to implement LAN Emulation, refer to Chapter 6.

ELANs in Relation to VLANs

An important feature of the VNswitch 900EA module is that it allows virtual LANs (VLANs) to be joined with other VLANs across an ATM network. This is accomplished by mapping VLANs from remote VNswitch 900EA modules to the same ELAN, effectively creating an extended VLAN. For more information, refer to the section Virtual LANs (VLANs).

Bridge Tunnels

The VNswitch 900EA module supports RFC 1483, allowing as many as 16 ATM *bridge tunnels* to be created on its ATM port. A bridge tunnel is a point-to-point connection between two legacy LAN devices across an ATM network. Either Ethernet or FDDI frames can be transported through a single bridge tunnel, and any mix of Ethernet and FDDI bridge tunnels is allowed. Figure 3-32 illustrates a bridge tunnel connecting a VNswitch 900EA module to an 802.3 multiport bridge.

Figure 3-32: Example of a Bridge Tunnel
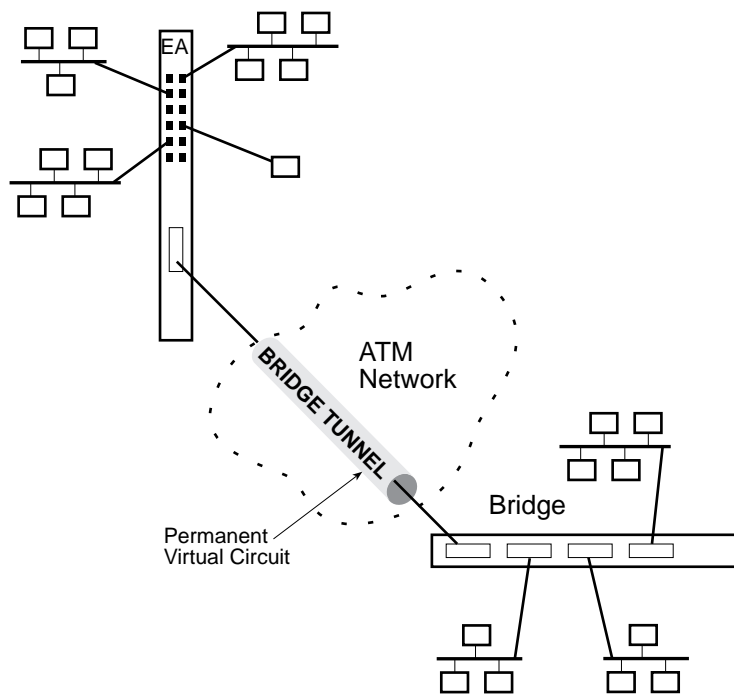
LKG-10217-96F

A bridge tunnel is established over a permanent virtual circuit (PVC). Each bridge tunnel requires a separate PVC. Generally, bridge tunnels are used only when communicating with non-LANE compliant devices, for example, between a VNswitch 900EA module and a standard 802.3 bridge. A bridge tunnel might also be used when only two VNswitch 900EA modules need to connect to each other but do not require the elaborate communication mechanisms provided by LAN Emulation (such as the ability to map VLANs across the ATM network or communicate with ATM hosts).

Using ELANs and Bridge Tunnels Together

The VNswitch 900EA allows LAN Emulation and bridge tunnels to be used simultaneously on its ATM port. Figure 3-33 depicts the 16 virtual ports implemented within the VNswitch 900EA module's ATM port. In this illustration, just a few of the many possible combinations of the virtual ports are shown.

Figure 3-33: Multiple Uses of ATM Virtual Ports
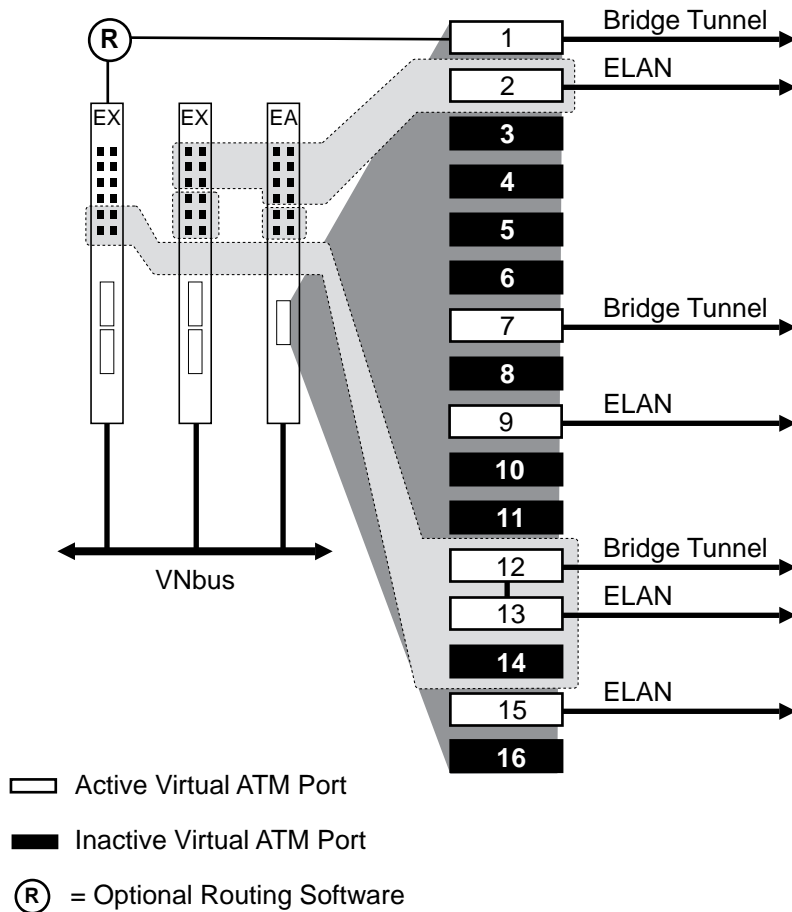
LKG-10223-96F

In Figure 3-33, only 7 of the possible 16 virtual ATM ports are used. The active virtual ATM ports (VAPs) in this illustration provide the following features and functionality:

- Virtual ATM port (VAP) 1 is connected to a VNswitch 900EX module equipped with the optional routing software. VAP 1 is configured as a bridge tunnel (RFC 1483) that connects to the VNswitch 900EX module to gain access to its routing facility and the LANs to which the router is logically attached. (Information about the optional routing software is provided later in this chapter.)

- VAP 2 is configured as a LANE client that connects LANs on two modules (the VNswitch 900EA module and the adjacent VNswitch 900EX module) to an ATM emulated LAN (ELAN).

- VAP 7 is configured as a bridge tunnel that provides a connection (across the ATM network) between the VNswitch 900EA module and another VNswitch 900EA module, or any other ATM edge-device that does not require (or is not capable of) LAN emulation services.

- VAPs 9 and 15 are configured as LANE clients that are not grouped with any other ports. These two VAPs connect the VNswitch 900EA module to two separate ELANs, each of which could be comprised of many ATMs hosts and edge-devices.

- VAP 12 is a bridge tunnel and VAP 13 is a LANE client, which are bridged together. This configuration allows some LANs on the VNswitch 900EX module to become members of an ELAN (through VAP 13), while at the same time providing (through VAP 12) a connection to an older device that does not yet support LANE.

As shown in Figure 3-33, the VNswitch 900EA module allows LANE and RFC 1483 bridge tunnels to coexist and interoperate on a single ATM edge-device. Another unique feature of the VNswitch 900EA module is its ability to interconnect ATM ELANs without having to dedicate any physical ports. For example, in Figure 3-33, only virtual ATM ports (VAPs 9 and 15) are used to connect the VNswitch 900EA module to ELANs. Any VNswitch module equipped with the optional routing software can be used to link the two remote ELANs together to provide connectivity between all members of the two ELANs.

NOTE

In the *VNswitch 900 Series Switch Management* guide, the virtual ATM ports are referred to as *ATM logical interfaces*.

Backplane LANs

Introduction

All VNswitch module ports, with the exception of the Fast Ethernet ports, can be redirected to the DEChub 900 backplane to create *backplane LANs*. A backplane LAN is an internal LAN segment created on one or more of the flexible channels located on the DEChub 900 backplane. Refer to Figure 3-34.

When VNswitch modules (and other network modules) are connected to these segments, they exchange data just as if they were on external LANs.

NOTES

DEChub 90 modules connect only to the ThinWire segment in the hub, and cannot connect to the flexible channels. The VNswitch modules support connection to the ThinWire segment, allowing VNswitch modules and DEChub 90 modules to be connected to the same backplane LAN.

This section does not discuss ATM on the backplane. ATM on the backplane is accomplished using dedicated channels that differ from the flexible channels (and ThinWire segment) used for backplane LANs.

Figure 3-34: Backplane LAN Segments or Backplane LANs

Slots 1   2   3   4   5   6   7   8

DEChub 900

ThinWire
Segment

Ethernet*
Backplane
LAN
Segments
1–12

0
1
2
3
4
5
6
7

Backplane
LAN
Segments
for FDDI,
and the
VNbus

8
9
10
1   13
12   14
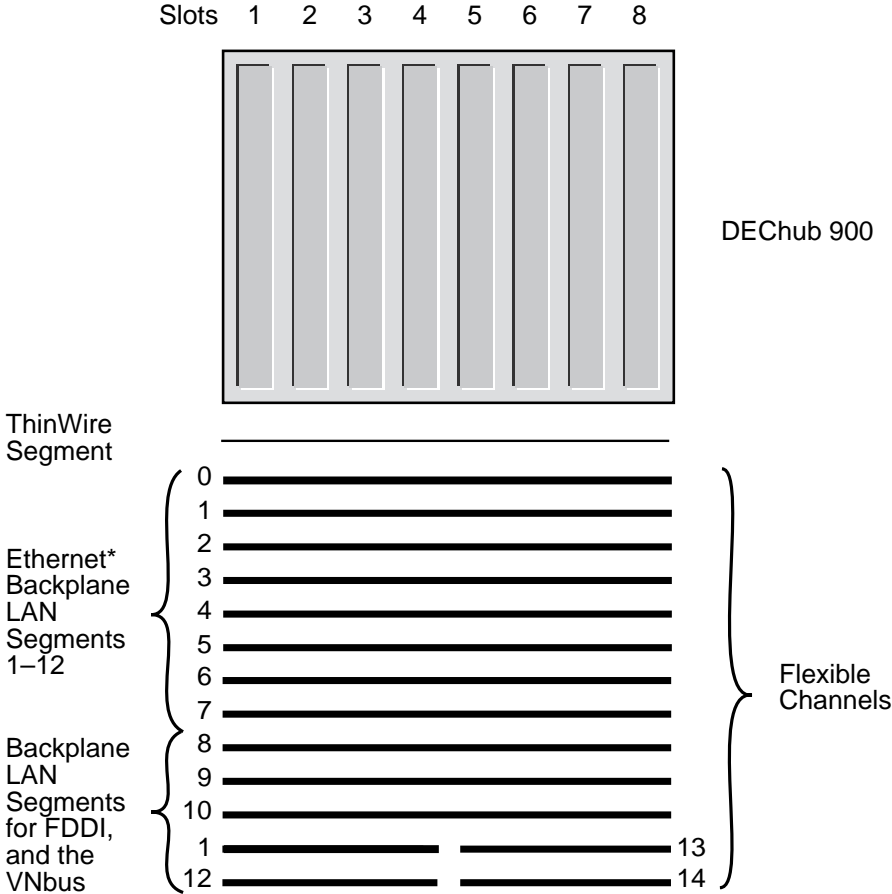
Flexible
Channels

*Channels 1-7 can also be used for FDDI LAN segments   LKG-10222-96F

As shown in Figure 3-34, flexible channels 11–14 span only half the width of the hub:

- Channels 11 and 12 span slots 1–4

- Channels 13 and 14 span slots 5–8

Only modules installed in slots 1–4 can access channels 11 and 12. Similarly, only modules installed in slots 5–8 can access channels 13 and 14.   Generally, these channels are reserved for VNbuses. However, if a VNbus is not being used, its associated flexible channels can be used for backplane LANs.

How Are Backplane LANs Managed?

Backplane LANs (or backplane LAN segments) are created, deleted, and managed using clearVISN software. These tasks are performed remotely without interacting directly with network hardware.

Using clearVISN, a network manager defines a backplane LAN by mapping selected ports from two or more modules to a specific flexible channel (LAN segment).

NOTE

When a port is directed to the backplane to create a backplane LAN, the associated front panel port becomes inactive.

Backplane LAN Examples

This section describes two uses of backplane LANs:

- Connecting a non-VNswitch module to a VNswitch module

- Connecting multiple non-VNswitch modules to a VNswitch module

Connecting a Non-VNswitch Module to a VNswitch Module

An example of using a backplane LAN to connect a non-VNswitch module to a VNswitch module is shown in Figure 3-35.

In this example, a backplane LAN segment (4) is created on flexible channel 3. A DECswitch 900EE is connected to this backplane LAN segment along with the VNswitch 900EX to create a backplane LAN. This configuration allows work groups connected to the DECswitch 900EE to access the two servers connected over a high-speed link attached to a VNswitch 900EX module.

Figure 3-35: Connecting a Non-VNswitch Module to a VNswitch

Printer

DECswitch 900EE

EX

To
10 Mb/s
LANs

Work Group

Work Group

Server

Server

100
Mb/s
Link

Server

Backplane
Ethernet LAN
Segments on
Flexible
Channels 0-5

1
2
3
4
5
6

100
Mb/s
Link

7
8
9
10
11
12

Backplane
Ethernet LAN
Segments on
Flexible
Channels
6 and 7
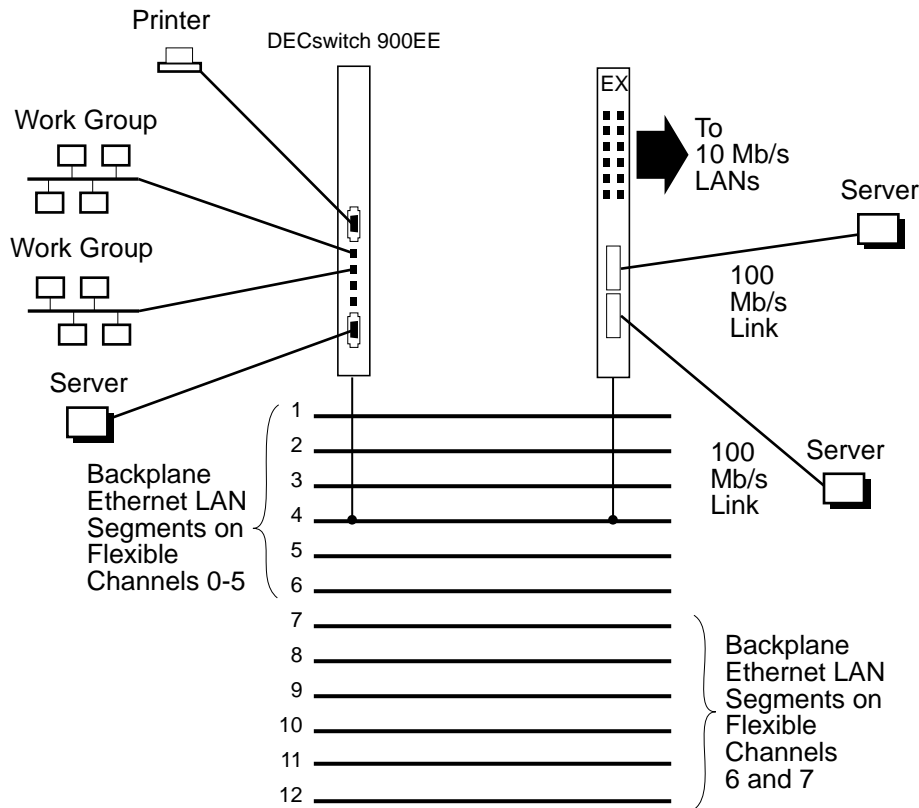
LKG-10205-96F

NOTE

The DECswitch 900EE and other earlier DEChub 900 modules support only the
first six flexible channels (0–5), corresponding to backplane LAN segments 1–6.
Therefore, these modules can only join backplane LANs on the first six backplane
LAN segments.

Connecting Multiple Non-VNswitch Modules to a VNswitch Module

An example of using multiple backplane LANs to connect non-VNswitch modules to a VNswitch module, and to each other, is shown in Figure 3-36.

In this example, a PORTswitch 900TP/12 module and a DECswitch 900EF module are connected to each other, and to a VNswitch 900EA module, across the backplane of the DEChub 900. The following backplane LANs are created:

- Three backplane LANs are created on segments 4, 8, and 11 to connect 3 ports on the VNswitch 900EA with three ports on the PORTswitch 900TP/12.

- Two backplane LANs are created on segments 2 and 3. Both of these backplane LANs include a single port from each module.

- A single backplane LAN is created on segment 6 to connect ports on the two non-VNswitch modules only.

This configuration allows LANs connected to the PORTswitch 900TP/12 module to be linked with the ports, LANs, and VLANs on the VNswitch 900EA module. These LANs can then access the high-speed ATM port on the VNswitch 900EA.

Connecting VNswitch Modules Using Backplane LANs

VNswitch modules are designed with the flexibility to allow the modules to be connected by means of either backplane LAN segments or the VNbus. However, the preferred method is for VNswitch modules to connect to each other by means of a VNbus. By using the VNbus to connect VNswitch modules, you do not need to redirect the front panel ports to the backplane, hence, the front panel ports remain active.

**Figure 3-36: Example of Using Multiple Backplane LANs to Connect
Non-VNswitch Modules to a VNswitch Module**
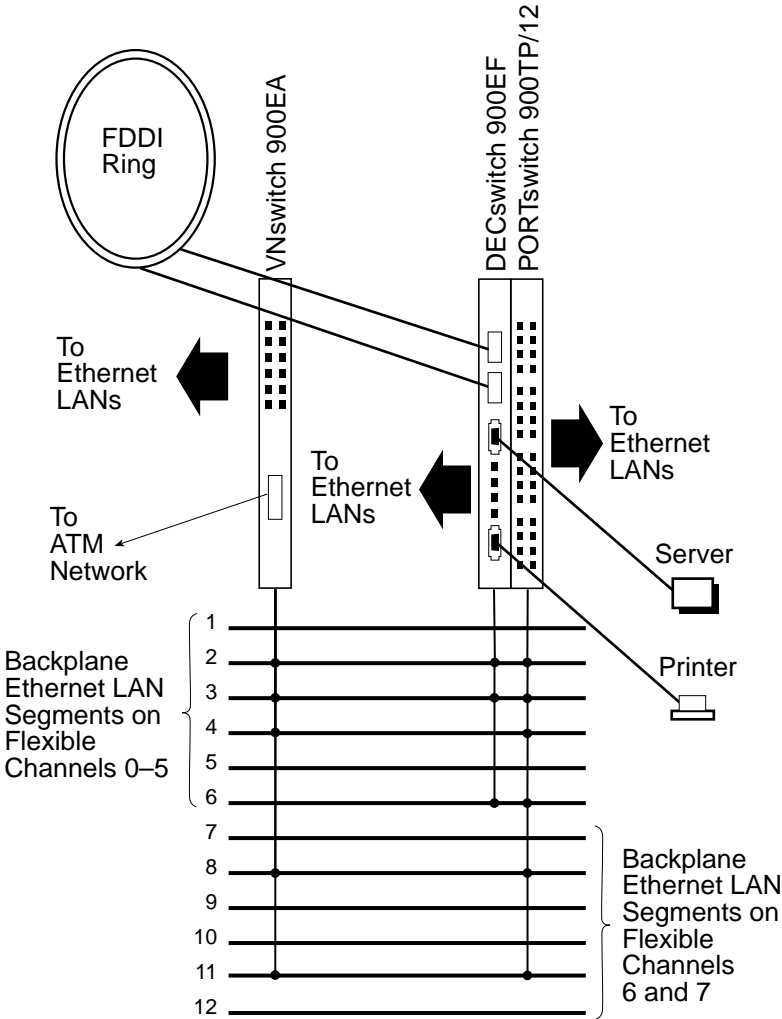
LKG-10204-96F

Virtual LANs (VLANs)

Introduction

A virtual LAN (VLAN) is a separate broadcast domain whose membership is based on logical groupings instead of physical media connections. This means that LAN membership can be defined by users' communication needs rather than by the physical connections of nodes and media. Members of a VLAN can communicate as though they were connected to the same physical LAN.

VLANs establish boundaries between LANs and isolate traffic to one or more LAN segments. A powerful feature of VLANs is the ability to group ports from different switches into the same VLAN. Each VNswitch can support as many as 32 VLANs.

Class 1 VLANs

Version 1.5 of the VNswitch 900 series switches supports Class 1 VLANs. A Class 1 VLAN is a group of ports on one or more VNswitch modules within a single hub (either a DEChub ONE docking station or a DEChub 900 MultiSwitch). The exception to this rule is the case of *standards-based VLANs over ATM*, described later in this section.

Figure 3-37 illustrates examples of several Class 1 (port-based) VLANs.

Figure 3-37: Class 1 VLANs
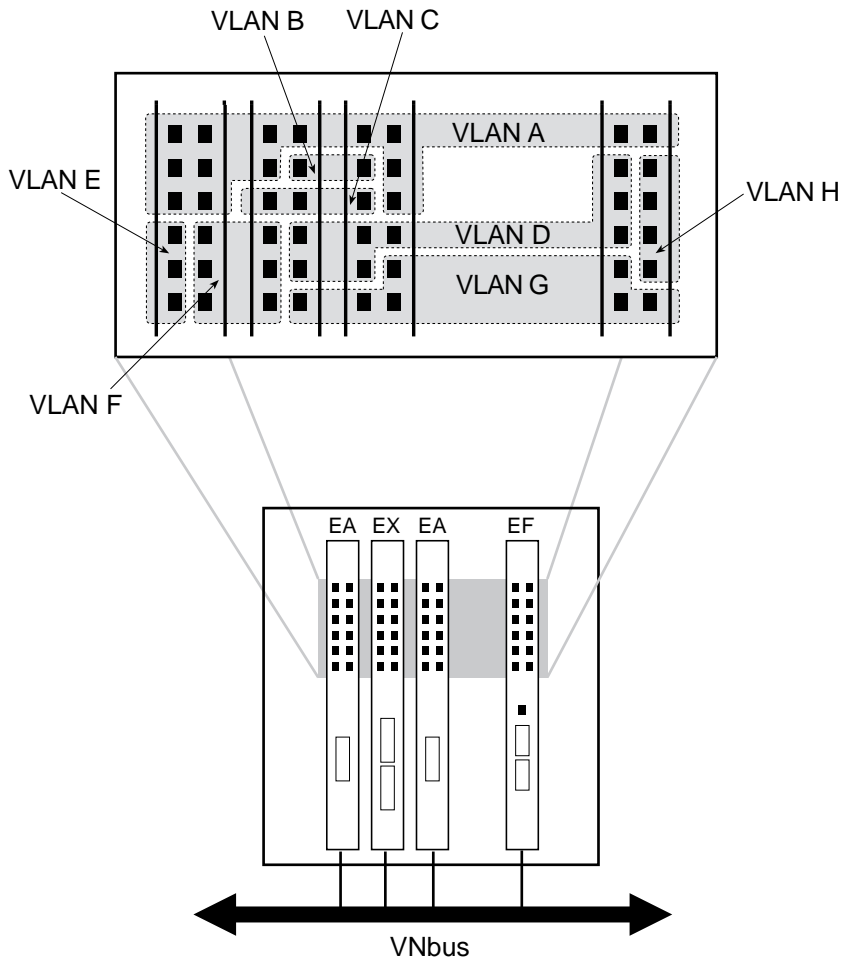
LKG-10237-96F

By default, all ports (including the high-speed ports) on the VNswitch modules are grouped into a single VLAN, called the *default VLAN*. Through software commands, the ports can be arbitrarily divided into port groups (as shown in Figure 3-37). Any ports not included in a port group (or VLAN) remain in the default VLAN. Thus, the high-speed ports in this illustration are all members of the default VLAN.

VLAN Secure Domain (VSD)

A VLAN Secure Domain (VSD) is a set of ports (on one or more switches) operating its own separate 802.1d spanning tree process. Version 1.5 of the VNswitch 900 series modules allows a single Class 1 (port-based) VLAN for each VSD. This arrangement of allowing a single VLAN per VSD ensures the highest degree of stability and security by ensuring that all VLANs are isolated from one another. Only physically separate LANs unconnected in any way would be more secure and isolated from each other. Refer to Figure 3-38.

Figure 3-38: VLAN Secure Domain

= VLAN

= VSD
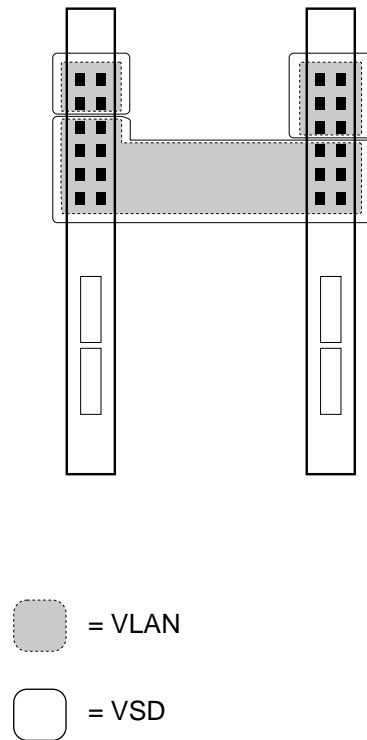
LKG-10213-96F

Since only port-based (Class 1) VLANs are supported in Version 1.5, and a VSD consists of a single set of ports (just like a port-based VLAN), the terms VSD and VLAN are sometimes used interchangeably. However, it is important to distinguish between a VLAN and VSD because future versions of the VNswitch firmware will support multiple VLANs per VSD (including VLANs based on ports and protocol type), with features such as overlapping VLANs.

Each VNbus can support as many as 63 VSDs. Together, the three VNbuses implemented on the DEChub 900 backplane allow as many as 189 VSDs. Thus, Version 1.5 of the VNswitch 900 series switches supports up to 189 port-based VLANs (one per VSD).

Each VLAN must reside wholly within a single VSD (it may not cross VSD boundaries). A VSD must be composed of a distinct set of ports (non-overlapping), and no traffic can flow between VSDs unless it is routed. Refer to Figure 3-39.

Virtual LANs (VLANs)

Figure 3-39: Traffic Flow Between VSDs
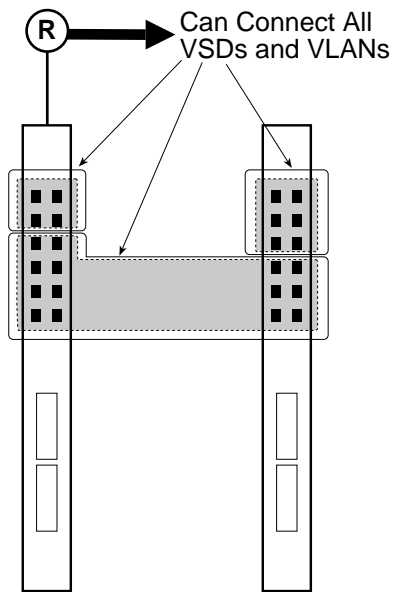
R = Optional Routing Software

= VLAN

= VSD

LKG-10214-96F

As shown in Figure 3-39, a module equipped with the optional routing software enables traffic to be routed between VSDs and VLANs. Without the routing software, no traffic can flow between VSDs or VLANs.

Standards-Based VLANs over ATM

ATM emulated LANs (LAN Emulation) provide a means for spanning VLANs across an ATM network.

The VNswitch 900EA module supports 16 emulated LANs (ELANs) on its ATM port. Using the Command Line Interface (CLI), you can include a VLAN in any of the 16 ELANs. An ELAN may consist of VLANs, LANs, ATM hosts, and other vendors' ATM edge-devices. By including VLANs from different switches or hubs into a single ELAN, you can create a single VLAN (or *extended VLAN*) that spans across the ATM network as illustrated in Figure 3-40.

Figure 3-40: Standards-Based VLANs over ATM

LKG-10315-96F
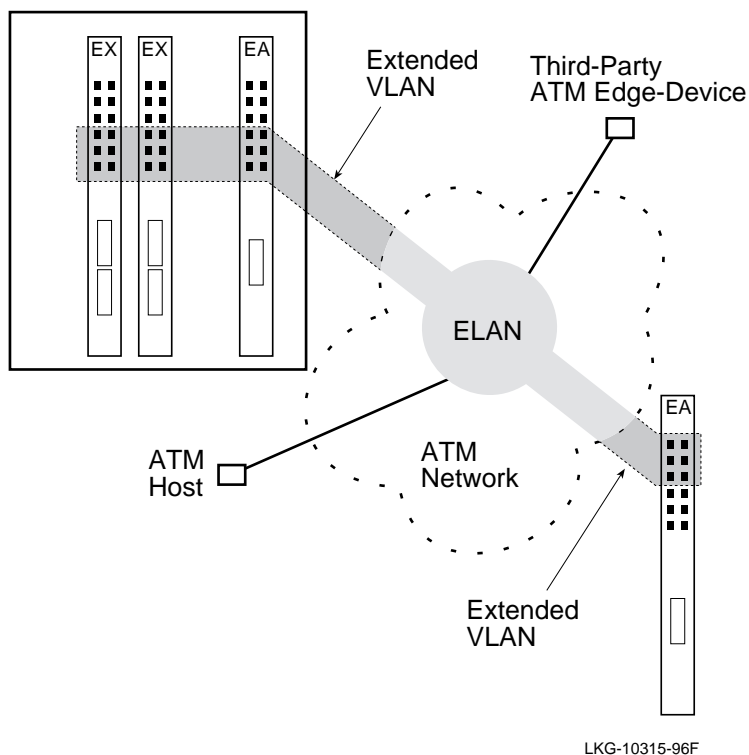
As illustrated in Figure 3-40, a VLAN comprised of ports from three different VNswitch modules in a DEChub 900 is grouped with a VLAN on a standalone VNswitch 900EA switch attached across an ATM network. This is accomplished by

mapping both VLANs to the same emulated LAN (ELAN). In this example, an ATM host and a third-party ATM edge-device are also included in the ELAN, and thus are also members of the extended VLAN.

As many as 16 VLANs can be mapped into the ATM network through a single VNswitch 900EA module. Since a single hub configuration of VNswitch modules can consist of almost 200 port-based VLANs, it may be desirable to map more than 16 VLANs into the ATM network. In this case, multiple VNswitch 900EA modules can be used, each providing support for 16 emulated LANs (ELANs). Eight VNswitch 900EA modules configured in a DEChub 900 can allow as many as 128 port-based VLANs to be joined to 128 ELANs connected to an ATM backbone, as illustrated in Figure 3-41.

Figure 3-41: High-Density VLANs over ATM

Up to 8
VNswitch 900EA
Modules Supporting
128 ELANs

16 VLANs
Over ATM

ATM
Network

ELANs

16 VLANs
Over ATM

16 VLANs
Over ATM

ATM
Network

ELANs

To Other
VNswitch 900EA
Modules, ATM Hosts,
or Edge-Devices

ATM
Network

ELANs

To Other
VNswitch 900EA
Modules, ATM Hosts,
or Edge-Devices
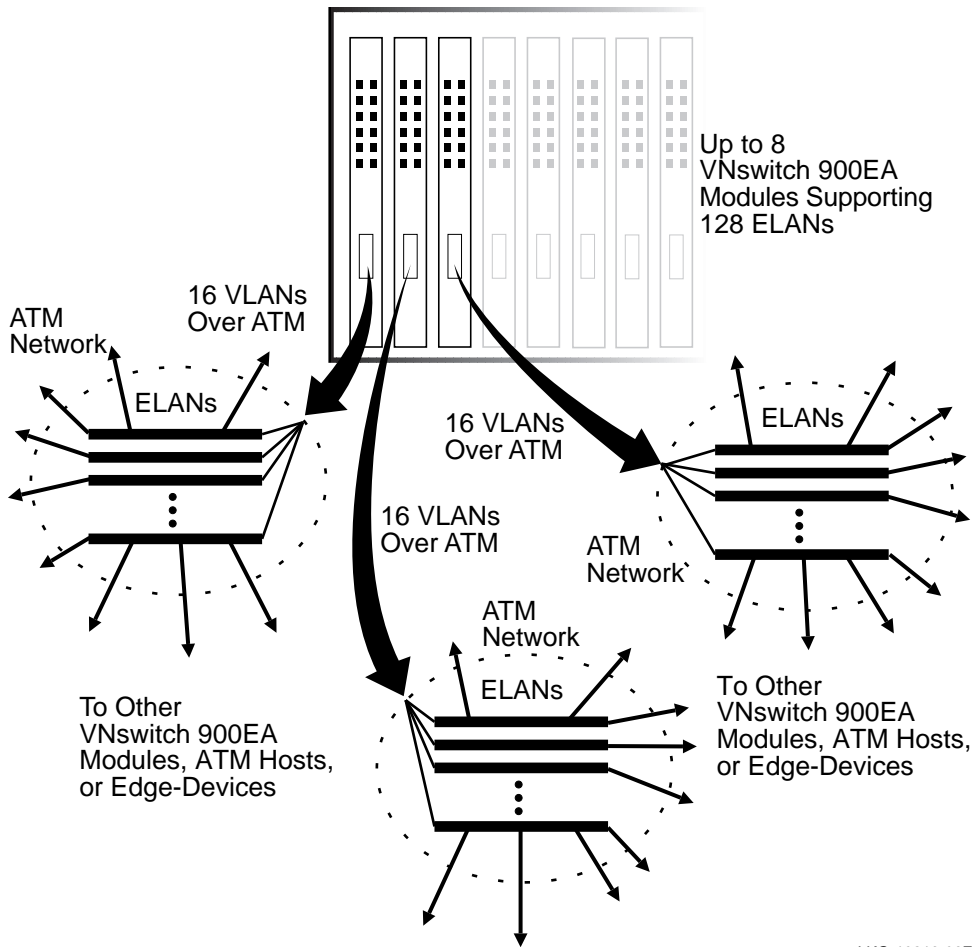
LKG-10216-96F

Optional Routing Software

Introduction

Version 2.0 of the VNswitch consists of an optional routing software package that adds IP routing support to the VNswitch 900 family of switches. The routing software provides routing between multiple VLANs.

The optional routing software can be installed on a single VNswitch to handle the routing needs of all the VNswitch modules in a DEChub 900 MultiSwitch, or multiple modules can be upgraded with the routing software if more capacity or redundancy is needed. Each module equipped with the routing software supports up to 32 virtual interfaces providing routing services for as many as 32 different VLANs on a single VNbus. No physical ports on the VNswitch need to be dedicated to the routing software.

In a typical installation, a single VNswitch module upgraded with the routing software can handle all the IP routing requirements of an entire local area network supporting thousands of users.

Multilayer Switching

By incorporating routing functionality into the VNswitch, the module is able to provide switching at both Layer 2 and Layer 3. This feature is referred to as *multilayer switching.* Equipped with the routing software, the VNswitch supports unicast LAN frame forwarding (Layer 2 bridging) for all protocols, while ensuring that only the inter-VLAN traffic that requires IP packet header handling (Layer 3 routing) is handled by the routing process. Layer 2 bridging is performed at line rate at up to 750,000 pps, while Layer 3 routing is performed at up to 100,000 pps.

Summary of Features and Functions

The following summarizes the features and functions of the optional VNswitch routing software:

- Performs Layer 3 routing without compromising unicast bridging performance.

- Eliminates the need for expensive external routers.

- Supports as many as 32 virtual interfaces.

- Incorporates routing functionality into the VNswitch without requiring the use of any of the switch's external ports.

- Supports both RIP and OSPF routing protocols.

- Supports multiple filtering options, including filtering based on network, host address, or application (for example, FTP).

- Includes BootP Relay for DHCP support.

- Provides quick configuration management tools.

- Relies on the base configuration of DRAM to provide support for more than 2000 routing table entries. DRAM can be upgraded to provide capacity for almost 10,000 networks.

- Provides both LAN and WAN routing capabilities over ATM through support for:

  — Bridge tunnels between two modules equipped with the routing software

  — Construction of an Emulated LAN (ELAN) whose only direct participants are router virtual interfaces on various VNswitch 900 modules

For more information about the VNswitch Routing Software, refer to the *VNswitch 900 Series Router Management* guide.

Switch Configuration and Management

Introduction

The VNswitch 900 series modules provide two available management interfaces:

- Command Line Interface (CLI) — Accessed through the setup port on either the DEChub 900 MultiSwitch or the DEChub ONE docking station. It is necessary to initially access the CLI to set up the module's IP address. Once the IP address is configured, the switch can be managed remotely using the CLI (over a Telnet connection) or the SNMP interface.

- Simple Network Management Protocol (SNMP) management interface — Accessed using Digital's clearVISN Element Manager. The ClearVISN Element Manager is a flexible, network configuration and management application that you access by means of a graphical user interface (GUI). To communicate configuration and management commands to the VNswitch, clearVISN sends SNMP commands to the SNMP management interface in the VNswitch 900 series module.

Both management interfaces include support for in-band operation as well as for out-of-band operation (using a serial PPP/SLIP connection) for high-security environments.

NOTE

Although it is possible to use SNMP commands with any vendor's management station, *clear*VISN makes it easy to manage the switches by use of customized window displays.

Plug-and-Play Functionality

All VNswitch 900 series modules are auto-configuring on power-up, using the factory set default settings described in the *VNswitch 900 Series Switch Management* guide. When initially powered on, all VNswitch modules are bridged into a single virtual LAN, as shown in Figure 3-42.

Figure 3-42: Default Configuration of VNswitch Modules
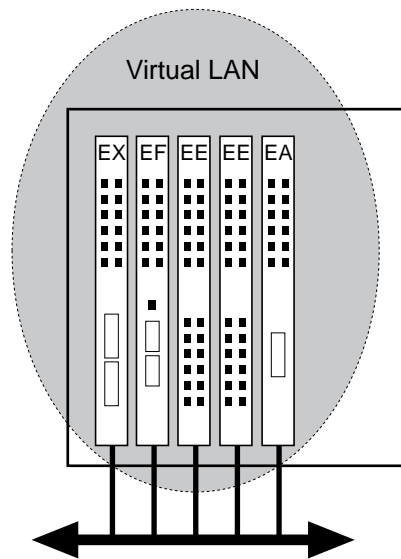
LKG-10193-96F

For many sites, the default configuration may be suitable. To modify the default configuration, you need to use either the CLI or SNMP management interface.

Using the Command Line Interface (CLI)

The CLI is a command line driven management tool for managing a VNswitch module. It includes a Telnet interface for remote command line control of all routing, bridging, and VLAN functions, plus access to a wide range of performance statistics and status information.

The CLI is comprised of a set of configuration programs that are referred to as *processes*. This section describes the following major CLI processes:

- Main process — Functions as the main control process and provides console services for one local or two remote console users.

- Configuration process — Provides online control of various configuration parameters, such as network addresses and event logging.

- Monitor process — Provides online monitoring of the various software components, such as protocols, network interfaces, and event logging.

- Events process — Receives system event messages and displays them on the console monitor.

Main Process

The Main process is the root level process of the CLI. Its primary function is to control the processes that are connected to the switch console. The Main CLI commands allow you to do the following:

- Configure a process.

- Change the intercept character.

- Reload the switch software.

- Monitor a process.

- Communicate with other processes.

- Manipulate the output from a process.

- Restart the switch software.

- Telnet to other switches or hosts.

- Display process statistics.

- Display information about switch memory usage.

- Dump the contents of the switch memory to a remote host.

The Main process allows you to log in to a VNswitch module from either a console directly connected to the docking station or hub, or through a Telnet session.

Configuration Process

You use the Configuration process (CONFIG) to modify configuration information for the switch. Using CONFIG commands, a console operator can perform the following tasks:

- Set, modify, and display various configuration parameters, including information stored in the switch's nonvolatile RAM (NVRAM).

- Configure a minimal set of parameters to boot the switch and run it with a minimal configuration or to download a complete configuration by means of TFTP.

- Clear, list, or update configuration information.

- Enable or disable console login and modem control.

- Communicate with protocol environments.

- Add, modify, or remove entries from the boot and dump configuration database.

- Disable or enable network memory dumping and assign a unique name to the dump files.

- Use the TFTP protocol to transfer configuration information between the switch's nonvolatile memory and remote hosts.

- View the current boot and dump configuration database.

NOTE

Some CONFIG commands are reserved for use by Digital Customer Support.

Monitor Process

The Monitor process is a set of programs that allow you to display switch statistics and modify the switch's operating environment. Using the Monitor process, you can execute commands to do the following:

- List the protocols and interfaces currently configured in the VNswitch.

- Display memory and network statistics.

- Set current Event Logging System (ELS) parameters.

- Test a specified network interface.

- Configure volatile configuration parameters.

- Enable and disable interfaces.

Events Process

The Events process receives messages from the Event Logging System (ELS) and displays them on the console monitor. The Event Logging System, which is accessed from the Main process, is a monitoring system that is an integral part of the VNswitch operating system. ELS manages event messages that are logged as a result of VNswitch activity. Using ELS commands, you can sort out only those messages that are important to you, and then display them on the console terminal screen through the Events process.

Events occur continuously while the VNswitch is operating. They can be caused by any of the following reasons:

- System activity

- Status changes

- Service requests

- Data transmission and reception

- Data and internal errors

When an event occurs, the ELS receives information that identifies the source and nature of the event. It then generates an event message that can be displayed on the console screen. The various ELS messages are described in the *Event Logging System Messages* guide.

For More Information

Topic	Where to Get More Information

Using the CLI to manage a VNswitch module	The *VNswitch 900 Series Switch Management* guide provides instructions for using the CLI to perform all switch management tasks, including:
	• Accessing CLI processes and prompts
	• Configuring transparent bridging parameters
	• Configuring VLANs
	• Configuring the VNbus
	• Upgrading and reinstalling switch software
	• Performing routine maintenance procedures such as:
	  — Configuring automatic image recovery
	  — Backing up configuration settings
	  — Checking available memory
	  — Capturing diagnostic data for problem analysis

Obtaining definitions of system event messages	The *Event Logging System Messages* guide provides a listing and description of all possible event messages.

Using *clear*VISN

The clearVISN SNMP-based management tool provides access to various statistics associated with the VNswitch 900 series modules, and provides graphical user interfaces for configuring backplane LANs on the DEChub 900. ClearVISN manages network modules installed in the hub as well as the hub itself.

The clearVISN management tool provides control over almost all VNswitch functions with the exception of advanced routing configuration options and user access lists.

NOTE

The optional routing software module includes its own GUI-based Router Configurator tool for conducting and maintaining simple routing configurations off line and then downloading as needed.

Using the clearVISN MultiChassis Manager, ATM, FDDI, and Ethernet backplane connections in the DEChub 900 to other switches, routers, or repeaters are easily established and visualized. VNbus connections between VNswitch modules are also managed in this way. The MultiChassis Manager provides front bezel views of the all modules in the DEChub 900 (including LEDs) and allows "double-click" access to module-specific configuration, status, and performance details.

With the clearVISN VLAN Manager, advanced VLAN configurations (including those involving ATM Emulated LANs) are easily created, visualized, and reconfigured.

Other clearVISN applications supporting the VNswitch 900 family include: Recovery Manager, Alarms Manager, Flash Loader, Net Designer and Router Manager. All clearVISN applications are integrated under a common database, and can easily share information with each other.

What You Can Do with clearVISN

You can use clearVISN software to:

- Configure VNswitch 900 series modules along with other DEChub 900-based modules

- View the status and activity of network modules and ports

- Configure LANs on the hub backplane

- Configure VNbus connections

- Configure port-based VLANs

- Monitor the flow and accuracy of network data

- Perform firmware updates

VNswitch 900 Series Windows

The clearVISN windows include the following information for managing the VNswitch 900 series modules:

clearVISN Window	Description
Hub Front Panel	Front panel representation of the hub and installed network modules, including status and port LEDs. Double click on a switch module in the front panel view to open the Switch Summary window.
Switch Summary	Provides key overall, as well as per-port, information for the switch. Allows you to set the root priority and port cost for the spanning tree.
Switch Port Information	Provides information related to counters (datalink-independent and datalink-dependent), Ethernet port, Fast Ethernet port, FDDI port, and ATM port information.
FDDI Summary	Provides detailed information specific to the FDDI interface.
ATM Summary	Provides detailed information specific to the ATM interface.
Forwarding Database	Used for locating where a specific MAC address, IP address, or DECnet address resides.
Address Filter	Allows configuration of filtering rules for destination address (DA) and source address (SA), manual or secure mode for each port (for disabling learning), and rate limiting for multicast addresses.
Protocol Filter	Allows configuration of filtering rules for Ethernet protocol types: 802.DSAP, 802 SNAP, or PID; rate limiting for any of these protocols when used in multicast frames; default protocol filtering action for each port.
Switch Configuration	Provides configuration settings of various switch parameters, such as: rate limit, spanning tree parameters, raw 802.3 IPX switch.

clearVISN Window	Description
Station Configuration	Used to configure the switch's ports to point to the hub backplane rather than the switch's front panel. This configuration must be done before connecting the port to the specific backplane LANs.
VLAN Manager	Used to configure and manage Class 1 port-based virtual LANs (VLANs).

For More Information

Topic	Where to Get More Information
Overview of clearVISN	The *clearVISN Product Overview* provides an overview of clearVISN and a description of each clearVISN application.
Using clearVISN windows	The *clearVISN Configuration and Use* guide provides instructions for configuring VLANs, VNbuses, and backplane LANs. It also provides instructions for setting up filters and enabling rate limiting.

Chapter 4

Ethernet to Fast Ethernet Switching

Overview

Introduction

Fast Ethernet is a LAN technology that allows standard 10 Mb/s Ethernet frames to be transmitted at 100 Mb/s.   This chapter provides a brief overview of Fast Ethernet technology and illustrates various configurations of the VNswitch 900EX module in an Ethernet to Fast Ethernet switching environment.

In This Chapter

The following topics are covered in this chapter:

Topic	Page
Fast Ethernet Overview	4-2
Ethernet to Fast Ethernet Switching Examples	4-6

Fast Ethernet Overview

Introduction

Fast Ethernet technology is defined in a set of ISO/IEC standards:

- 100BaseT — Designed as a 100 Mb/s extension to 10BaseT Ethernet, this specification maintains the original CSMA/CD (carrier sense multiple access with collision detection) medium access control mechanism of a Standard 10BaseT Ethernet system.

- 100VG-AnyLAN — Originally designed with the intention of supporting token ring frames, this specification defines a new medium access control layer that implements a "demand priority" mechanism similar to the MAC layer in a hub.

The 100BaseT standard has been formally designated as IEEE 802.3u. The 100VG-AnyLAN standard has been designated as IEEE 802.12.

Because 100BaseT uses the CSMA/CD access scheme, it preserves the media access control (MAC) layer functions of 10 Mb/s Standard Ethernet. For this reason, 100BaseT has become the most widely implemented Fast Ethernet standard. The VNswitch 900EX implements the 100BaseT standard and its subsets, including 100BaseTX and 100BaseFX.

100BaseTX

The 100BaseTX standard is a Fast Ethernet media standard that specifies baseband operation over two pairs of Category 5 unshielded twisted-pair (UTP) cable, or two pairs of Type-1 shielded twisted-pair (STP) cable. One pair of wires is used for receiving data signals and the other pair for transmitting data signals.

100BaseFX

100BaseFX standard is a Fast Ethernet media standard that specifies baseband operation over two strands of fiber optic cable.

NOTE

The TX and FX medium standards are collectively known as 100BaseX. These standards were adopted from the ANSI physical media standards originally developed for the Fiber Distributed Data Interface (FDDI) LAN standard.

100BaseT4

100BaseT4 is a Fast Ethernet media standard that specifies baseband operation over four pairs of telephone-grade twisted-pair wire (Category 3, 4, or 5 UTP).

Comparison to 10 Mb/s Ethernet

Compared to a standard 10 Mb/s Ethernet LAN, the amount of time it takes to transmit a bit on the Ethernet channel (called the bit-time) is decreased by ten times with Fast Ethernet. This produces a tenfold increase in the speed of the packets over the media system.

Fast Ethernet Repeaters

Two kinds of repeaters are specified in the 100BaseT system: Class I and Class II.

- Class I repeaters receive line signals on an incoming port and translate the signals to digital form. Before transmitting the signals, the signals are retranslated into line signals. This allows a Class I repeater to repeat signals between media segments that use different signaling techniques, such as 100BaseTX/FX segments and 100BaseT4 segments.

- Class II repeaters connect only to segment types that use the same signaling technique, such as 100BaseTX and 100BaseFX segments. When maximum cable lengths are used, no more than two Class II repeaters can be used within a single collision domain.

According to the Fast Ethernet standard, all repeaters in a Fast Ethernet system must be labeled with the Roman numeral "I" or "II" enclosed in a circle.

Full-Duplex Ethernet Links

Both Fast Ethernet and 10 Mb/s Ethernet support full-duplex operation on all individual links (LAN segments). Because there is no contention for the communication channel, devices can send and receive data simultaneously. Theoretically, this can provide twice the bandwidth of normal (half-duplex) Ethernet. Because they have transmit and receive signal paths that can be simultaneously active, 10BaseT, 100BaseTX, and 100BaseFX all support full-duplex operation.

Full-duplex fiber-optic links (100BaseFX) have the advantage of longer allowable segment lengths than normal half-duplex 100BaseFX. A full-duplex 100 Mb/s fiber link segment can be as long as two kilometers.

Link Integrity Test

Fast Ethernet transceiver circuits constantly monitor the link media to ensure that the link is functioning properly. The signaling mechanism used for Fast Ethernet segments is based on the ANSI FDDI signaling mechanism, which sends signals continually, even during idle periods when no network traffic is present. This ensures that the communications link between DTEs is always operational.

What Is Auto-Negotiation?

Auto-negotiation is a feature of Fast Ethernet that allows a device to detect the mode of operation (10 Mb/s or 100 Mb/s, and half- or full-duplex) supported by another device at the other end of a LAN segment. The VNswitch 900EX supports auto-negotiation on its Fast Ethernet ports that are configured with 100BaseTX.

Auto-negotiation enables a device to automatically configure itself to achieve the best possible mode of operation over a link. For example, auto-negotiation can detect whether a device is operating in half- or full-duplex mode, and can provide automatic speed matching for multispeed devices at each end of a link.

Using the auto-negotiation protocol, the VNswitch 900EX can automatically configure itself to use full-duplex mode when connected to a similarly equipped device.

Auto-Negotiation Priority

The following table list the priority assigned to operating modes:

Priority (Highest to Lowest)	Operating Mode
1	Full-duplex
2	Half-duplex
3	100 Mb/s
4	10 Mb/s

Fast Link Pulse (FLP) Signals

Auto-negotiation takes place using Fast Link Pulse (FLP) signals, which are also used for link integrity testing. The FLP signals are generated automatically at power-up, or may be selected manually using clearVISN management software. FLP signals are a modification of the Normal Link Pulse (NLP) signals used by 10 Mb/s Ethernet LANs.

The FLP signals can exist alongside the NLP signals, enabling an Ethernet device that uses NLP signals to continue to detect the proper link integrity even when attached to an auto-negotiation hub that sends FLP signals. Like the NLP signals, the FLP signals are generated during idle times on the communications link, and do not interfere with normal traffic.

FLP signals are specified for 100BaseTX only. Devices or ports connected to fiber-optic segments do not support FLP signals or auto-negotiation.

Ethernet to Fast Ethernet Switching Examples

Introduction

The VNswitch 900EX is a powerful, high-density switch that can satisfy the backbone switching requirements of various network configurations. Used as either a standalone module installed in a DEChub ONE docking station, or configured in a DEChub 900 MultiSwitch, the module provides a flexible solution to the bandwidth needs of various Ethernet LANs.

Like all VNswitch 900 series modules, the VNswitch 900EX switch is also easily incorporated into a multi-technology environment. This section is focused on Ethernet to Fast Ethernet switching examples. For examples of using the VNswitch 900EX in multi-technology environments, refer to Chapter 7.

NOTE

This section provides configuration examples for Version 1.5 of the VNswitch 900 family of switches. For examples of configurations using the optional VNswitch routing software, refer to Chapter 7.

Multiple Standalone VNswitch 900EX Modules

Typically, Fast Ethernet technology is used to provide a high-speed backbone over which 10 Mb/s Ethernet traffic is switched. Figure 4-1 illustrates a configuration in which both Standard and Fast Ethernet segments are attached to multiple standalone VNswitch 900EX switches.

Figure 4-1: Multiple Standalone VNswitch 900EX Modules
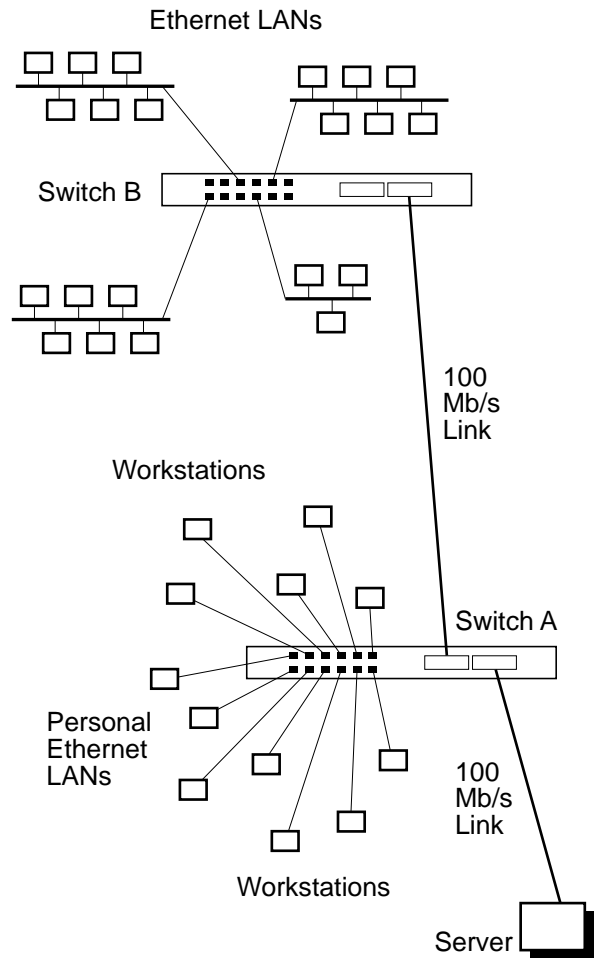
LKG-10220-96F

In this configuration, two VNswitch 900EX modules are connected by a Fast Ethernet link segment. Switch A is used as a hub for 12 high-performance workstations that are connected to private (or personal) Ethernet LANs. Each workstation has a dedicated 10 Mb/s connection to the switch and connects to the server over a 100 Mb/s high-speed link attached to one of the VNswitch 900EX's Fast Ethernet ports.

Switch B in this example connects four separate Ethernet LANs (or segments a single Ethernet LAN into four LANs) and enables all attached workstations to access the server connected to switch A. All traffic between switch A and switch B is transferred at 100 Mb/s.   Switch B provides room for network growth, allowing for the addition of more LAN segments and another high-speed 100 Mb/s link (possibly to another server, or to a Fast Ethernet repeater connecting multiple Fast Ethernet workstations).

Multiple Shared Ethernet

The following sections provide examples of using the VNswitch 900EX installed in a DEChub 900 along with non-VNswitch modules. Two multiple shared Ethernet environments are illustrated:

- Low-throughput configuration

- High-throughput configuration

Low-Throughput Configuration

Figure 4-2 illustrates a low-throughput configuration in which a single VNswitch 900EX module shares the hub's resources with multiple port switching repeaters. In this configuration, all front panel ports on the VNswitch 900EX module are directed to the backplane LAN segments, allowing the port switching repeaters to connect to the VNswitch 900EX and access the servers connected to its Fast Ethernet ports.

Figure 4-2: Low-Throughput Fast Ethernet Switching Example
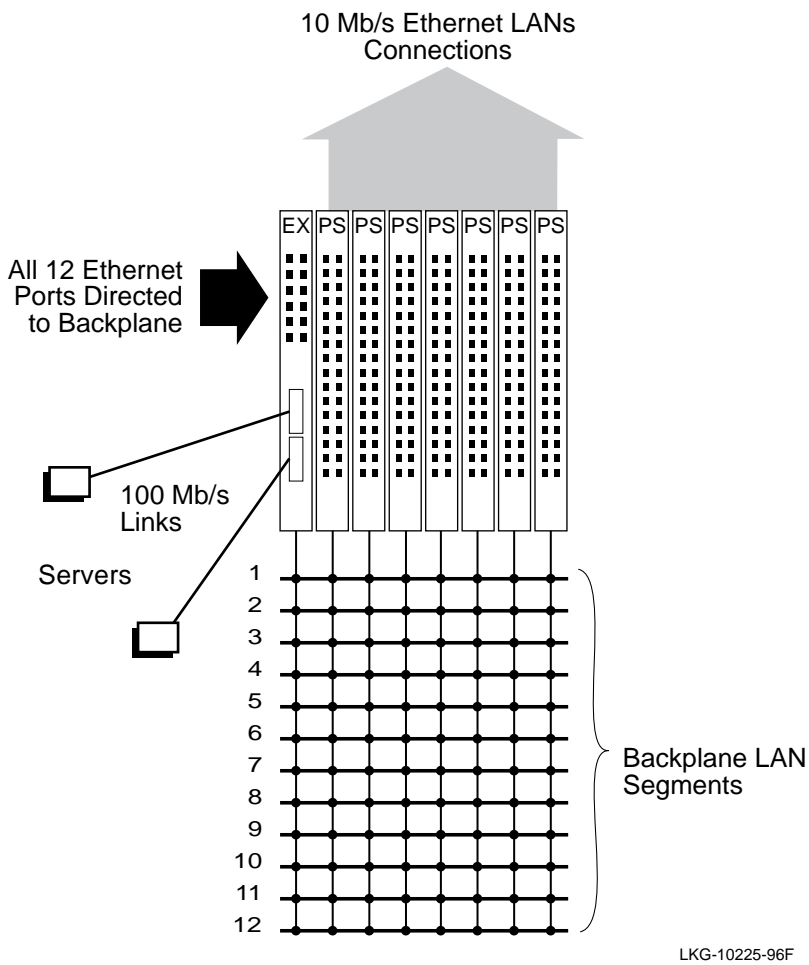
LKG-10225-96F

The PORTswitch 900TP/12 modules used in this example allow connections to all 12 backplane Ethernet LAN segments. Seven PORTswitch 900TP/12 modules configured in a single hub provide 224 (7 x 32) 10BaseT ports that can connect to the backplane LAN segments (18 per segment) to use the resources of the servers attached to the VNswitch 900EX module.

High-Throughput Configuration

Figure 4-3 illustrates a high-throughput configuration in which two VNswitch 900EX modules share the hub's resources with port switching repeaters. In this example, the VNswitch 900EX modules share the traffic load from the three PORTswitch 900TP/12 port switching repeaters.

Figure 4-3: High-Throughput Fast Ethernet Switching Example
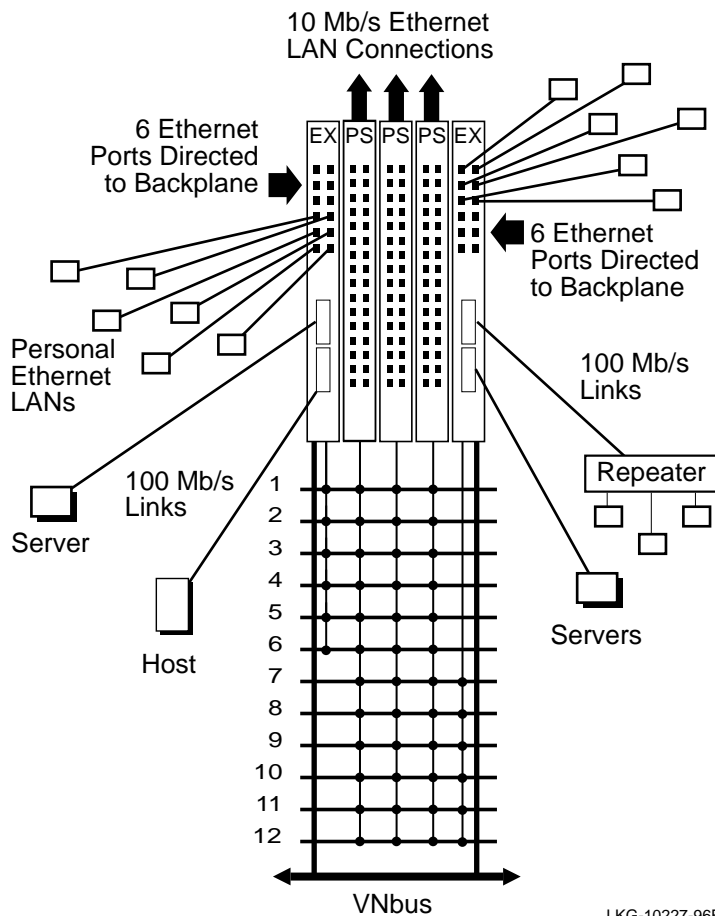
Dual High-Speed Ethernet Backbone

Figure 4-4 illustrates a network configuration in which VNswitch modules are used to create two Fast Ethernet collapsed backbones, each providing as many as 84 personal Ethernet connections to the network servers attached to the 100 Mb/s Fast Ethernet ports. In this example, the VNswitch modules also provide high-speed connections to a Fast Ethernet workgroup and an individual workstation.

Figure 4-4: Dual Collapsed Backbones
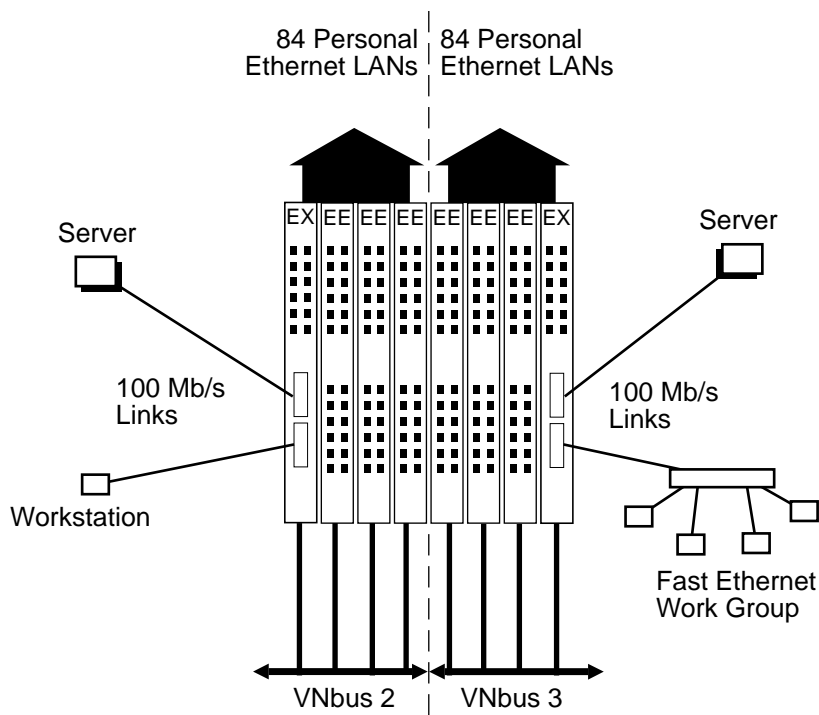
LKG-10229-96F

Shared and Switched Mixed Environment

Figure 4-5 further illustrates the versatility of the VNswitch 900EX module in both shared and switched Ethernet environments. In this example, multiple VNswitch 900EX modules and EE modules are configured with port switching repeaters in a DEChub 900 MultiSwitch. Servers are attached to the high-speed ports on all three VNswitch 900EX modules. All of the LANs connected to the VNswitch modules and port switching repeaters can access any of the five servers. Some of the ports on the VNswitch 900EE modules are used for personal Ethernet LANs, while others can be configured for multiple connections.

Figure 4-5: VNswitch 900EX in a Shared and Switched Environment

96 Ports of
Shared 10BaseT

24 Personal
Ethernet LANs

Mix of 24 10 Mb/s LANs
and Personal Ethernet LANs

EX EE PS PS PS EE EX EX

6 Ports on Each
EX Module Directed
to Backplane

All Ethernet
Ports Active
on Front Panel

6 Ports on Each
EX Module Active
on Front Panel

100 Mb/s
Links

100 Mb/s
Links

Servers

Servers

VNbus

LKG-10228-96F
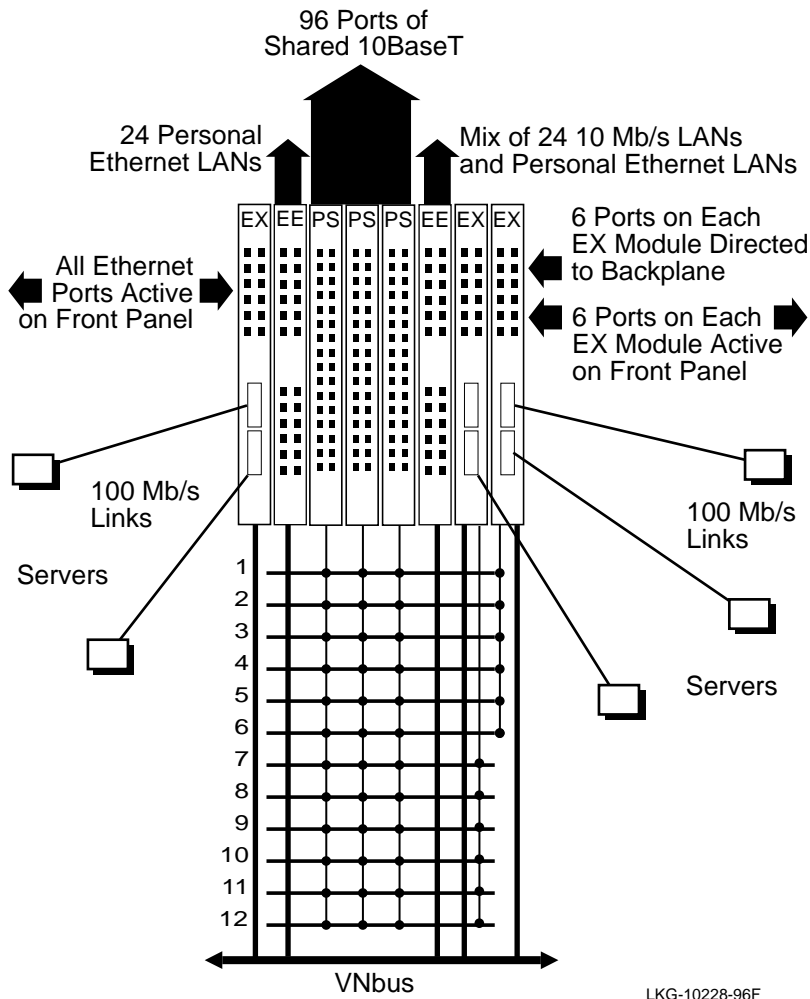
The VNswitch modules in this example are connected to a single VNbus, allowing traffic between the VNswitch modules to flow at 400 Mb/s, while also allowing VLANs to span across the modules.

All 96 ports on the 32-port PORTswitch 900TP/12 repeaters are connected to the backplane Ethernet LAN segments. Several ports on the VNswitch 900EX modules are redirected to the backplane to join backplane LANs with the port switching repeaters, while others are used for external LAN connections.

When a port that is a member of a VLAN is directed to the backplane, any other LANs connected to the same backplane LAN segment automatically become members of the VLAN. For example, any stations connected to repeater ports that are linked to a LAN segment occupied by a VLAN become members of the same VLAN.

NOTE

This differs from the VLANs implemented on the VNbus because VLANs that span across modules connected to a VNbus do not require any redirection of front panel ports. Also note that the VNbus is reserved for use by VNswitch modules.

Summary

In this chapter, we illustrated just a few of the many possible uses of the VNswitch 900EX module in an Ethernet to Fast Ethernet switching environment. Though only the PORTswitch 900TP/12 repeater was illustrated in the examples of shared Ethernet environments, the VNswitch 900EX switch is fully interoperable with all DEChub 900 modules. Additionally, the VNswitch 900EX switch can interoperate with other VNswitch modules (and other DEChub 900 modules) in multi-technology environments. For examples of using the VNswitch 900EX module in multi-technology environments, refer to Chapter 7.

For More Information

Topic	Where to Get More Information
Installing the VNswitch 900EX module	Refer to the *VNswitch 900EX Installation and Configuration* guide.
Configuring and using the VNswitch 900EX module	Refer to the *VNswitch 900 Series Switch Management* guide.

Chapter 5

Ethernet-to-FDDI Switching

Overview

Introduction

Fiber Distributed Data Interface (FDDI) is a set of ANSI/ISO standards for a 100 Mb/s token-passing ring that uses multimode fiber, single-mode fiber, unshielded twisted-pair, screened/shielded twisted-pair, or any combination of the four as the transmission medium. This chapter provides a brief overview of FDDI, discusses the FDDI configuration capabilities of the VNswitch 900EF, and illustrates various configurations of the VNswitch 900EF module in an Ethernet-to-FDDI switching environment.

In This Chapter

The following topics are covered in this chapter:

| Topic	Page
FDDI Overview	5-2
FDDI Configuration Capabilities	5-8
Ethernet-to-FDDI Configuration Examples	5-19

FDDI Overview

Introduction

All FDDI networks operate as logical token rings, where the right to transmit is granted by the possession of a token. There is one token per ring, and it is passed from station to station, according to a set of rules known as the timed token protocol. A station wishing to transmit on the ring first captures the token. It then transmits frames for a period of time determined by the timed token rules, and then releases the token immediately after completing its transmission.

A transmitting station is also responsible for removing the frames it transmitted from the ring once they have circled the ring and returned to the station. This process is called *frame stripping*.

Dual Ring

The most basic FDDI network is constructed of two independent rings and is commonly referred to as a dual ring (Figure 5-1). A dual ring connects to each FDDI station in the network. Each ring can span up to 100 km, which allows for a distance of 200 km (2 x 100 km/ring) for the FDDI dual ring. As many as 500 stations can be attached to the dual ring: typical configurations usually have no more than 200 stations.

**Figure 5-1: FDDI Dual Ring**
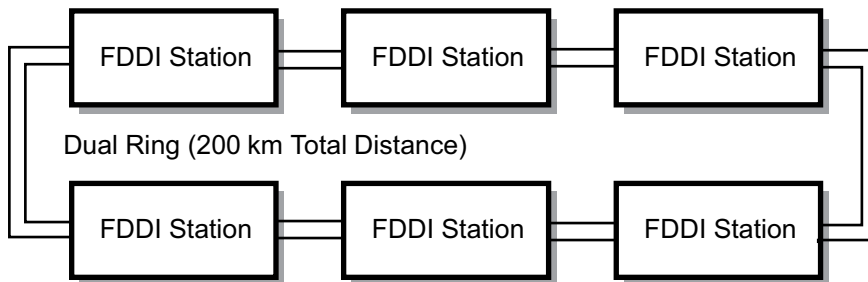
LKG-10095-96

FDDI Devices

FDDI networks are constructed using the following types of devices:

- Stations, of which there are two types:

  — Dual Attachment Station (DAS), which attaches directly to the FDDI ring

  — Single Attachment Station (SAS), which attaches to the ring through a concentrator

- Concentrators, of which there are two types:

  — Dual Attachment Concentrator (DAC), which attaches directly to the FDDI ring

  — Single Attachment Concentrator (SAC), which attaches to the ring through another concentrator

Figure 5-2 illustrates these devices in an FDDI ring. The VNswitch 900EF functions as either a DAS or SAS in an FDDI network.

Figure 5-2: FDDI Device Types
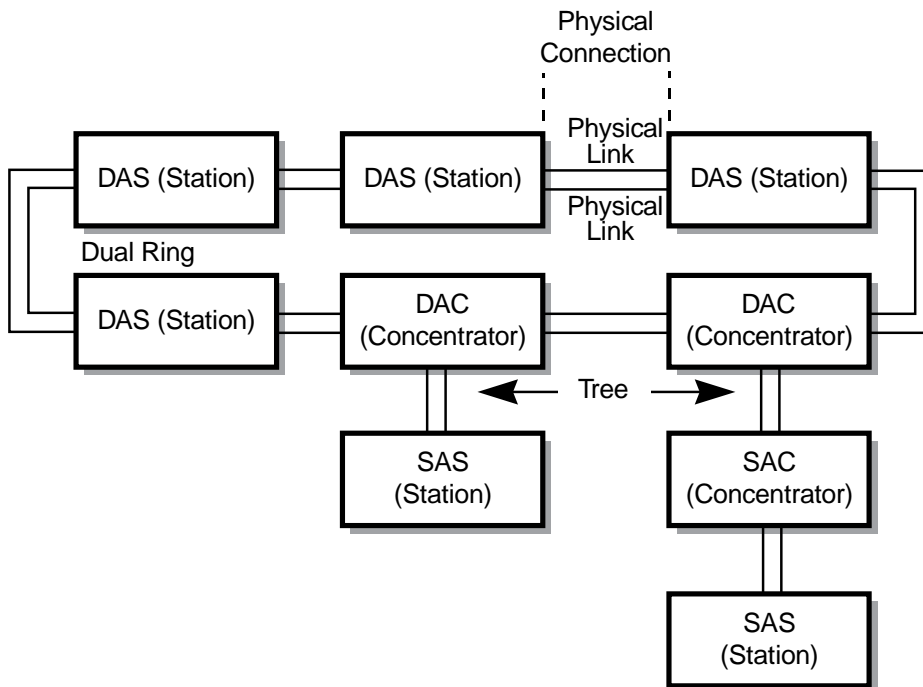
As shown in Figure 5-2, all devices connect to each other by means of a full-duplex connection, called a physical connection.

Each physical connection comprises two physical links (a single fiber-optic cable for each physical link):

- Single Attachment Stations and Single Attachment Concentrators connect to a concentrator, or to another station, via one physical connection.

- Dual Attachment Stations and Dual Attachment Concentrators connect to each other via two physical connections.

Media Types and Maximum Distances

The following tables describes the different media types that can be used to construct FDDI physical links:

Fiber Media	Fiber Type (Microns)	Power Budget (Decibels)	Maximum Link Distance (Kilometers)
Multimode Fiber (MMF)	62.5/125	11	2.0
Single-Mode Fiber (SMF)	8-to-10/125	22	60

| Copper Media	Copper Type	Maximum Link Distance (Meters)
Unshielded Twisted-Pair (UTP)	100-ohm Category 5 Twisted-Pair	100
Screened Twisted-Pair	100-ohm Category 5 Twisted-Pair	100
Shielded Twisted-Pair	150-ohm Category 5 Twisted-Pair (equivalent to IBM Type 1 cable)	100

NOTE

The VNswitch 900EF module supports MMF, SMF, and UTP media types.

FDDI Port Types

Stations and concentrators (Figure 5-3) can be identified by the types of ports that they use to attach to other stations. Following are the different types of FDDI ports:

- A and B ports — Dual Attachment Stations have A and B ports for attachment to other stations in the dual ring, or to concentrator M ports.

- M port — All concentrators (SAC or DAC) are identifiable by the presence of M ports. The M port allows the attachment of other stations or concentrators.

- S port — If the concentrator also has an S port, it is a SAC; if it has A and B ports it is a DAC. A Single Attachment Station has an S port for connection to a concentrator M port.

Figure 5-3: FDDI Port Types
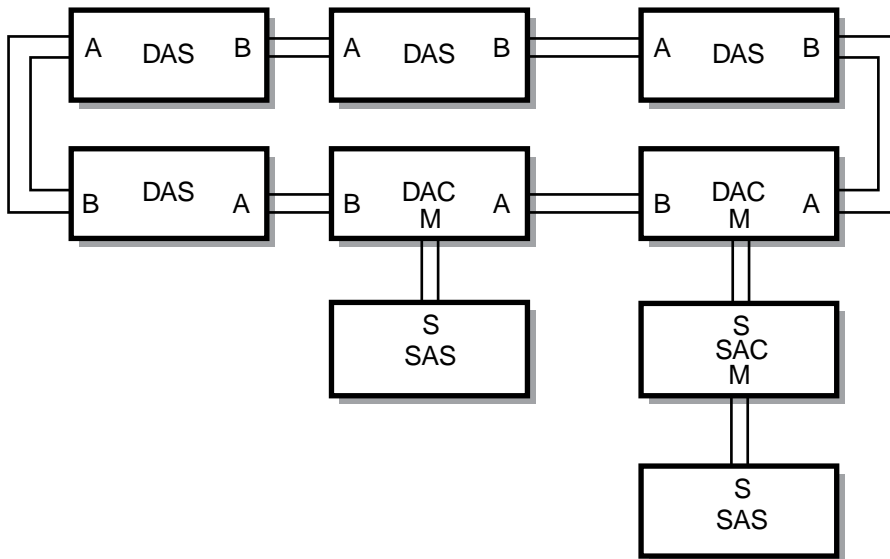
LKG-10097-96

Physical Topologies

Although all FDDI networks are logical rings, the following physical topology types can be created:

- Dual rings

- Trees

- Dual ring of trees

Dual Ring Topologies

Dual ring topologies (as shown earlier in Figure 5-1) are created by connecting only Dual Attachment Stations. The A port of one Dual Attachment Station (DAS) station connects to the B port of the next DAS, forming a physical connection between the two stations. The connections can continue until up to 500 stations are included in the dual ring.

Independent Rings

Because there are two physical links per physical connection (see Figure 5-2), two operating links are created between each DAS when the ring is formed. This design allows for two independent rings, the primary ring and the secondary ring, to be formed.

The primary ring enters the DAS at the A port, and exits the DAS at the B port. The secondary ring enters the DAS at the B port and exits the DAS at the A port.

Backup Benefit

In general, all data traffic between stations is carried on the primary ring, and the secondary ring is idle during this time. The FDDI standards allow both rings to be used for carrying data traffic, but in practice this is rarely done. The common practice is to use the secondary ring as a backup for the primary ring.

A benefit of having dual rings is that, in case of a failure, the secondary ring can be used to patch the primary ring. If a station detects that one of its A or B ports has failed, it wraps the primary and secondary rings together, restoring operation to the other stations in the ring. Wrapping occurs quickly, usually in well under a half a second.

Tree Topologies

FDDI also allows for the creation of tree topologies. Tree topologies start with a standalone concentrator at the top of the tree. Stations (DAS or SAS) or other concentrators (DACs) connect to this concentrator, and branch out from the top of the tree.

How to Identify Trees

The distinguishing feature of a tree topology is the connection to an M port. Single Attachment Stations (S ports) typically connect to a concentrator, but Dual Attachment Stations (A and B ports) can also connect to the M ports.

Benefit

A major advantage of the use of trees built with concentrators is the ability of the concentrator to electronically disconnect stations from the ring in case of failure, or by management control.

Dual Ring of Trees Topology

The third, and most common, FDDI configuration (Figure 5-4) is referred to as the *dual ring of trees*. This configuration type is a robust topology and is a hybrid of the dual ring and tree topologies.

How to Identify Dual Ring of Trees

Dual Attachment Concentrators (DACs) and Dual Attachment Stations (DASs) attach in a dual ring, with treed stations connected to the concentrators' M ports. The concentrator connects stations attached to these M ports into the flow of the primary ring.

Figure 5-4: Dual Ring of Trees Topology
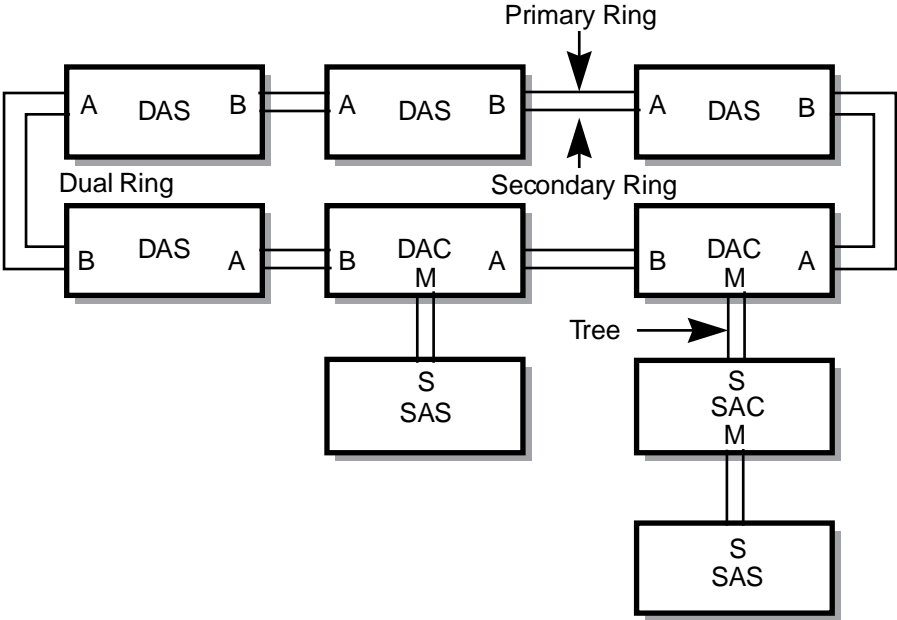
LKG-10098-96

FDDI Configuration Capabilities

Introduction

This section describes FDDI configuration capabilities that are supported by the VNswitch 900EF.

The VNswitch 900EF can be configured within individual or multiple FDDI networks:

- An individual FDDI network can be configured as either a dual ring or a tree. This allows compliancy with FDDI configuration rules that operate at the module level and are independent of DEChub 900 management considerations.

- Multiple FDDI networks can be constructed by creating FDDI networks on the backplane (either dual rings *or* trees), and then interconnecting the various networks across the front panel to create a dual ring of trees.

NOTE

The terms *dual ring* and *tree* are physical descriptions of the network topology. However, in all cases, the FDDI is a logical ring that enables a token to be passed from station to station in the ring. The physical implementation can be configured as a tree or a dual ring, but in all cases a logical ring exists.
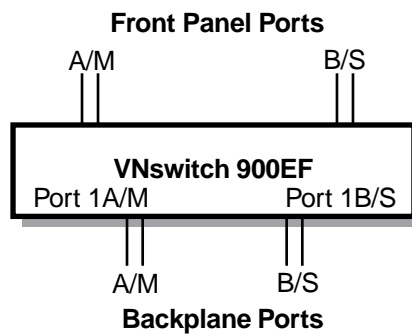

Backplane Configurations

By default, all ports (including the FDDI ports) on the VNswitch 900EF module are active on the front panel. The network manager can issue the appropriate clearVISN commands to establish FDDI networks across the backplane of the DEChub 900 MultiSwitch by redirecting (or assigning) one or both FDDI ports to the backplane.

When an FDDI port is assigned to the backplane, it is referred to as a backplane port. An FDDI backplane port connects to two backplane LAN segments (flexible channels) because two point-to-point connections are required between any two FDDI ports.

Port Assignments

Figure 5-5 illustrates the various ways the VNswitch 900EF module's FDDI ports (port 1A/M and port 1B/S) can be assigned as two front panel ports, two backplane ports, or a front panel port and a backplane port.

Figure 5-5: VNswitch 900EF Port Configuration Capabilities

Front Panel Ports

A/M   B/S

```
┌─────────────────────────────────┐
│        VNswitch 900EF           │
│ Port 1A/M           Port 1B/S   │
└─────────────────────────────────┘
```

A/M   B/S

Backplane Ports

LKG-10238-96

Types of Backplane Configurations

Technically, any type of configuration (dual ring, tree, or dual ring of trees) can be created in the DEChub 900 backplane. However, for ease of use, only dual ring and tree configurations are supported.

Up to four independent FDDI networks can be created across the DEChub 900 backplane. Each of the created FDDI networks can be either a dual ring or a tree on the backplane, but not a mixture of both.

A treed FDDI DEChub 900 backplane can connect to a dual ring that is external to the DEChub 900, and can be part of an overall dual ring of trees configuration.
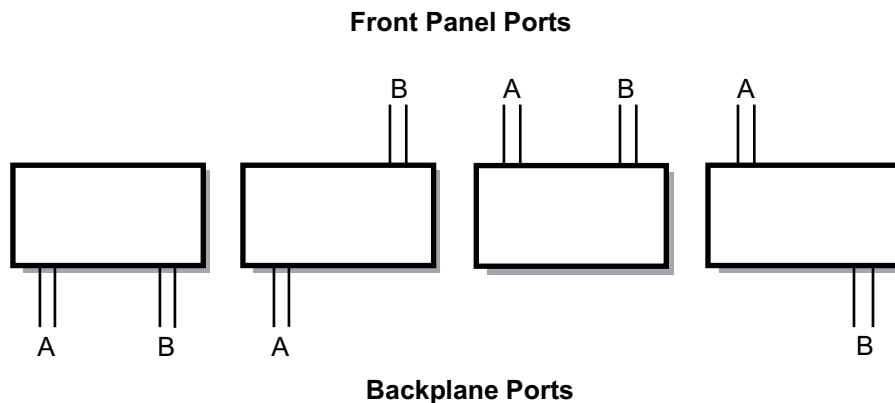
Dual Ring Configurations

This section provides several examples of dual ring configurations. When configured in a dual ring topology, ports are assigned a ring port (A port or B port) by management.

Figure 5-6 shows the four basic dual ring port configurations that can be assigned via management software.

Figure 5-6: Dual Ring Port Configurations

Front Panel Ports
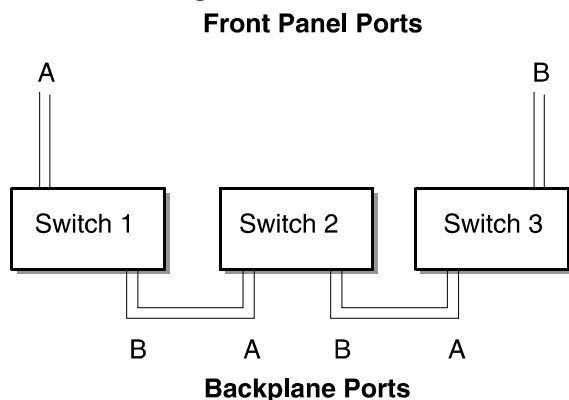

Backplane Ports

LKG-10063-96

Multiple VNswitch 900EF Modules with Connection to External Dual Ring

As shown in Figure 5-7, all of the VNswitch 900EF modules (switches 1 through 3) can connect directly to the dual ring for this dual ring configuration:

• Switch 3 has its B port assigned to the front panel port connection and the A port assigned to the backplane port connection.

• Switch 2 has its A and B ports assigned to the backplane ports only (the module's front panel ports are not active, nor usable).

• Switch 1 has its B port assigned to a backplane port and its A port is assigned to the front panel port.

You can also assign both A and B ports to the front panel ports, which allows a single VNswitch 900EF or DECconcentrator in a DEChub 900 to connect to an external FDDI backbone network.

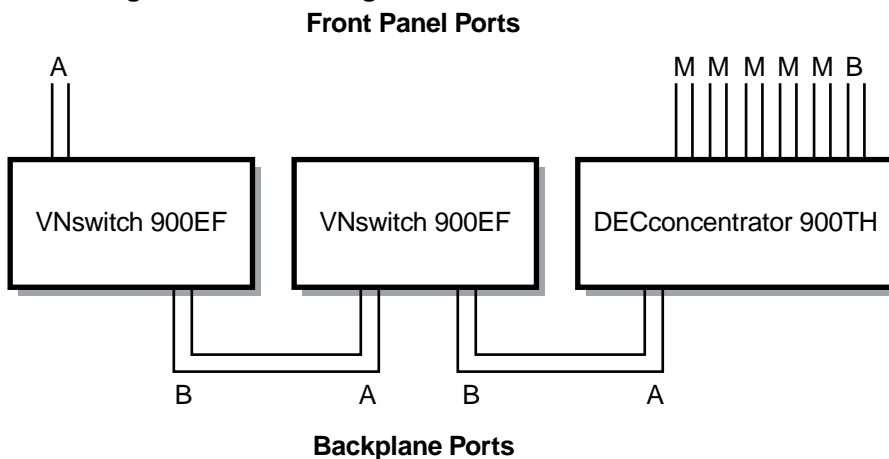**Figure 5-7: Dual Ring Connections for VNswitch 900EF Front Panel Ports**

A   B

┌──────────┐   ┌──────────┐   ┌──────────┐
│ Switch 1 │   │ Switch 2 │   │ Switch 3 │
└──────────┘   └──────────┘   └──────────┘

         B   A   B   A
              Backplane Ports

LKG-10086-96

Two VNswitch 900EF Modules with an FDDI Concentrator

In this configuration, two VNswitch 900EF modules are configured into a dual ring, with a DECconcentrator making the connection to an external FDDI network. As shown in Figure 5-8, the FDDI dual ring connections are made by means of the DEChub 900 backplane: (1) the dual ring enters through the DECconcentrator 900TH; (2) it then passes through the DEChub 900 backplane to the first VNswitch 900EF; and (3) from the first VNswitch 900EF, it again passes through the DEChub 900 backplane to exit through the second VNswitch 900EF. Note that although both front panel FDDI ports on the first VNswitch are disabled, all 12 front panel Ethernet ports can still be used.
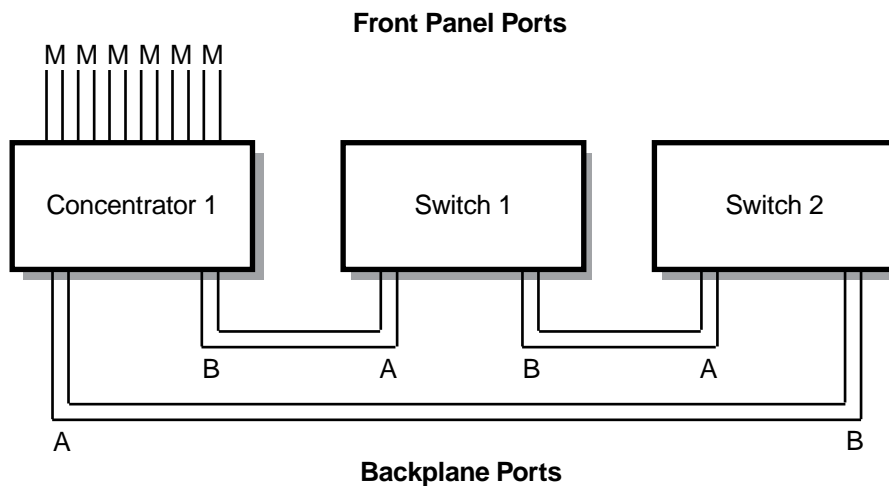
Figure 5-8: Dual Ring Connections for VNswitch 900EF Front Panel Ports

A   M M M M M B

┌──────────────┐ ┌──────────────┐ ┌────────────────────┐
│ VNswitch 900EF│ │ VNswitch 900EF│ │ DECconcentrator 900TH│
└──────────────┘ └──────────────┘ └────────────────────┘

        B   A   B   A
              Backplane Ports

LKG-10242-96

Self-Contained Dual Ring in the Backplane

You can create a dual ring that is completely contained in the DEChub 900 backplane by connecting all the backplane ports of the FDDI network modules (Figure 5-9).

Note that the DECconcentrator in Figure 5-9 is providing M port tree connections to externally treed FDDI stations (such as bridges, concentrators, or end user stations that use S, A, or B ports). The concentrator is not connected to an external M port such as an external concentrator's M port.

Figure 5-9: Self-Contained Dual Ring in the Backplane

Front Panel Ports

M M M M M M

| | | | | | | | | | | |

Concentrator 1   Switch 1   Switch 2

B   A   B   A

A   B

Backplane Ports

LKG-10243-96

Maintaining Connectivity in Dual Rings

When a VNswitch 900EF module in a dual ring is removed (hot swapped), the FDDI ring wraps according to FDDI standards. If the removed module has both A and B ports that are connected to the DEChub 900 backplane, the clearVISN MultiChassis Manager detects the removal and reconnects the two neighboring (upstream and downstream) FDDI stations.

If a VNswitch 900EF module is detected as being unavailable, the MultiChassis Manager also patches out the module (as long as Auto Healing has been enabled).

This is a distinct advantage that VNswitch 900EF modules (configured into a dual ring in a DEChub 900) have over individual nonhub-based FDDI network modules.

NOTE

Because ring wrap occurs much faster than backplane healing, it always occurs first. Thereafter, the MultiChassis Manager firmware automatically reconnects the disconnected DAS stations through the DEChub 900 backplane, allowing the dual ring to unwrap. Fault recovery (unwrapping of the dual ring in the DEChub 900 backplane) is supported by the VNswitch 900EF module.


Example of Dual Ring Backplane Healing

As shown in Figure 5-10, if concentrator 1 is removed, concentrator 2 and switch 1 continue communications with each other and with all other FDDI stations on the DAS backbone. This occurs because concentrator 2 (sensing the change) wraps the FDDI dual ring.

If concentrator 2 fails, concentrator 1 and switch 1 wrap the dual ring and communications can continue. Within seconds of this particular ring wrap, where concentrator 2 has become unavailable, the MultiChassis Manager automatically reconnects switch 1 to concentrator 1 and the dual ring unwraps.

If, at a later time, a replacement concentrator of the same type is installed into the same slot that held concentrator 2, the replacement concentrator will be automatically reinserted into the ring through the Auto Healing feature.

Figure 5-10: Ring Configuration
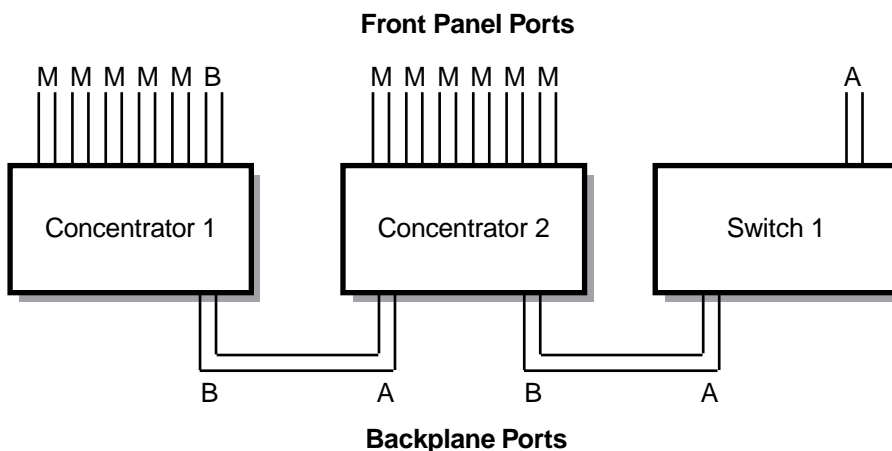
LKG-10093-96

Tree Configurations

When connected in a tree configuration, front panel ports and backplane ports on the VNswitch 900EF are assigned a tree port (M or S port). Specific instructions on how to assign a tree port are provided in the *clearVISN Configuration and Use* guide.

Figure 5-11 shows the five different tree port configurations that can be assigned via management software.

Figure 5-11: Tree Port Configurations

Front Panel Ports
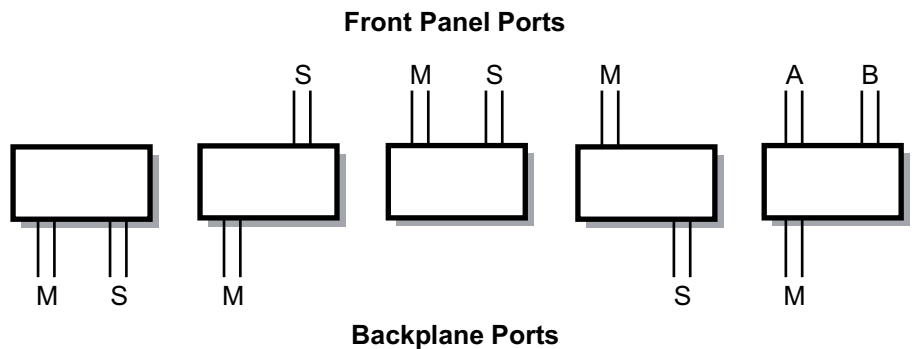
Backplane Ports

LKG-10064-96

The following sections provide examples of FDDI tree configurations with the VNswitch 900EF installed in a DEChub 900 Multiswitch. The DEChub 900 backplane can support many types of FDDI tree configurations. FDDI trees can be contained within the hub backplane and can connect an individual tree in the hub backplane to a higher level of a tree, or can be used to interconnect other hubs.

Self-Contained Tree of Switches and Concentrators

In this example (Figure 5-12), an FDDI tree is constructed entirely within the hub backplane; however, a mixture of FDDI network module types (switches and concentrators) are used. The backplane ports of the VNswitch 900EF modules are enabled as M ports and S ports. The front panel ports of the concentrators become M ports and the front panel FDDI ports of the switches become unusable.

The module that is installed in the highest slot number of the hub (the VNswitch 900EF labeled A) is designated as the *top of the tree*, and its M port connects to the S port of the module with the second highest slot number (the VNswitch 900EF labeled B).

Two hub backplane channels are required for each connection. In Figure 5-12, a total of four backplane channels are in use.

**Figure 5-12: Tree Connections with Switches and Concentrators**

Front Panel Ports

M M M M M M

DECconcentrator 900TH   Module B   VNswitch 900EF   Module A   VNswitch 900EF

M   S   M   S   M   S

Backplane Ports
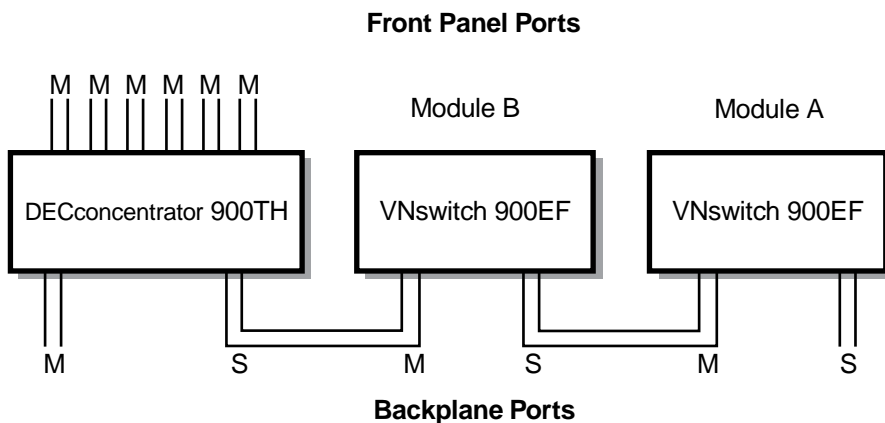
LKG-10239-96

Multiple VNswitch 900EF Modules Used to Create Tree Connections to an External Concentrator

To connect FDDI trees (constructed within the DEChub 900 backplane) to an external higher level concentrator, you can configure one of the VNswitch 900EF modules with a front panel S port, and all other modules with backplane M and S ports.

Figure 5-13 illustrates a tree of VNswitch 900EF modules connected to an external FDDI ring.

Figure 5-13: Tree Connections to an External FDDI Network

Front Panel Ports

S

VNswitch 900EF   VNswitch 900EF   VNswitch 900EF

M   S   M   S   M

Backplane Ports
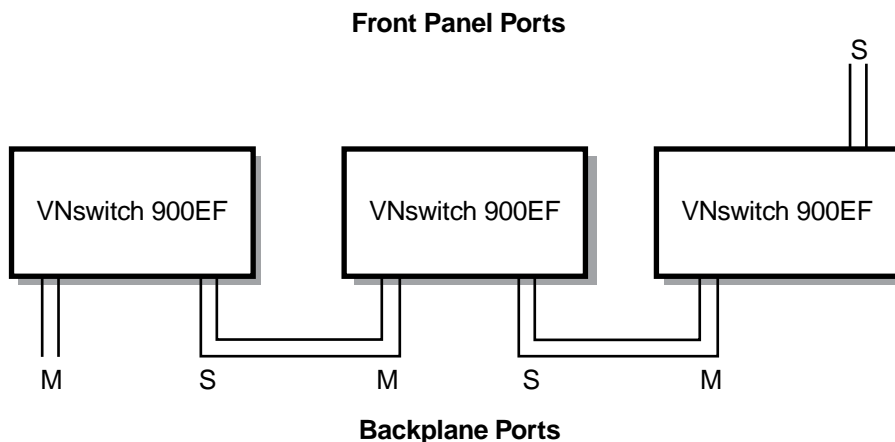
LKG-10240-96

**Dual Homing and Connection to an External Dual Ring**

In some cases, it is desirable to connect the hub to an external dual ring, and configure the network modules as FDDI trees within the hub backplane. In other cases, it is also desirable to have a redundant connection from the hub to higher level concentrators.

This second case (Figure 5-14) is commonly referred to as dual homing. Either of these configurations requires the use of a concentrator at the top level of the FDDI tree configuration.

NOTE

The VNswitch 900EF modules that are connected through the DEChub 900 backplane are Single Attachment Station (SAS) modules. The term *dual homing,* in this case, refers to the connection between the hub-based tree and the external concentrators.

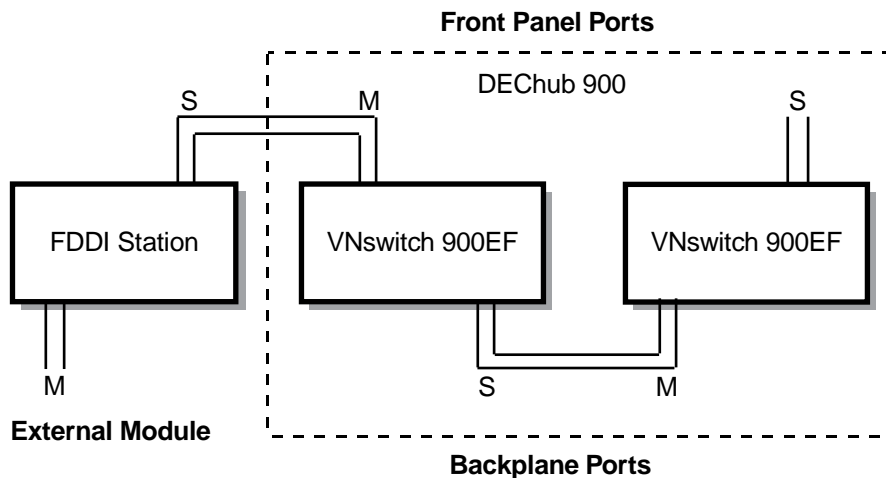**Figure 5-14: Dual Homed Connections to an FDDI Network**
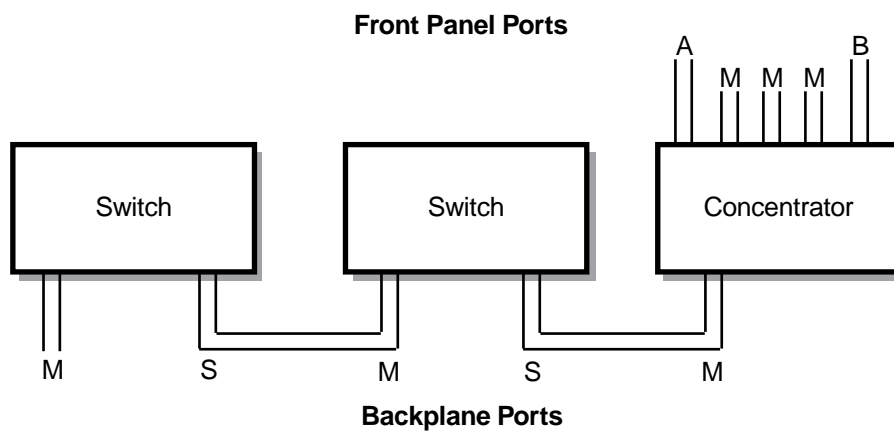
LKG-10080-96

VNswitch 900EF in an FDDI Tree Extension Configuration

You can configure a VNswitch 900EF module that is at the end of a branch of an FDDI tree, with front panel M ports. The M port can then be used to connect to an external FDDI network module or to an additional hub (Figure 5-15).

This configuration is used in situations where it is necessary to configure multiple hubs or VNswitch 900EF modules into trees (for example, when more than eight modules are required in the same closet).

Figure 5-15: Hub-Based Tree Connections to an External FDDI Network

Front Panel Ports

DEChub 900

| S | M | | S |

| FDDI Station | VNswitch 900EF | VNswitch 900EF |

| M | S | M |

External Module

Backplane Ports

LKG-10241-96

Maintaining Connectivity in Trees

When a VNswitch 900EF module in a tree is removed (hot swapped), the FDDI tree splits into two independent networks, per the FDDI standard.

If the removed module has both M and S ports that are connected to the DEChub 900 backplane, the clearVISN MultiChassis Manager detects the removal and reconnects the two neighboring (upstream and downstream) FDDI stations.

If the switch is detected as being unavailable, the MultiChassis Manager also patches out the module (as long as Auto Healing has been enabled).

This is a distinct advantage that the VNswitch 900EF module (configured into a tree in a DEChub) has over nonhub-based FDDI network modules.
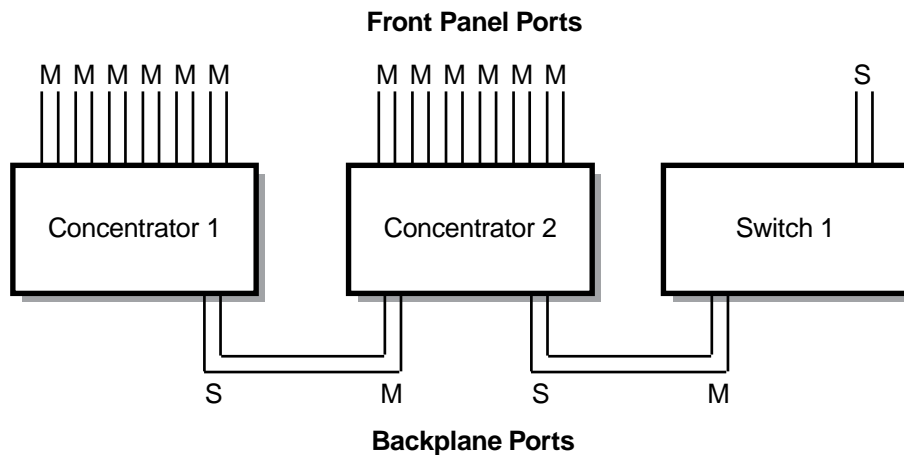
Example of Tree Backplane Healing

As shown in Figure 5-16, if concentrator 1 is removed, concentrator 2 and switch 1 continue communications with each other and with all other FDDI stations in the tree.

If concentrator 2 fails, concentrator 1 and switch 1 can no longer communicate with each other. Within seconds of when concentrator 2 becomes *unavailable,* the MultiChassis Manager automatically reconnects switch 1 to concentrator 1 and the tree is reestablished.

If, at a later time, a replacement concentrator of the same type is installed into the same slot that held concentrator 2, the replacement concentrator will be reinserted automatically into the tree through the Auto Healing feature.

Figure 5-16: Tree Configuration

Front Panel Ports

M M M M M M   M M M M M M   S

| Concentrator 1 | Concentrator 2 | Switch 1 |

S   M   S   M

Backplane Ports

LKG-10125-95

For More Information

Topic	Where to Get More Information
Features, topologies, and components of the FDDI local area network standard	Refer to *A Primer on FDDI: Fiber Distributed Data Interface.*
Digital Equipment Corporation's DECconcentrator 900 FDDI network modules	Refer to the *DEChub 900 Network Modules, Concentrator Reference* manual.

Ethernet-to-FDDI Configuration Examples

Introduction

The VNswitch 900EF module offers many flexible configuration options and is designed for use in different switching environments, including Ethernet-to-FDDI, Fast Ethernet to FDDI, and FDDI-to-ATM (when used with the VNswitch 900EX and EA, respectively). This section provides several examples of how the VNswitch 900EF module might be used in an Ethernet-to-FDDI switching environment.

NOTE

This section provides configuration examples for Version 1.5 of the VNswitch 900 family of switches. Examples of configurations using the optional VNswitch routing software are provided in Chapter 7.

Standalone VNswitch 900EF Module

Figure 5-17 illustrates a configuration in which a single standalone VNswitch 900EF module (installed in a DEChub ONE docking station) is used as a hub to provide dedicated switching between multiple Ethernet LANs and a high-speed server.

This configuration may be suitable for a small- to medium-sized facility that needs to provide several users with personal Ethernet connections, while at the same time grouping other users into multiple virtual LANs.

With the VNswitch 900EF module's filtering features, the network can be configured to allow only the personal Ethernet users access to the server, or could be configured to prevent only certain users or VLANs from accessing the server.

Figure 5-17: Standalone VNswitch 900EF Module
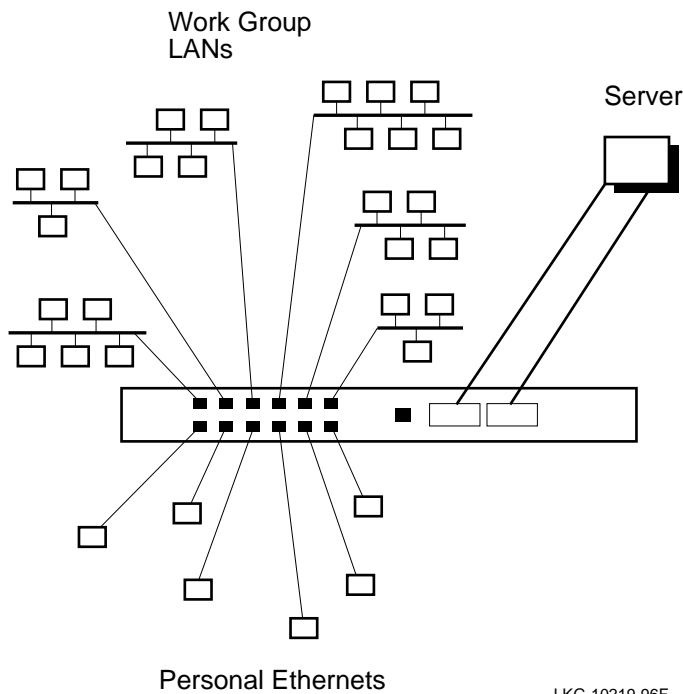
Work Group LANs

Server

Personal Ethernets   LKG-10219-96F

Multiple Shared Ethernet

The following sections provide examples of using the VNswitch 900EF installed in a DEChub 900 along with non-VNswitch modules. Three multiple shared Ethernet environments are illustrated:
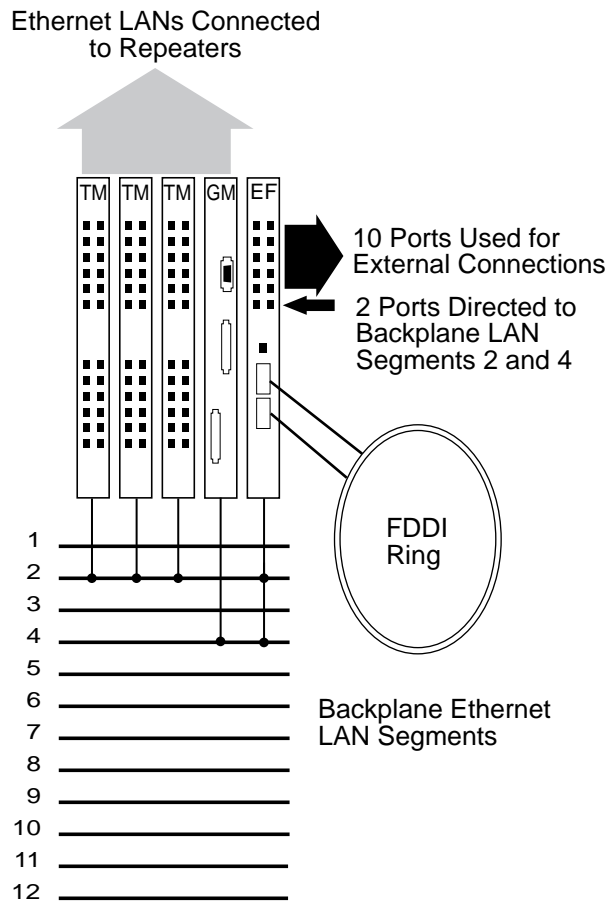
- Low-throughput ring configuration

- Medium-throughput dual-homed configuration

- High-throughput configuration

Low-Throughput Ring Configuration

Figure 5-18 illustrates a low-throughput configuration in which a single VNswitch 900EF module shares the hub's resources with several repeaters. In this configuration, two of the VNswitch module's front panel Ethernet ports are directed to backplane LAN segments, allowing the repeaters to connect to the VNswitch 900EF and access an FDDI ring.

Figure 5-18: Low-Throughput Ethernet-to-FDDI Switching Examples

Ethernet LANs Connected
to Repeaters

| TM | TM | TM | GM | EF |

10 Ports Used for
External Connections

2 Ports Directed to
Backplane LAN
Segments 2 and 4

FDDI
Ring

1
2
3
4
5
6   Backplane Ethernet
7   LAN Segments
8
9
10
11
12

LKG-10233-96F

The DECrepeater 900GM and 900TM modules in this example support connections to the first six backplane Ethernet LAN segments only (flexible channels 0–5). Thus, as shown in Figure 5-18, backplane LANs are created on LAN segments 2 and 4 (flex channels 1 and 3). The three DECrepeater 900TM modules are connected on LAN segment 2 along with a single port from the VNswitch 900EF. The DECrepeater 900GM is connected to LAN segment 4 along with another port from the VNswitch to form a second backplane LAN.

In this configuration, all LANs connected to the repeaters (and the VNswitch 900EF module) have access to the FDDI ring attached to the VNswitch 900EF module's high-speed FDDI ports. If the VNswitch ports directed to the backplane are members of a VLAN, the repeater LANs connected to the same backplane LAN automatically become members of the VLAN.
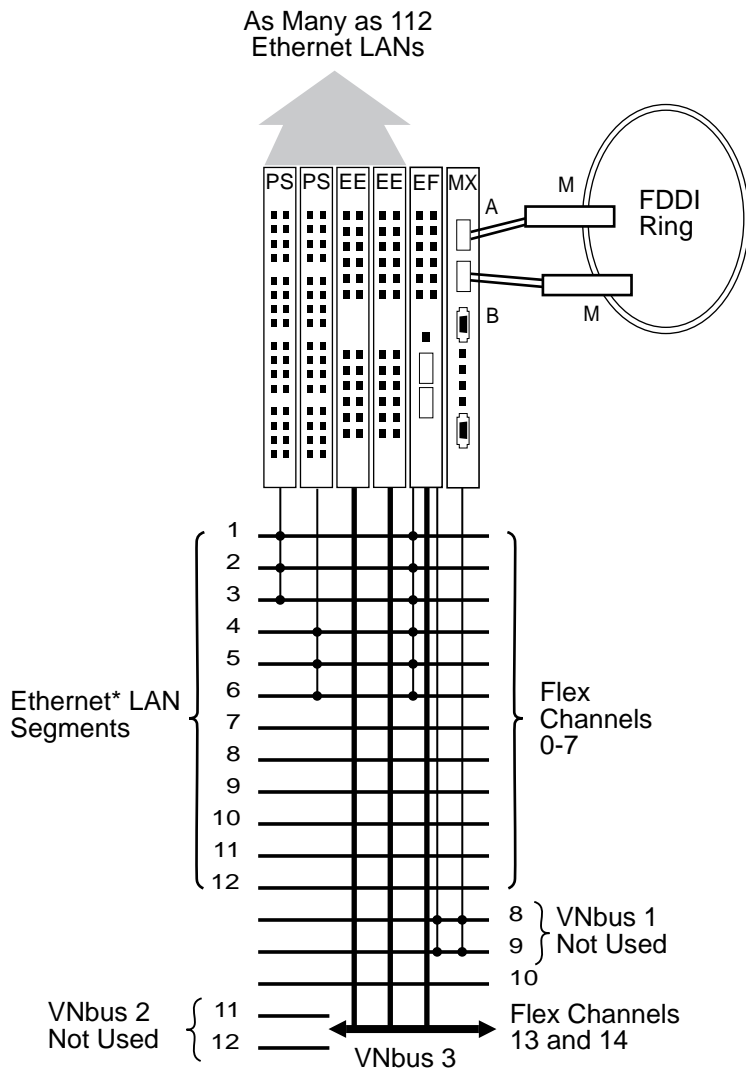
Medium-Throughput Dual-Homed Tree Configuration

Figure 5-19 illustrates a medium throughput configuration in which the FDDI ports on a VNswitch 900EF module are directed to the backplane. In this example, both FDDI ports (A and B) are connected to a backplane LAN with an FDDI concentrator in a dual-homed tree configuration. The DECconcentrator 900MX and VNswitch 900EF FDDI ports (in this illustration) use two flexible channels (8 and 9) to form a single backplane FDDI LAN.

All 12 Ethernet ports on the VNswitch 900EF module are directed to the backplane to provide FDDI support for the port switching repeaters. The VNswitch 900EE modules connect to the VNswitch 900EF over a single VNbus. This configuration supports as many as 112 Ethernet LANs (any mix of personal Ethernets or Ethernet work groups).

Figure 5-19: Medium-Throughput Ethernet-to-FDDI Switching Example

As Many as 112
Ethernet LANs

PS PS EE EE EF MX

A
M

FDDI
Ring

B
M

Ethernet* LAN
Segments

1
2
3
4
5
6
7
8
9
10
11
12

Flex
Channels
0-7

8 } VNbus 1
9 } Not Used
10

VNbus 2
Not Used

11
12

Flex Channels
13 and 14

VNbus 3

*Can also be used for FDDI LANs

LKG-10232-96F

High-Throughput Multiple Network Configuration

Figure 5-20 illustrates a high throughput configuration in which two VNswitch 900EF modules are installed in a DEChub 900 with multiple VNswitch 900EE modules and PORTswitch 900TP/12 repeaters. In this example, multiple servers are attached to two distinct FDDI networks, and hundreds or thousands of users can access the servers through the VNswitch 900EF modules.

All three VNbuses are used in this configuration; hence, three separate networks are formed. Although the FDDI networks in this example are separate from each other (because the EF modules are connected to different VNbuses), this configuration allows the repeater ports to be dynamically allocated to either FDDI network. As shown in Figure 5-20, six Ethernet ports on each VNswitch 900EF module are directed to the backplane to join backplane LANs with the port switching repeaters, enabling the repeaters to access both high-speed FDDI links.

Adjacent to each EF module, a VNswitch 900EE module is configured on the same VNbus (VNbus 2 and VNbus 3, respectively). Two more VNswitch 900EE modules are connected to VNbus 1 to form an all-switched Ethernet network that is completely separate from the networks on VNbus 2 and VNbus 3.

This multiple network configuration provides a total of 172 Ethernet ports for external connections:
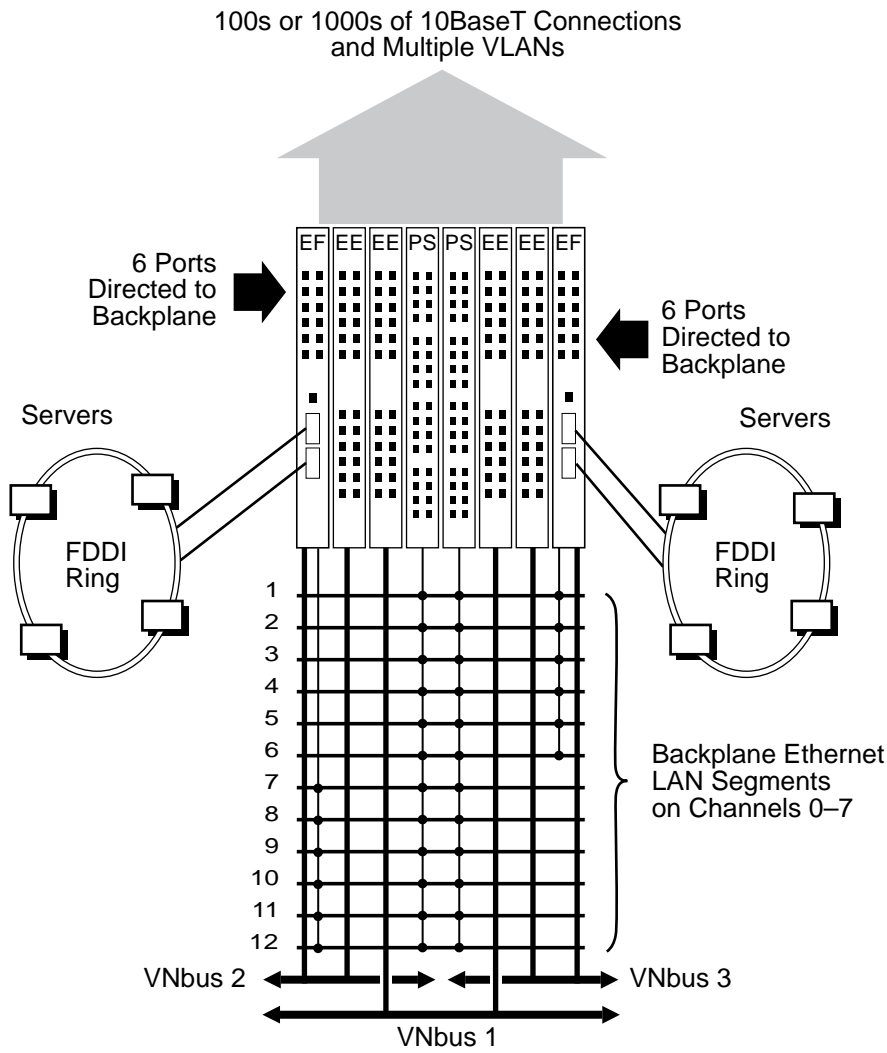
- 12 ports on the two VNswitch 900EF modules (6 front panel Ethernet ports from each module)

- 64 ports on the PORTswitch 900 repeaters (32 ports on each repeater)

- 96 ports on the VNswitch 900EE modules (4 x 24 ports)

To achieve optimum throughput, the LANs attached to the 172 ports could be configured into 10–20 port-based VLANs by grouping their associated ports into port groups or Virtual Secure Domains (VSDs) within each network.

Each VNbus supports a total of 63 VSDs (or Class 1 port-based VLANs). By employing all three VNbuses and configuring multiple VLANs based on users' networking needs (such as how often a particular work group or department requires access to the high-speed servers), you can design a flexible, high-throughput, multiple network configuration.

Figure 5-20: High-Throughput Ethernet-to-FDDI Configuration Example

100s or 1000s of 10BaseT Connections
and Multiple VLANs

6 Ports
Directed to
Backplane

6 Ports
Directed to
Backplane

Servers

Servers

FDDI
Ring

FDDI
Ring

Backplane Ethernet
LAN Segments
on Channels 0–7

VNbus 2

VNbus 3

VNbus 1

LKG-10231-96F

Personal Ethernet to FDDI Backbone

Figure 5-21 illustrates a network configuration in which a VNswitch 900EF module functions as a high-speed collapsed backbone connecting 180 personal Ethernets to multiple servers attached to an FDDI ring. In this example, all of the modules are connected by a single VNbus.

Figure 5-21: High-Speed Personal Ethernet to FDDI Backbone Example

LKG-10230-96F
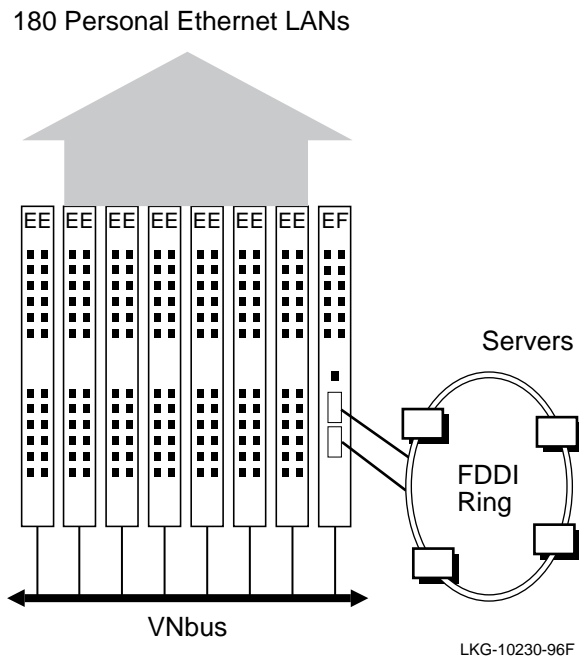
High-Performance FDDI Switched Backbone

Figure 5-22 illustrates multiple VNswitch 900EF modules used with a GIGAswitch/FDDI to create a high-performance FDDI switched backbone. In this example, two VNswitch 900EF modules are configured with VNswitch 900EE modules and connected over the VNbus. By keeping all the modules in the default VLAN, all the LANs connected to the VNswitch modules have access to the GIGAswitch/FDDI.

Figure 5-22: High-Performance FDDI Switched Backbone
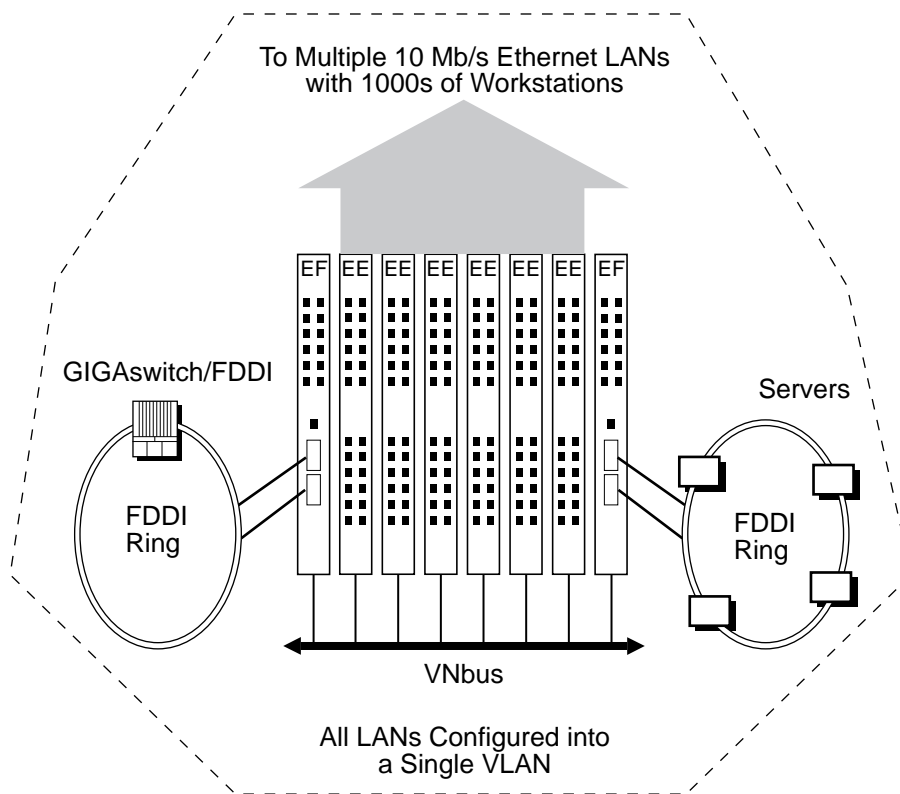
LKG-10298-96F

Summary

In this chapter, we illustrated just a few of the many possible uses of the VNswitch 900EF module in an Ethernet-to-FDDI switching environment. However, the VNswitch 900EF module is designed for use with all four technologies (Ethernet, Fast

Ethernet, FDDI, and ATM) and can be used in any combination of Ethernet-to-FDDI, Fast Ethernet to FDDI, and FDDI-to-ATM environments. For examples of using the VNswitch 900EF module in these other environments, refer to Chapter 7.

For More Information

Topic	Where to Get More Information
Installing the VNswitch 900EF module	Refer to the *VNswitch 900EF Installation and Configuration* guide.
Configuring and using the VNswitch 900EF module	Refer to the *VNswitch 900 Series Switch Management* guide.

Chapter 6

Ethernet-to-ATM Switching

Overview

Introduction

Asynchronous Transfer Mode (ATM) is a switch-based technology that provides for the transmission of data over virtual circuits at very high speeds.   This chapter provides a brief overview of ATM and illustrates various configurations of the VNswitch 900EA module in an Ethernet-to-ATM switching environment.

In This Chapter

The following topics are covered in this chapter:

Topic	Page
ATM Overview	6-2
Ethernet-to-ATM Configuration Examples	6-6

ATM Overview

Introduction

ATM is a layered architecture allowing multiple services such as voice, data, and video to be mixed and transmitted over a network. Information is transmitted by switching data over virtual circuits. By providing connectivity through a switch (instead of a shared bus), ATM offers the following benefits:

- Dedicated bandwidth per connection

- Higher aggregate bandwidth

- Flexible access speeds

- Well-defined connection procedures

How ATM Works

ATM works by setting up virtual point-to-point connections between two end points. When information needs to be sent, the sender negotiates a requested path with the network for a connection to the destination. When setting up the connection, the sender specifies the type, speed, and other attributes of the call.

Once the connection is set up, data is segmented into short, fixed-length cells and transported over a virtual circuit (VC), based on the destination address in the cell. Cells can handle any digital information—voice, data, images, or video. At the other end of the point-to-point connection, cells are reassembled into the original information.

NOTE

For a description of the ATM cell format, refer to Chapter 3.

ATM Layers

ATM consists of the following layers:

- ATM Adaption Layer — The ATM Adapation Layer (AAL) specifies the packet format and method of segmenting packets into cells. It is also responsible for reassembling the packets at the destination.

- ATM Layer — The ATM Layer (AL) specifies how cells are passed through the network and managed.

- Physical Layer (PHY) — The PHY layer is responsible for converting information to the appropriate electrical or optical format for transmission over the network.

Figure 6-1 illustrates the different layers in the ATM architecture. Note that LAN emulation is built on top of the AAL5 layer, which is the most widely implemented version of AAL.

Figure 6-1: ATM Layers

LAN Emulation (LANE)
Is Built on Top of the
AAL5 Layer

| AAL5 |
| AL |
| PHY |
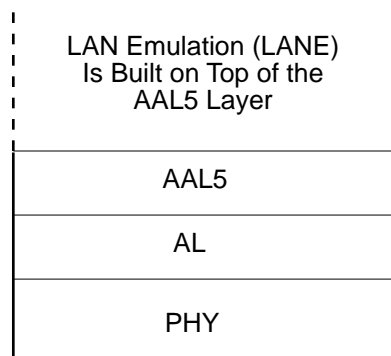
LKG-10299-96F

ATM Virtual Circuits

For ATM communication to take place, a virtual circuit must be established. A virtual circuit is similar to a physical circuit, but with a variable capacity. There are two types of virtual circuits:

- Permanent virtual circuit (PVC) — a virtual circuit that is preassigned on a permanent basis or statically configured via network management

- Switched virtual circuit (SVC) — a virtual circuit that is dynamically set up using a signaling protocol

An ATM network relies on numerous virtual circuits and virtual paths (VPs). A virtual path is the route over which a virtual circuit is established.

When a request is made to set up a virtual circuit, ATM switches (such as the GIGAswitch/ATM) respond by selecting a route capable of meeting the request, and then set aside the appropriate amount of network resources. Once a virtual circuit has been defined and a route selected, cells begin to flow from the source, through the intermediate ATM switches, to the destination. All cells in the virtual circuit travel over the same route, maintaining cell sequence.

LAN Emulation

LAN Emulation (LANE) exploits the switching properties of ATM technology to provide point-to-point, non-shared communications channels between legacy LANs over an ATM network or backbone.

LAN Emulation is performed by converting legacy LAN frames into ATM cells and vice versa. An edge-device (such as the VNswitch 900EA), residing on the edge of the ATM network, translates legacy LAN frames and transmits them as cells onto the ATM network. When the VNswitch 900EA module (or other LANE-compliant edge-device) receives cells from the ATM network, it reassembles the cells into frames and passes them to the legacy LAN. Both the legacy LAN and the ATM Emulated LAN (ELAN) perceive that they are connected to a single bridged LAN.

NOTE

For a description of the Emulated LAN frame format, refer to Chapter 3.

LAN Emulation Components

The LAN Emulation (LANE) service is comprised of the following components:

- LAN Emulation clients (LEC), which reside on each edge-device (such as the VNswitch 900EA switch)

- LAN Emulation Configuration Server (LECS), which provides configuration and virtual LAN information

- LAN Emulation servers (LES), which typically reside on an ATM host, switch, or router, and are responsible for establishing ATM connections

- Broadcast and Unknown Server (BUS), which performs broadcasting and multicasting functions

Though not a LANE component, ATM hosts (which are workstations or other computers equipped to attach directly to the ATM network) are also usually included in a LANE configuration. Figure 6-2 illustrates the various LANE components in an Ethernet-to-ATM switching environment.

Figure 6-2: Components Used In LAN Emulation

GIGAswitch/ATM

LECS

LES   BUS

ATM Host or
Another LEC

Switched
Virtual
Circuits
(SVCs)

LEC

VNswitch 900EA

10 Mb/s Ethernet LAN
Connections and
400 Mb/s VNbus Connection
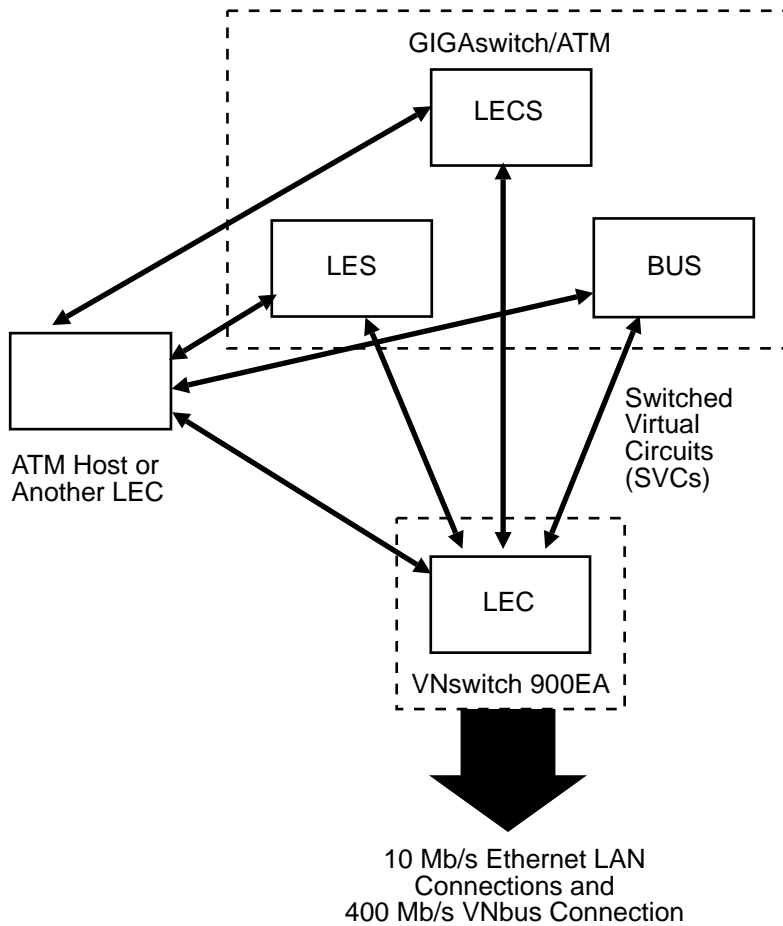
LKG-10300-96F

As shown in this figure, the VNswitch 900EA module implements the LAN Emulation
Client (LEC) and the other LANE components are implemented by the
GIGAswitch/ATM (or by any other LANE-compliant device).

Ethernet-to-ATM Configuration Examples

Introduction

The VNswitch 900EA is a powerful, high-density Ethernet-to-ATM switch that can provide a cost-effective solution for legacy LAN-based networks that need to evolve to ATM technology. The VNswitch 900EA module offers many different network configuration possibilities, ranging from simple tunneling of LAN frames between two legacy LAN devices, to the segmentation of multiple LANs that span across an ATM cloud.

The VNswitch 900EA module is particularly well suited for use in multi-technology environments, and can be used as a hub to connect dissimilar networks. Used as either a standalone module installed in a DEChub ONE docking station, or configured in a DEChub 900 MultiSwitch, the module provides a flexible solution to the bandwidth needs of various Ethernet LANs.

This section provides several examples of how the VNswitch 900EA module might be used in an Ethernet-to-ATM switching environment. For examples of using the VNswitch 900EA module with the VNswitch 900EX and EF modules, refer to Chapter 7.

NOTE

This section provides configuration examples for Version 1.5 of the VNswitch 900 family of switches. For examples of configurations using the optional VNswitch routing software, refer to Chapter 7.

Two Standalone VNswitch 900EA Modules

Figure 6-3 illustrates a configuration in which two VNswitch 900EA modules are connected by a bridge tunnel through an ATM network. In this example, each switch is used as a hub to connect both workgroup LANs and personal Ethernet LANs across the ATM network. Each LAN has a dedicated 10 Mb/s connection to the local switch, but connects to the ATM network through the switch's high-speed ATM port.

In this configuration, only a single bridge tunnel is implemented. However, the VNswitch 900EA module supports as many as 16 bridge tunnels (or any combination of tunnels and ELANs) on its ATM port.

Figure 6-3: Two Standalone VNswitch 900EA Modules
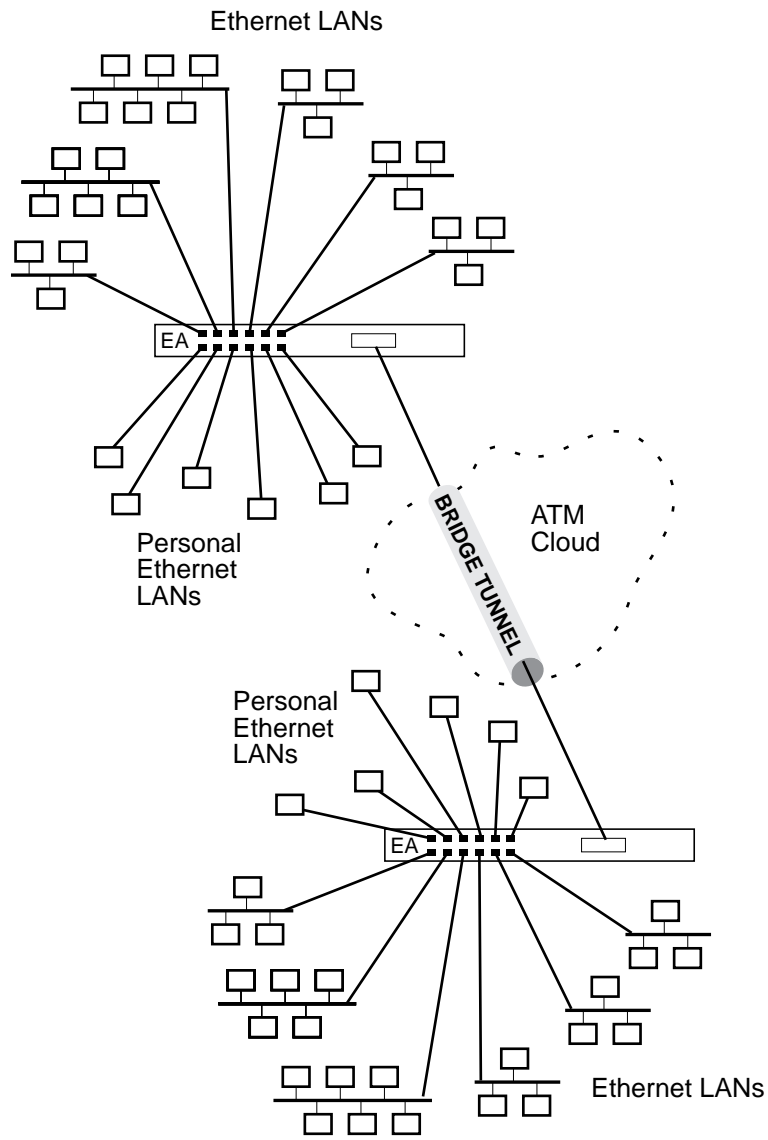
Ethernet LANs

EA

Personal
Ethernet
LANs

ATM
Cloud

BRIDGE TUNNEL

Personal
Ethernet
LANs

EA

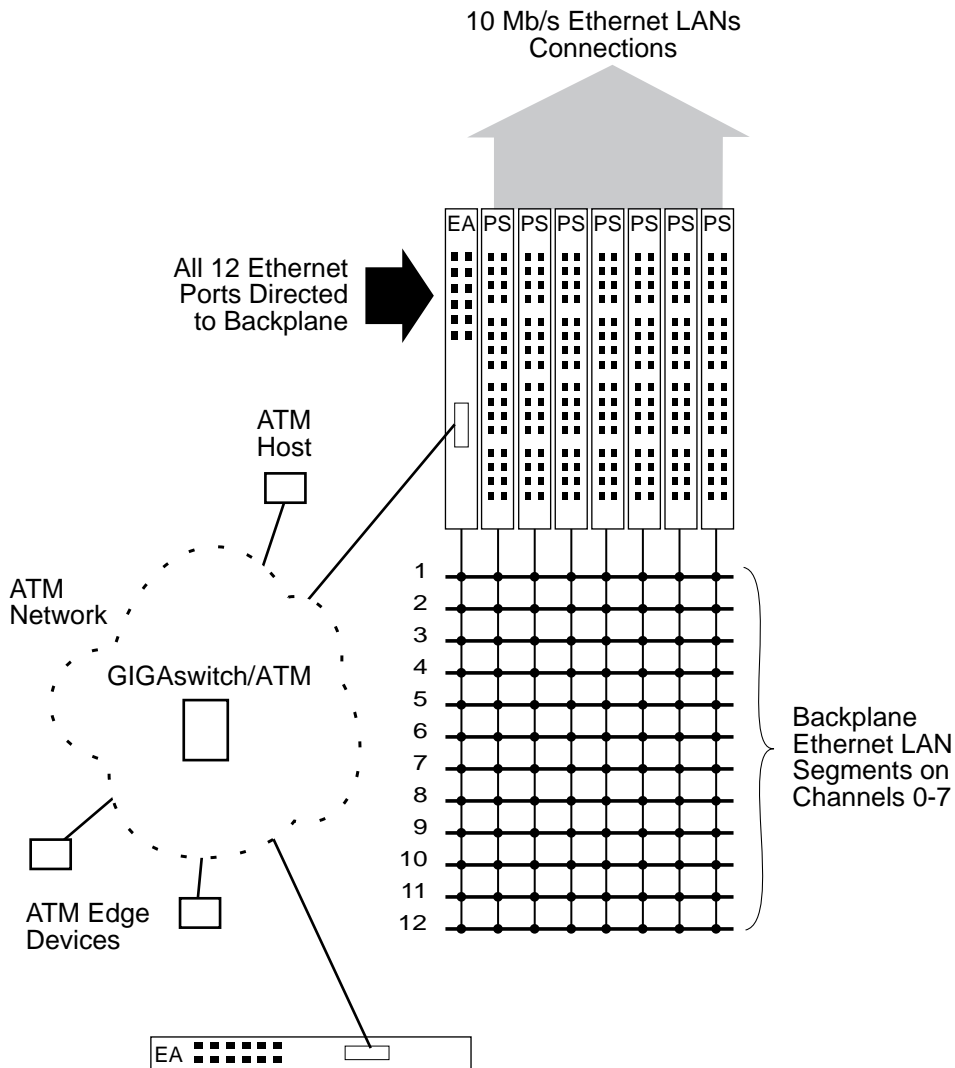Ethernet LANs

LKG10254-96F

Multiple Shared Ethernet

The following sections provide examples of using the VNswitch 900EA installed in a DEChub 900 MultiSwitch along with other non-VNswitch modules. Three multiple shared Ethernet environments are illustrated:

- Low-throughput configuration

- Medium-throughput configuration

- High-throughput configuration

Low-Throughput Configuration

Figure 6-4 illustrates a low-throughput configuration in which a single VNswitch 900EA module shares the hub's resources with multiple port switching repeaters. In this configuration, all front panel Ethernet ports on the VNswitch 900EA module are directed to the backplane LAN segments, allowing the port switching repeaters to connect to the VNswitch 900EA. The LANs connected to the port switching repeaters can be incorporated into Emulated LANs (ELANs) attached to the VNswitch 900EA module's 16 virtual ATM ports.

Figure 6-4: Low-Throughput ATM Switching Example

LKG-10255-96F

The PORTswitch 900TP/12 modules used in this example allow connections to all 12 backplane Ethernet LAN segments. Seven PORTswitch 900TP/12 modules configured in a single hub provide 224 (7 x 32) 10BaseT ports that can connect to the backplane LAN segments (18 per segment) to be joined with ELANs that span across

the ATM network. In this example, a GIGAswitch/ATM provides essential LANE services (such as the LES and BUS) to the VNswitch 900EA module, which supports 16 LAN Emulation Clients (LECs) on its single high-speed ATM port.

Medium-Throughput Configuration

Figure 6-5 illustrates a medium-throughput configuration in which two VNswitch 900EA modules share the hub's resources with DECswitch 900EE switches and port switching repeaters. A single GIGAswitch/ATM provides LANE services to both VNswitch 900EA modules. Flow control implemented by the GIGAswitch/ATM allows the VNswitch modules to achieve their full 155 Mb/s capacity, even after adjusting for Sonet and ATM protocol traffic.

Figure 6-5: Medium-Throughput Ethernet-to-ATM Switching Example

Servers, Printers, and Ethernet LANs

148 Ports of Shared 10BaseT

DECswitch 900EEs

EA | PS | PS | PS | PS | EE | EE | EA

All 12 Ethernet ports directed to backplane

6 ports directed to backplane

6 ports connected to Ethernet LANs

GIGASwitch/ATM

ATM Cloud

ATM Edge Devices

1
2
3
4
5
6
7
8
9
10
11
12

Backplane Ethernet LAN Segments on Channels 0-7
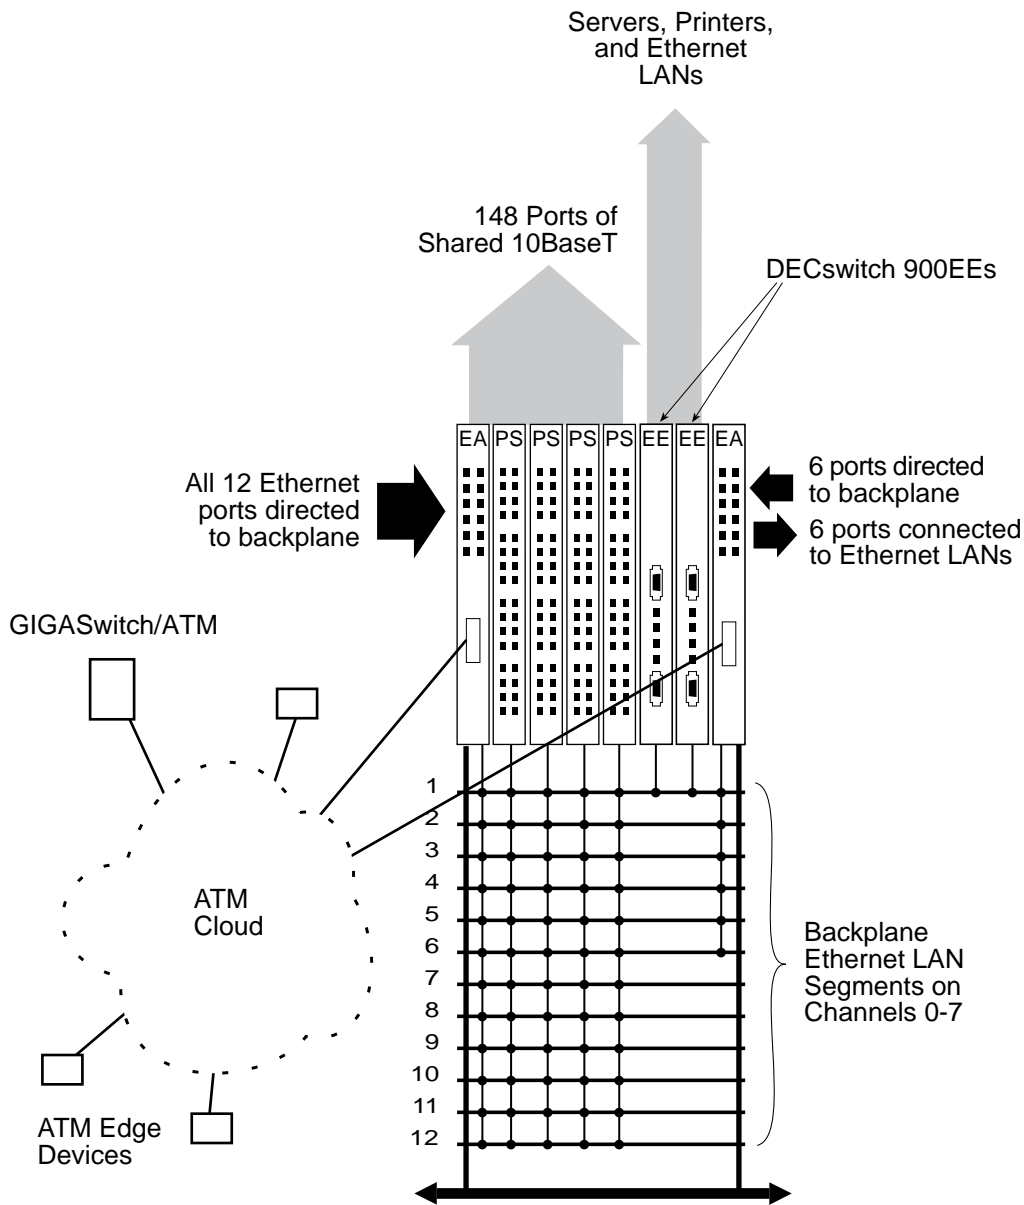
LKG-10256-96F

High-Throughput Configuration

Multiple VNswitch 900EA modules installed in a DEChub 900 with multiple VNswitch 900EE modules provides a viable solution for high-density, high-throughput, switched Ethernet environments comprised of thousands of users.

Figure 6-6 illustrates a high-density, high-throughput Ethernet-to-ATM configuration. In this example, four VNswitch 900EA modules are configured in the DEChub 900 MultiSwitch with four VNswitch 900EE modules.

This robust configuration could be used as part of the backbone of a large enterprise network. In this configuration, all 48 Ethernet ports on the VNswitch 900EA modules are active (none are directed to the backplane). The VNswitch 900EE modules are employed as expansion modules, adding an additional 96 10BaseT ports. Together, the VNswitch EA modules and EE modules provide a total of 144 switched Ethernet ports for personal Ethernet connections (for example, servers and high-performance workstations) and Ethernet work groups.

In this example, all three VNbuses are used to provide connectivity between the EA modules and the EE modules:

- VNbus 1 connects two EA modules and two EE modules.

- VNbus 2 and VNbus 3 each connect a single EA module and EE module.

As a result, three separate network segments are created on the VNbuses. However, all segments are connected to the same ATM network via the high-speed ATM port on each VNswitch 900EA module. By connecting all the EA modules to the same ATM network, all of the LANs connected to the Ethernet ports are connected over the ATM network (unless the ports are divided into multiple VLANs).

Figure 6-6: High-Throughput Ethernet-to-ATM Configuration Example
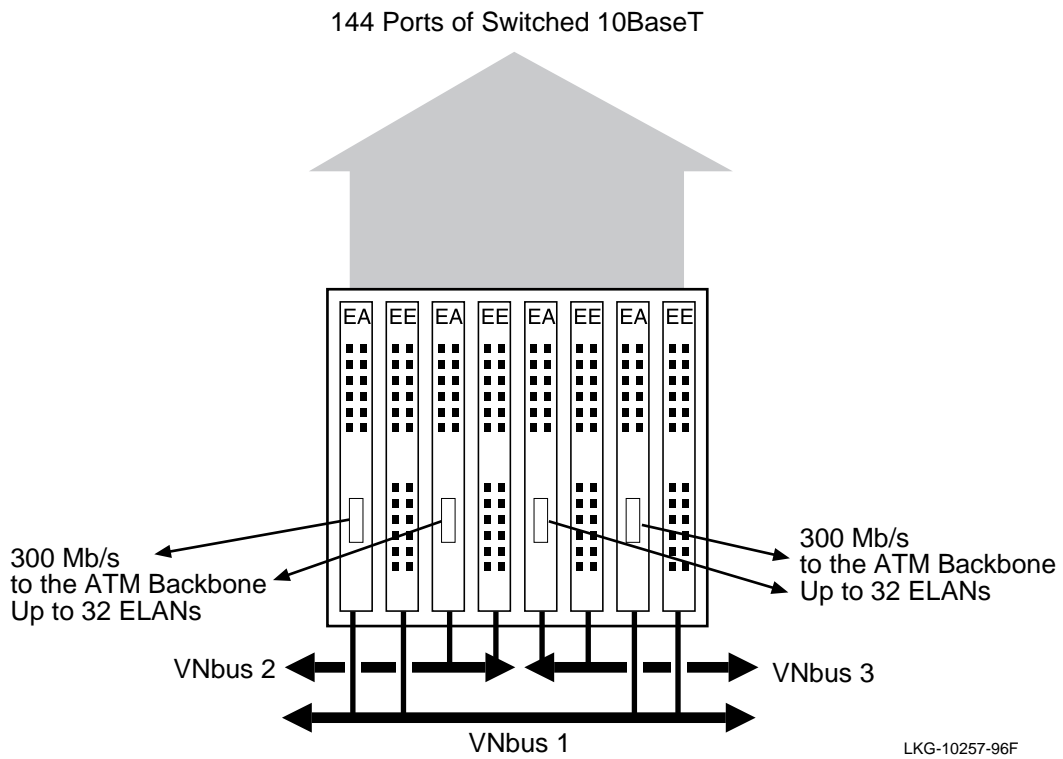
144 Ports of Switched 10BaseT

EA EE EA EE EA EE EA EE

300 Mb/s
to the ATM Backbone
Up to 32 ELANs

300 Mb/s
to the ATM Backbone
Up to 32 ELANs

VNbus 2

VNbus 3

VNbus 1

LKG-10257-96F

High-Density Personal Ethernet-to-ATM Backbone

Figure 6-7 illustrates a network configuration in which VNswitch 900EA modules provide two ATM uplinks for multiple VLANs comprised of personal Ethernets and workgroup LANs. As indicated in the illustration, the personal Ethernet connections could include several servers for local file access.

In this example, 168 personal Ethernet LANs (including the LANs used for the servers) are configured on a single VNbus. The personal Ethernets could be combined into a number of VLANs (or left in the default VLAN) and joined with VLANs on other VNswitch modules across the ATM network by means of the VNswitch 900EA module's LAN Emulation (LANE) features.

The dual ATM links provide connections for as many as 32 separate ELANs (16 for each VNswitch 900EA by means of the 16 virtual ATM ports on each module). Redundancy can be provided by attaching a virtual ATM port from each module to the same ELAN. In this case, spanning tree will place one of the connections into backup.

NOTE

In the *VNswitch 900 Series Switch Management* guide, the virtual ATM ports are referred to as *ATM logical interfaces*.

Figure 6-7: High-Density Personal Ethernet-to-ATM Backbone Example
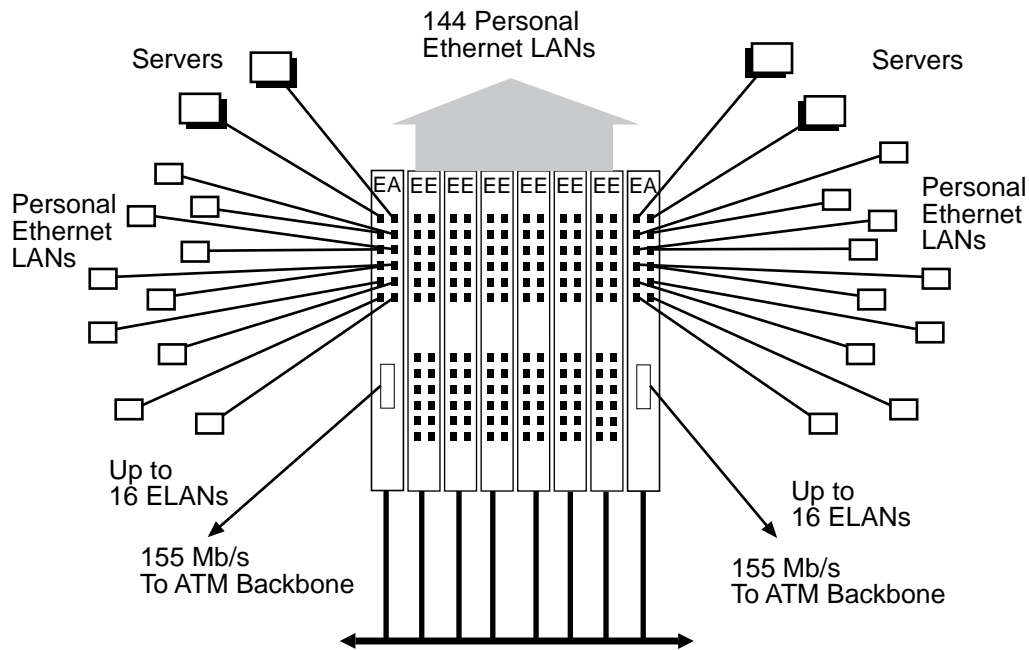
LKG-10258-96F

Shared and Switched Mixed Environment

Figure 6-8 illustrates the flexible configuration options of the VNswitch 900EA module in both shared and switched Ethernet-to-ATM environments. In this example, multiple VNswitch 900EA modules and EE modules are configured with port switching repeaters in a DEChub 900 MultiSwitch. Some of the ports on the VNswitch 900EE modules are used for personal Ethernet LANs, while others can be configured for both small and large work groups.

The VNswitch modules in this example are connected to two distinct VNbuses, allowing traffic between the VNswitch 900EA and its adjacent VNswitch 900EE module to flow at 400 Mb/s, while also allowing VLANs to span across the modules.

All 128 ports on the 32-port PORTswitch 900TP/12 repeaters are connected to backplane Ethernet LAN segments. Several ports on the VNswitch 900 modules are redirected to the backplane to join backplane LANs with the port switching repeaters, while others are used for external LAN connections.

When a port that is a member of a VLAN is directed to the backplane, any other LANs connected to the same backplane LAN segment automatically become members of that VLAN. For example, any stations connected to repeater ports that are linked to a LAN segment occupied by a VLAN become members of the same VLAN. Furthermore, by mapping the VLAN to an ELAN (on the VNswitch 900EA module's ATM port), the LAN connected to the repeater port is incorporated into the ELAN and can share information with other members of the ELAN, including ATM hosts.

Figure 6-8: VNswitch 900EA in a Shared and Switched Environment
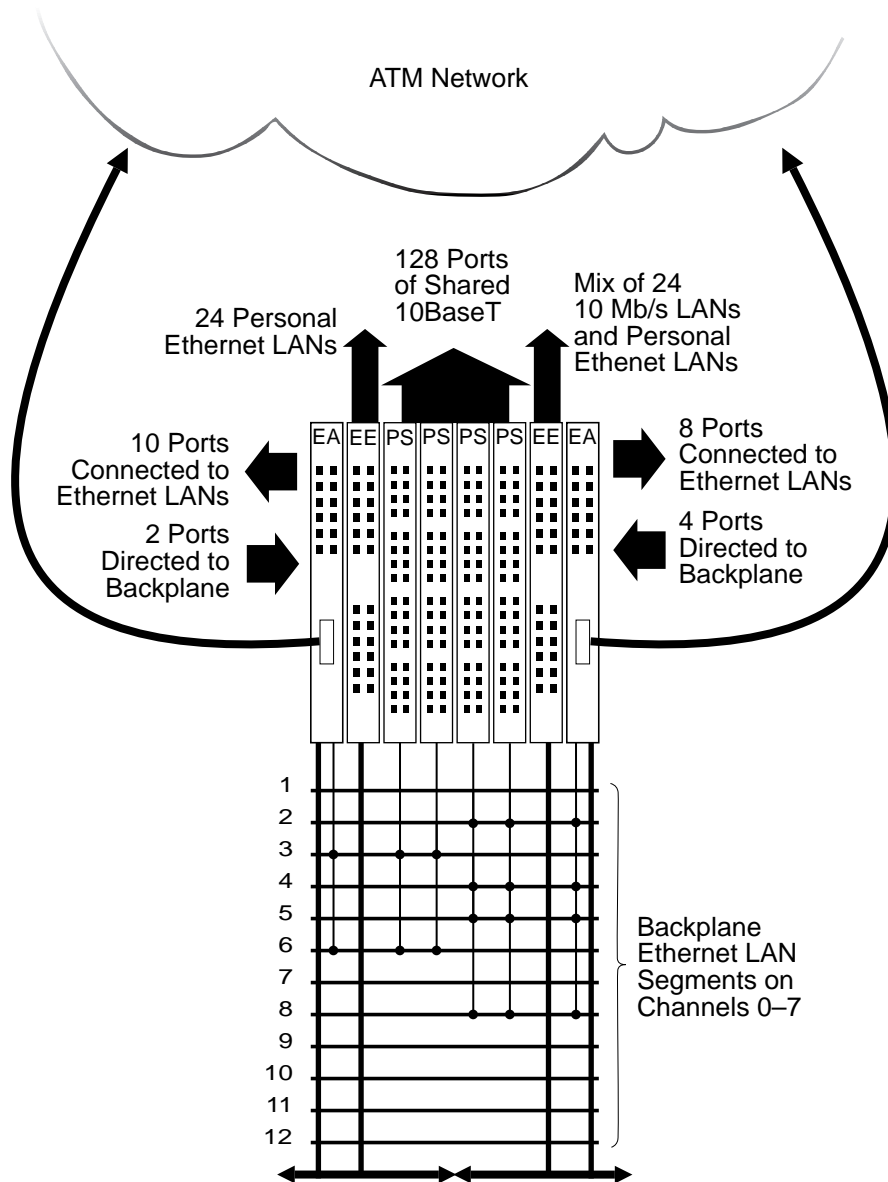
LKG-10224-96F

Summary

In this chapter, we illustrated just a few of the many possible uses of the VNswitch 900EA module in an Ethernet-to-ATM switching environment. However, the VNswitch 900EA module is designed for use with all four technologies (Ethernet, Fast Ethernet, FDDI, and ATM) and can be used in any combination of Ethernet-to-ATM, Fast Ethernet to ATM, and FDDI-to-ATM environments. For examples of using the VNswitch 900EA module in these other environments, refer to Chapter 7.

For More Information

Topic	Where to Get More Information
Installing the VNswitch 900EA module	Refer to the *VNswitch 900EA Installation and Configuration* guide.
Configuring and using the VNswitch 900EA module	Refer to the *VNswitch 900 Series Switch Management* guide.

Chapter 7

Switching Between Multiple Technologies

Overview

Introduction

The VNswitch 900 series network modules are fully interoperable with each other when installed in a DEChub 900 MultiSwitch. Through the DEChub 900 VNbus, the modules are automatically preconfigured to pass 802.3 frames between each other across the backplane, regardless of the technology implemented on the modules' front ports.

This chapter describes some of the possible configurations that can be used to switch LAN traffic between multiple technologies like Fast Ethernet, FDDI, and ATM.

In This Chapter

The following topics are covered in this chapter:

Multi-Technology Collapsed Backbone Configuration

Introduction

This section provides an example of how several VNswitch modules might be configured in a DEChub 900 MultiSwitch backplane to create a multi-technology collapsed backbone for a medium-sized network.

Configuration Example

Figure 7-1 illustrates several VNswitch modules configured in a DEChub 900 to create a multi-technology collapsed backbone. In this example, networks (or network segments) based on all four technologies (Ethernet, Fast Ethernet, FDDI, and ATM) are integrated into a single multi-technology network.

Figure 7-1: Multi-Technology Collapsed Backbone Configuration
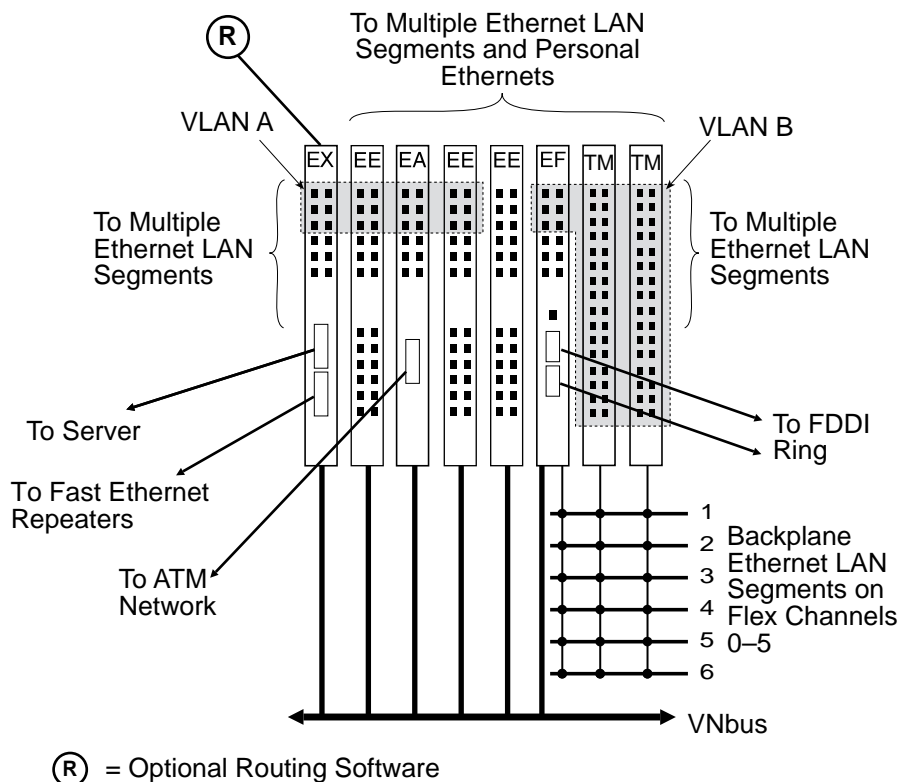
LKG-10285-96F

Multi-Technology Collapsed Backbone Configuration

The DECrepeater 900TM modules in this example are connected to the VNswitch
900EF module using backplane LANs. To accomplish this, the first six 10BaseT ports
on the VNswitch 900EF module are directed to the backplane to join backplane LANs
with the repeaters. These six ports on the VNswitch 900EF module are configured into
a VLAN; hence, when the repeater ports are joined with VNswitch 900EF module via
the same backplane LAN(s), these ports become members of the VLAN.

As shown in Figure 7-1, a single VNbus connects the VNswitch modules. The
VNswitch 900EX module in this example is equipped with the optional routing
software, which provides routing services (over the VNbus) for all the VNswitch
modules. For example, this configuration enables routing between VLAN A and B.

If additional routing capacity is needed, or redundancy is desired, any or all of the
VNswitch modules can be upgraded with the optional routing software. In the example
shown in Figure 7-1, as many as six routers may be configured.

Edge-Switch for Multi-Gigabit Backbone Configurations

Introduction

This section provides several examples of VNswitch modules used as edge-switches in multi-gigabit backbone configurations.

Configuration Examples

Figure 7-2 illustrates a single VNswitch 900EA module and two VNswitch 900EF modules configured as edge-switches attached to a GIGAswitch/ATM and two GIGAswitch/FDDIs, respectively. This example also includes a VNswitch 900EX module and four VNswitch 900EE modules.

Figure 7-2: Multi-Technology Edge-Switch Configuration

In this configuration, all the VNswitch modules are connected to the same VNbus, illustrating how all four technologies can co-exist in the same network. For example, when the modules are configured into a single VLAN (the default VLAN), the LANs attached to the VNswitch ports function as a single LAN, allowing information to flow freely between all attached devices regardless of the different LAN technologies used.

However, for increased throughput and/or security, the modules can be grouped into separate clusters by attaching the modules to different VNbuses. As shown in Figure 7-3, the modules are grouped into three different networks that are fully isolated from each other.

Figure 7-3: Grouping Modules into Separate Clusters or Networks

LKG-10286-96F

Networks formed by attaching the VNswitch modules to separate VNbuses can be joined by bridging the networks together. In Figure 7-4 for example, networks B and C are joined by bridging the modules over the ATM link.

Figure 7-4: Bridging Between VNbuses

LKG-10287-96F
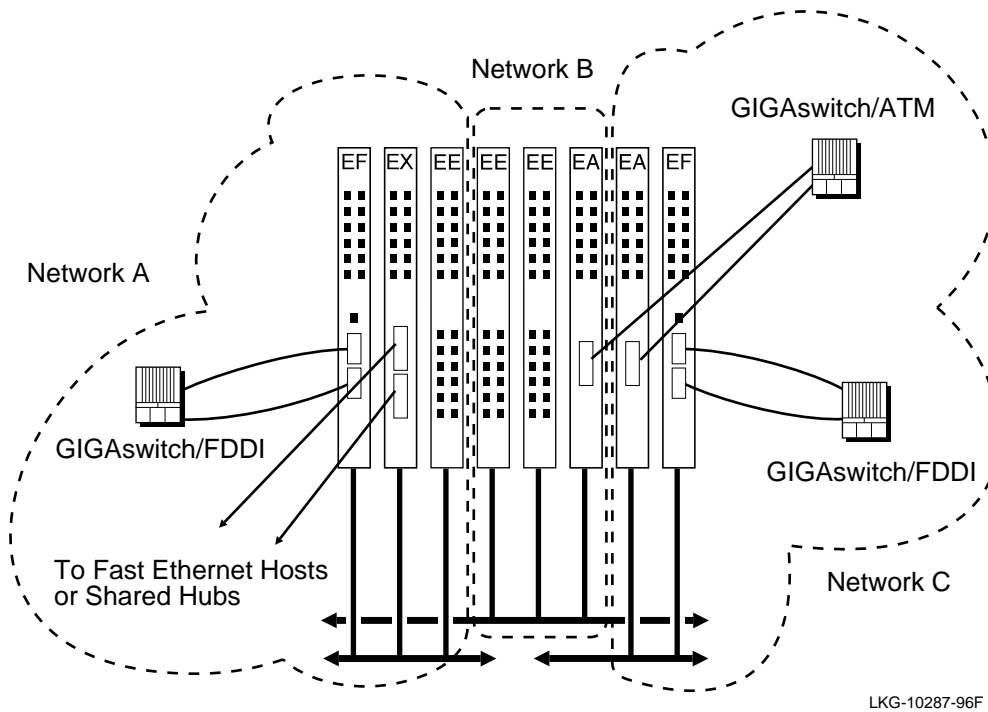
In Figure 7-5, the optional routing software provides routing between VLANs that reside on one or more modules attached to the same VNbus. In this illustration, three modules (one from each cluster, or network) are equipped with the optional routing software. More modules can be upgraded with the optional routing software if additional capacity or redundancy is needed.

Figure 7-5: Routing Between VLANs
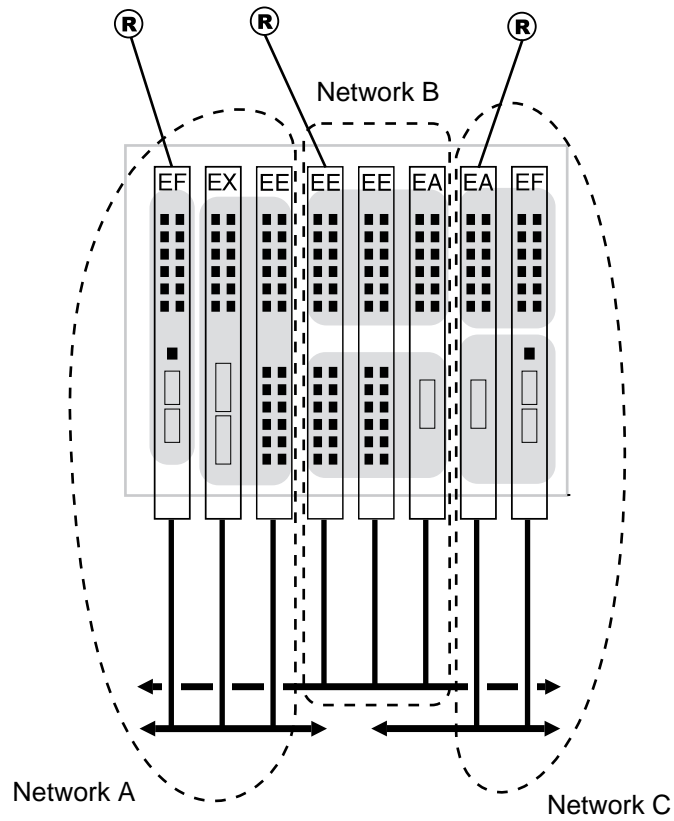
R = Optional Routing Software

LKG-10288-96F

Distributed Switched Backbone Configurations

Introduction

This section provides examples of several multi-technology distributed switched backbone configurations. The examples in this section show how the VNswitch modules can be used to join two or more DEChub 900 MultiSwitch chassis without relying on an external switch such as a GIGAswitch/FDDI or GIGAswitch/ATM. Two configurations are illustrated:

- Distributed switched backbone configuration using VNswitch 900EX modules

- Distributed switched backbone configuration using VNswitch 900EA modules

Distributed Switched Backbone Configuration Using VNswitch 900EX Modules

Figure 7-6 illustrates a distributed switched backbone configuration in which two VNswitch 900EX modules in separate hubs are connected over two full-duplex links, providing 400 Mb/s of inter-chassis capacity (100 Mb/s in each direction per link).

Figure 7-6: Using VNswitch 900EX Modules to Connect Multiple Hubs
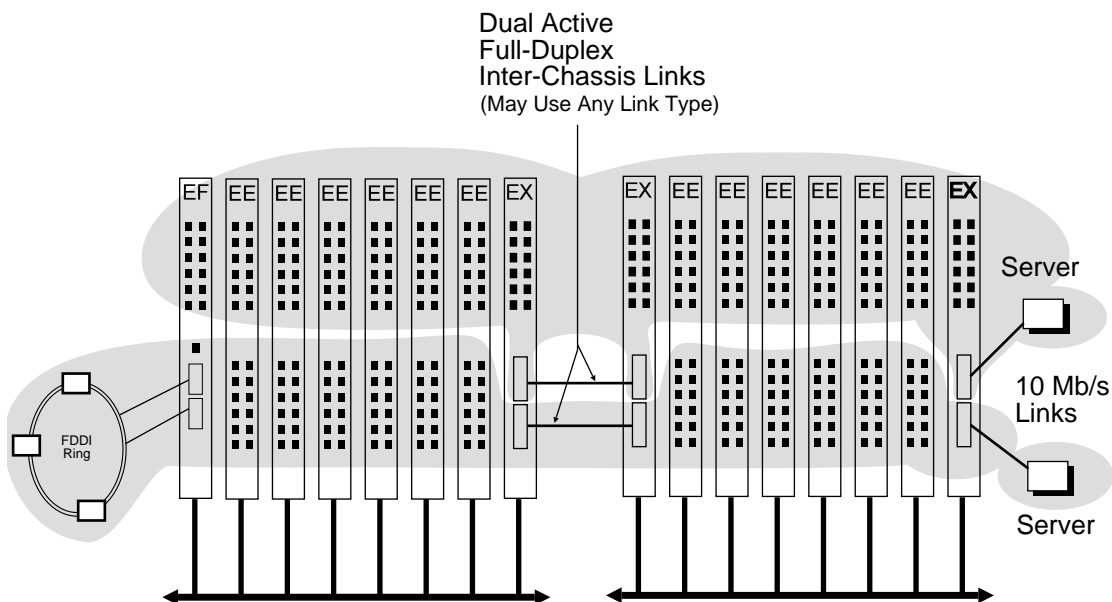
LKG-10289-96F

Although this configuration supports as many as 63 VLANs, only two VLANs are shown: all the lower ports on the modules in both hubs are configured into a single VLAN and the upper ports are configured into another VLAN (the VNswitch 900EF FDDI port-pair is configured as part of the lower VLAN).

Distributed Switched Backbone Configuration Using VNswitch 900EA Modules

Figure 7-7 illustrates a distributed switched backbone configuration in which two VNswitch 900EA modules in separate hubs are connected over a single full-duplex link, providing approximately 310 Mb/s of inter-chassis capacity (155 Mb/s in each direction per link).

Figure 7-7: Using VNswitch 900EA Modules to Connect Multiple Hubs

LKG-10291-96F
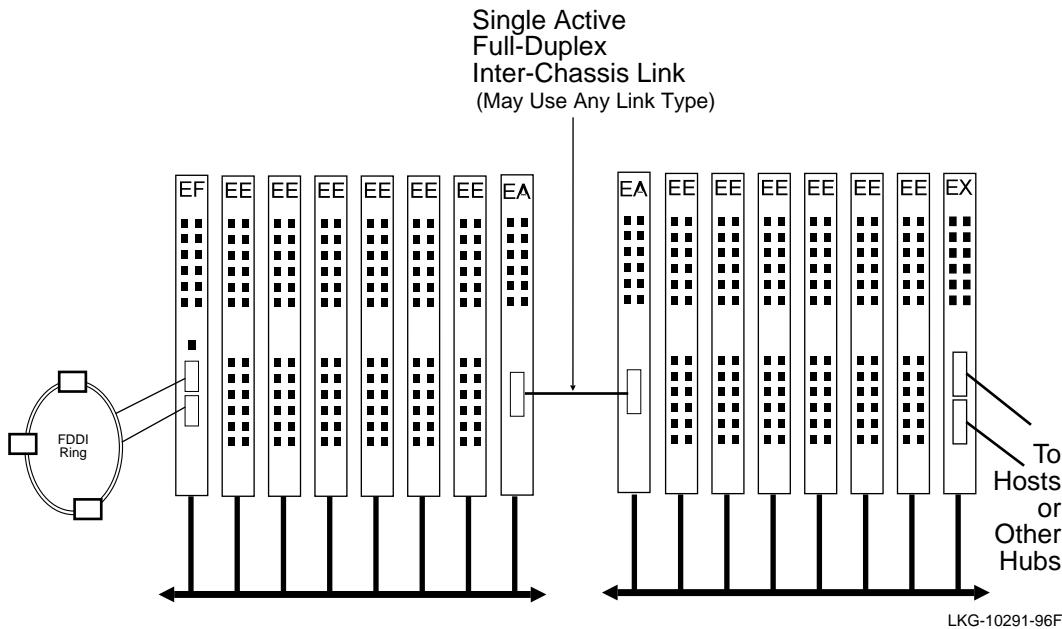
Though not shown in this figure (or in Figure 7-6), more than two parallel links can be configured. For example, two VNswitch 900EA modules could be used in each hub to provide full-duplex links for multiple VLANs, with a separate bridge tunnel configured for each VLAN. In this case (using VNswitch 900EA modules), the dual full-duplex links would provide approximately 620 Mb/s of inter-chassis capacity (155 Mb/s in each direction per link).

Multiple VLANs Across an ATM LAN/WAN Backbone

Introduction

The VNswitch modules adapt themselves to a wide range of both LAN and WAN configurations, even without the use of routers. This section provides an example of how multiple VLANs on different VNswitch modules (and other hub modules) can be transported across an ATM LAN/WAN backbone.

Configuration Example

Figure 7-8 illustrates two hubs connected to an ATM LAN/WAN backbone. In this example, one hub is configured with three VNswitch modules: an EF, EE, and EA. If the three modules are connected to the same VNbus, all the LANs attached to the modules can access the ATM port (on the EA module) to join VLANs on the other hub attached to the ATM LAN/WAN.

The other hub is configured with both VNswitch modules and non-VNswitch modules. As shown in the illustration, this hub includes a DEChub 90 module, a PortSwitch 900, and a DECswitch 900EF. Some of the ports on these non-VNswitch modules are connected to backplane LANs with VNswitch modules and other ports may reside in backplane LANs that are separate from the VNswitch modules, while others are not connected to any backplane LANs.

All the VNswitch modules in the second hub are connected to the same VNbus, allowing VLANs to span across the modules. If a port from a non-VNswitch module is connected to a backplane LAN, and the backplane LAN includes a VNswitch port that belongs to a VLAN, the non-VNswitch module port also becomes a member of the same VLAN. Thus, in this example, ports from each of the modules can be included in the same VLAN, even though the LANs attached to these ports are based on different technologies. Furthermore, by configuring an Emulated LAN (ELAN) between the two VNswitch 900EA modules across the ATM LAN/WAN backbone, VLANs on the two hubs can be joined to form a single VLAN that spans across the ATM LAN/WAN and includes LANs from all the various modules in both hubs, irrespective of the LANs' underlying technologies.

Figure 7-8: Multiple VLANs Across an ATM LAN/WAN Backbone
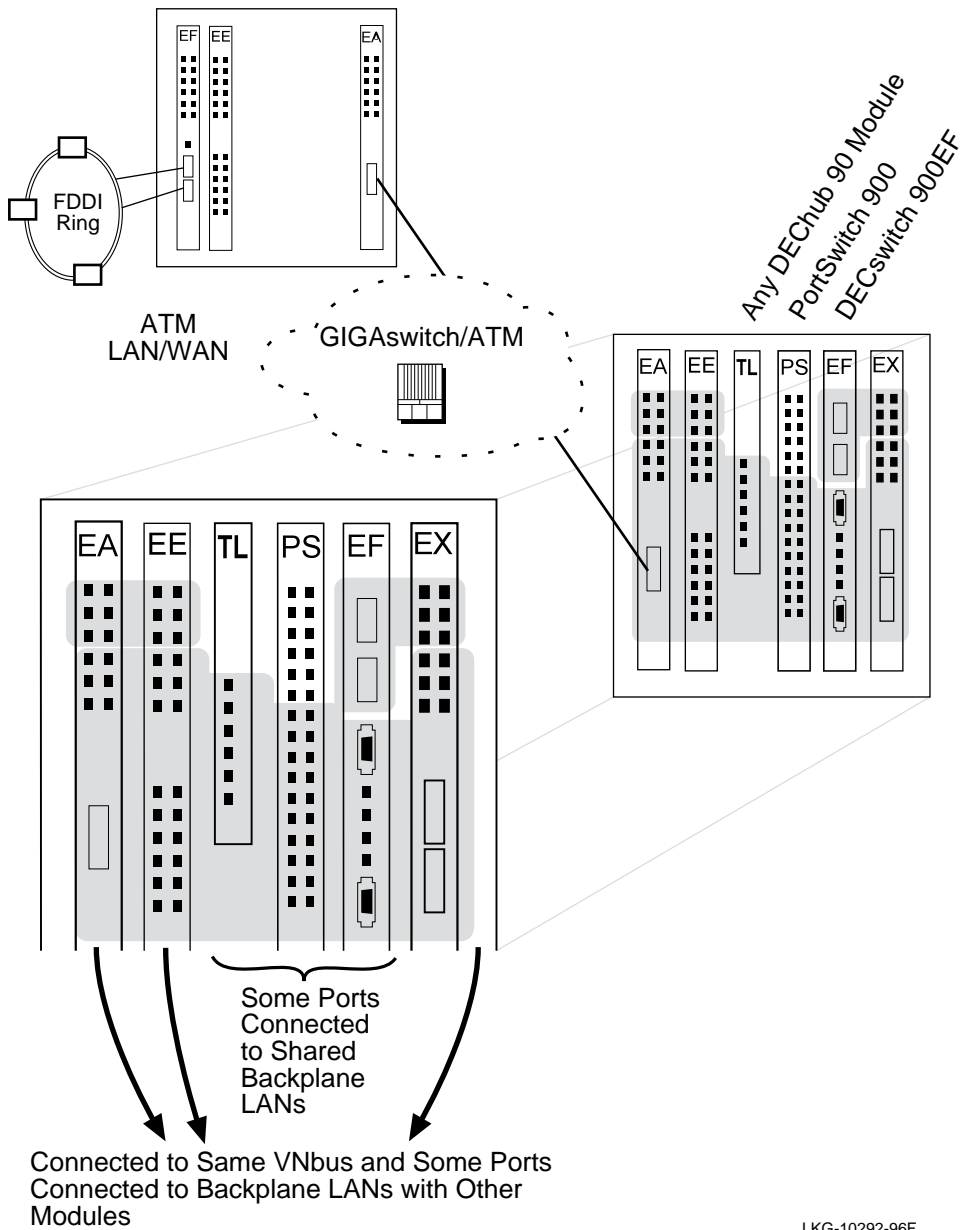
LKG-10292-96F

Summary

Virtually any configuration of multiple VLANs comprised of LANs from all four technologies (Ethernet, Fast Ethernet, FDDI, and ATM) can be designed using the VNswitch modules. By using the 16 virtual ATM ports on the VNswitch 900EA module, along with the powerful VNbus and VLAN features that are standard on all VNswitch modules, it is possible to create highly efficient and flexible network configurations that meet the demands of almost any enterprise or institution.

The VNswitch modules are designed with built-in flexibility, to be used as standalone key components in small- to medium-range networks, or as integral members of an enterprise-wide network. This evolving family of high-performance, cost-effective, network switches is set to become the standard by which all other network switching products are measured. For information about future product offerings, consult your local Digital sales representative.

Appendix A

Accessing MIBs and RFCs

Overview

Introduction

This appendix describes how to access this product's online release notes, public MIBs, Digital's private MIBs, firmware images, and requests for comments (RFCs).

In This Appendix

The following topics are covered in this appendix:

Topic	Page
Supported MIBs and RFCs	A-2
Accessing Online Information	A-3

Supported MIBs and RFCs

Following is a complete list of the MIBs and RFCs supported by the VNswitch 900 family of switches:

- SNMP Management   (RFC 1157)

- MIB-II   (RFC 1213)

- ifExtensions   (RFC 1573)

- ICOM   (DEC Proprietary)

- PCOM   (DEC Proprietary)

- LAN Emulation Client   (DEC Proprietary)

- dot1d   (RFC 1493)

- dot3   (RFC 1398)

- FDDI   (RFC 1512)

- SONET   (RFC 1595)

- AToM   (RFC 1695)

- DS3/E3   (RFC 1407)

- Ether-like   (RFC 1643)

- DEC VLAN V1   (DEC Proprietary)

- 4mic   (DEC Proprietary)

- 100baseX   (DEC Proprietary)

- ATM Bridge Tunnels   (DEC Proprietary)

- LAN Emulation Client   (DEC Proprietary)

- ILMI   (DEC Proprietary)

Accessing Online Information

MIBs, product release notes, and firmware update files can be accessed from Digital's Network Products Business (NPB) World Wide Web site. You can also access this information from the NPB Bulletin Board System (BBS).

RFCs can be accessed using electronic mail (email).

Using the World Wide Web

To access product information from Digital's NPB World Wide Web site, use one of the following Universal Resource Locators (URLs):

Americas Network Products Business Home Page	http://www.networks.digital.com/
Europe Network Products Business Home Page	http://www.networks.europe.digital.com/
Australia Network Products Business Home Page	http://www.digital.com.au/networks/
Digital Equipment Corporation Home Page	http://www.digital.com/

To get firmware and MIB information, choose the "Products and Technology" link, and from there choose the "Technical Data" link.

Using the NPB Bulletin Board System (BBS)

To access product information from the NPB Bulletin Board System (BBS), set your modem to 8 bits, no parity, 1 stop bit, and dial 508-486-5777 (U.S.A.).

Using Electronic Mail

The DDN Network Information Center (NIC) of SRI International provides automated access to NIC documents and information through electronic mail. This is especially useful for people who do not have access to the NIC from a direct Internet link, such as BITNET, CSNET, or UUCP sites.

To use the mail service, send a message to:

```
SERVICE@NIC.DDN.MIL
```

Accessing Online Information

In the SUBJECT field, request the type of service that you want, followed by any needed arguments. Normally, the message body is ignored, but if the SUBJECT field is empty, the first line of the message body is taken as the request.

The following are example SUBJECT lines to obtain DDN NIC documents:

```
HELP

RFC 822

RFC INDEX

RFC 1119.PS

FYY 1

IETF 1IETF-DESCRIPTION.TXT

INTERNET-DRAFTS 1ID-ABSTRACTS.TXT

NETINFO DOMAIN-TEMPLATE.TXT

SEND RFC: RFC-BY-AUTHOR.TXT

SEND IETF/1WG-SUMMARY.TXT

SEND INTERNET-DRAFTS/DRAFT-IETF-NETDATA-NETDATA-00.TXT

HOST DII
```

Requests are processed automatically once a day. Large files are broken into separate messages.

Appendix B

Spanning Tree

Overview

Introduction

This appendix provides detailed information about the spanning tree algorithm, a process implemented by all transparent bridging switches that conform to the IEEE 802.1d MAC Bridge specification. In this section, the terms switch and bridge imply the same meaning.

For more information about the spanning tree, see the *Bridge and Extended LAN Reference* manual.

NOTE

The terms transparent bridging and spanning tree have the same meaning and are used interchangeably.

In This Appendix

The following topics are covered in this appendix:

Spanning Tree Algorithm

Introduction

Every extended 802.3 LAN of any mesh complexity is logically configured into a network topology called a spanning tree. This is accomplished by a continuous, distributed process that is determined by the spanning tree algorithm. The algorithm ensures that the configuration contains no loops (that there is only one path between any two nodes) and that all LANs are connected.

Although the spanning tree algorithm is continuous and self-maintaining, some parameters can be set by management software. Other parameters cannot be set but are determined by the spanning tree computation process and can be displayed by management software.

This section discusses the spanning tree algorithm used by the VNswitch 900 family of switches, the properties of the algorithm, and how the algorithm computes the spanning tree.

Properties of the Spanning Tree Algorithm

The spanning tree algorithm has the following properties:

- Loop detection — If bridges (or switches) are accidentally or deliberately configured into a loop, the algorithm computes a loop-free topology that still spans the entire network.

- Automatic backup (using redundant bridges) — Bridges can be deliberately configured in a redundant path so that one of the bridges in the loop can serve as the backup for another. The process automatically configures a redundant bridge as a backup bridge. The backup bridge does not forward frames.

- Determinism — A fixed set of rules controls the process so that when variables change, the results are predictable.

- Low network overhead — The messages that control the spanning tree are usually transmitted at 1-second intervals (default), thus using a very small percentage of the available network bandwidth.

- Management — The algorithm allows tuning of parameters by management software to control the topology.

Spanning Tree Computation Process

Introduction

The spanning tree computation process is a continuous process that sets up the spanning tree when bridges are initialized and maintains the spanning tree thereafter.

Establishing the Spanning Tree

Establishing the spanning tree involves these steps:

Step	Action
1	Bridges in the extended LAN elect a unique root bridge.
2	Bridges in the extended LAN elect a designated bridge for each LAN.
3	Redundant paths are removed from the logical spanning tree.

The spanning tree is self-maintaining, and performs the following functions after it is established:

- Replaces a failed forwarding bridge with a backup bridge.

- Removes a redundant bridge when a loop is detected.

- Maintains address timers that control the aging of forwarding database address entries.

How Bridges (or Switches) Communicate with Each Other

The spanning tree algorithm is a distributed process in which all bridges in the extended LAN participate. Each bridge maintains information about itself and the spanning tree in databases: one database for a set of spanning tree parameters associated with the bridge, and another database for each bridge or switch port. These parameters are used for computing the spanning tree and for providing results of the spanning tree computation.

Bridges communicate with each other with a minimum-size packet called a Hello message, referred to in the IEEE 802.1d specification as a Configuration Bridge Protocol Data Unit (BPDU).

The Hello messages provide the following information:

Attribute	Meaning
Designated Root (Best Root)	The Designated Root parameter indicates the bridge ID that this bridge believes is the root bridge.
My Cost	The cost of the path to the root from the bridge that is transmitting the Hello messages.
Bridge ID	The unique identifier of the bridge that is transmitting the Hello messages.
Port ID	The unique identifier of the port on which the Hello message was transmitted.
Root Age (Message Age)	Indicates the age, in seconds, of the last Hello message sent by the designated root. This parameter is referred to as the Message Age in the 802.1d specification.
Actual Forward Delay	The length of time spent by the bridge in each of the Listening and Learning states before transitioning into the Forwarding state
Actual Listen Time (Max Age)	Specifies the age of a Hello message (in seconds), after which the bridge considers the message to be stale. This parameter is referred to as the Max Age in the 802.1d specification.
Hello Flags	Flags used by the designated bridge to notify other bridges of an impending topology change.

Per-Port Control of the Spanning Tree Algorithm

The VNswitch 900 series switches provide the ability to turn off the spanning tree selectively on individual ports. The switches accomplish this by using an SNMP MIB object that can be set by management software.

When the spanning tree algorithm is turned off on a port, that port no longer participates in the spanning tree computation process. No Hello messages are transmitted on the port, and any Hello messages that are received on the port are discarded. This feature helps in setting up isolated spanning tree domains (where topology changes in one domain do not affect the other domains, and connectivity betwen the domains is still maintained).

For example, to set up two isolated spanning tree domains, start with two extended LANs with no connectivity between them. Next, turn off the spanning tree algorithm on one switch port on each of the extended LANs. Establish connectivity between these two ports by connecting them, while ensuring that there are no other paths between the two extended LANs.

Refer to the *clearVISN Configuration and Use* guide for specific instructions on how to toggle on and off the spanning tree algorithm on switch ports.

Spanning Tree Parameters

Introduction

Spanning tree parameters describe and determine the shape of the spanning tree. Several spanning tree parameters can be set with management software; others cannot be set but are determined by the spanning tree computation process. All spanning tree parameters can be displayed using management software.

This section describes all of spanning tree parameters, including the global parameters (called bridge spanning tree parameters) and the per-port parameters (called port spanning tree parameters).

Bridge Spanning Tree Parameters

Bridge spanning tree parameters describe the spanning tree from the module's (bridge or switch) perspective. A database of bridge spanning tree parameters is associated with each module.

Actual Forward Delay

The Actual Forward Delay parameter indicates the Forward Delay currently in use by the root bridge. The Forward Delay for a bridge may be set with bridge management, but once the spanning tree computation process is complete, the bridge uses the root bridge's Forward Delay. If the bridge becomes the root bridge, its Forward Delay value becomes the Actual Forward Delay for all bridges in the network.

During the first half of the Forward Delay, bridges send and listen to Hello messages, participating in the spanning tree computation process. During the second half, bridges examine frames received on both ports, adding station address entries in their forwarding databases (this is the Learning state).

Actual Hello Interval

The Actual Hello Interval parameter indicates the Hello interval currently in use by the root bridge. The Hello Interval parameter for a bridge may be set with bridge management, but once the spanning tree computation process is complete, each bridge uses the root bridge's Hello Interval parameter value. If the bridge becomes the root bridge, its Hello Interval parameter value becomes the Actual Hello Interval for all bridges in the network.

Actual Listen Time

The Actual Listen Time parameter indicates the Listen Time currently in use by the root bridge. The Listen Time for a bridge may be set with bridge management but, once the spanning tree computation process is complete, the bridge uses the root bridge's Listen Time. If the bridge becomes the root bridge, its Listen Time value becomes the Actual Listen Time for all bridges in the network.

Bad Hello Limit

The Bad Hello Limit parameter specifies the number of successive Hello intervals during which a bridge may receive bad Hello messages before the bridge performs a link test on the port. A bad Hello message may indicate a port problem. The Bad Hello Limit parameter works with other spanning tree parameters.

Bad Hello Reset Interval

The Bad Hello Reset Interval parameter specifies how many Hello intervals without bad Hello messages a bridge will wait before it resets the Bad Hello Count for a port.

This parameter indicates how long a bridge will hold the current value of the Bad Hello Count, even though the bridge is not receiving bad Hello messages. The bridge automatically restarts this timer each time it receives another bad Hello message. The timer is expressed in Hello intervals.

Best Root (Designated Root)

The Best Root parameter indicates the bridge ID that this bridge believes is the root bridge. In the 802.1d specification, the name for this parameter is Designated Root. The Best Root Age parameter indicates the age, in seconds, of the most recent Hello message from the Best Root. When the value of the Best Root Age exceeds the value of the Listen Time parameter, the bridge assumes the root has expired and sends out Hello messages on all of its ports, declaring itself to be the root bridge and the designated bridge on its LANs.

Forwarding Database Normal Aging Time

The Forwarding Database Normal Aging Time parameter specifies how long a bridge retains learned station address entries in its forwarding database during normal network operation. If an address does not appear in the source field of a frame for a period of time defined by this parameter, its entry in the forwarding database is marked inactive and may be removed.

Forwarding Database Short Aging Time

The Forwarding Database Short Aging Time parameter specifies how long a bridge retains learned station address entries in its forwarding database when a topology change is in progress.

Spanning Tree Parameters

Forward Delay

The Forward Delay parameter specifies the period of time that a bridge's ports stay in the Preforwarding state before entering the Forwarding state. This value is significant only for the root bridge, because it administers the Actual Forward Delay for all bridges in the spanning tree.

Hello Interval

The Hello Interval parameter specifies the time interval between the transmission of Hello messages when this bridge attempts to become the root, or *is* the root bridge.

Inlink

The Inlink parameter indicates the port number of this bridge's port on the path to the root bridge. In the IEEE 802.1d specification, this parameter is called the Root Port.

Listen Time

The Listen Time parameter specifies the age of a Hello message (in seconds), after which the bridge considers the message to be stale. This value is significant only for the root bridge, since it administers Actual Listen Time for all bridges in the spanning tree. In the IEEE 802.1d specification, the name for this parameter is Max Age.

My Cost

The My Cost parameter indicates the bridge's current path cost to the root bridge.

No Frame Interval

The No Frame Interval parameter specifies the number of seconds that a bridge waits without receiving a frame on a port before the bridge suspects a problem and runs a link test on the port.

Root Priority

The Root Priority parameter is the most significant byte of the bridge ID. It can be used to establish the root bridge or designated bridge.

During the spanning tree computation process, Hello messages from all bridges in the network are compared so that the root bridge and designated bridges can be determined. The bridge with the lowest bridge ID becomes the root bridge, with Root Priority values compared first and hardware addresses second.

Tell Parent Flag

The Tell Parent Flag parameter indicates that the bridge needs to send a Topology Change Notification on its inlink to its parent bridge, the next closest bridge in the path to the root.

Topology Change Flag

The Topology Change Flag parameter indicates that the root bridge has been notified of a topology change in the network and that bridges are to use the Forwarding Database Short Aging Time.

Topology Change Timer

The Topology Change Timer indicates the number of seconds that the root bridge sends Hello messages with the Topology Change Flag set. The value of this parameter is equal to the duration one half the Forward Delay parameter value plus the Listen Time parameter value.

Port Spanning Tree Parameters

Port spanning tree parameters consist of one settable parameter and several nonsettable parameters. Port spanning tree parameters describe the spanning tree from each port's perspective. The Line Cost parameter can be modified with management software; all others are non-settable but can be displayed.

A database of port spanning tree parameters is associated with each port.

Acknowledgment Flag

The Acknowledgment Flag indicates that the port has received the Topology Change Notification from a bridge lower in the spanning tree (farther from the root bridge). The acknowledgment is sent in the next Hello message. In the IEEE 802.1d specification, this parameter is called the Topology Change Detected flag.

Bad Hello Count

The Bad Hello Count parameter indicates the number of consecutive Hello intervals during which the bridge received a bad Hello message on a port. When the value of the Bad Hello Count reaches the Bad Hello Limit set for the bridge, the bridge resets this counter, increases the Bad Hello Limit Exceeded Count by one, and performs a link test on the port.

Note that if the Clear Time Count parameter value reaches the value of the Bad Hello Reset Interval bridge parameter before the Bad Hello Count reaches the Bad Hello Limit, the port resets the Bad Hello Count.

Bad Hello Limit Exceeded Count

The Bad Hello Limit Exceeded Count indicates the number of times that this bridge's Bad Hello Limit has been exceeded since its initialization.

Spanning Tree Parameters

Clear Time Count

The Clear Time Count parameter indicates the number of consecutive Hello intervals during which the bridge has received no bad Hello messages on this port. When the Clear Time Count reaches the Bad Hello Reset Interval bridge parameter, the bridge resets the Clear Time Count and the Bad Hello Count.

Designated Bridge ID

The Designated bridge ID parameter indicates the bridge ID of the designated bridge on this LAN (the LAN connected to this port).

Designated Bridge Link Number

The Designated Bridge Link Number parameter indicates the port number of the designated bridge on this LAN (the LAN connected to this port). This parameter is referred to as the Designated Port in the 802.1d specification.

Designated Root Age

The Designated Root Age parameter indicates the age, in seconds, of the last Hello message sent by the designated root.

Designated Root ID

The Designated Root ID parameter indicates the bridge ID of the bridge considered to be the root bridge on this port.

Forward Delay Timer

The Forward Delay Timer parameter indicates the time remaining before the port will leave the PREFORWARDING state and enter the FORWARDING state.

Line Cost

The Line Cost parameter specifies the cost value for the port, which is used to determine the path cost to the root bridge. The Line Cost parameter may be set by network management. This parameter is referred to as the Path Cost in the 802.1d specification.

Port Address

The Port Address parameter indicates the hardware address of the port, which may differ from the bridge address.

Possible Loop Flag

The Possible Loop Flag parameter indicates whether the bridge has detected a loop condition in a situation where the Bad Hello Count is not zero.

Root Path Cost

The Root Path Cost parameter indicates the cost to the root bridge for the designated bridge on this LAN. This parameter is referred to as the Designated Cost in the 802.1d specification.

The spanning tree protocol specifies that all participating bridges in the network exchange HELLO Bridge Protocol Data Units (BPDUs) which provide configuration information about each bridge. BPDUs include information such as the bridge ID, root ID, and root path cost. This information helps the bridges to determine unanimously which bridge is the root bridge and which bridges are the designated bridges for LANs to which they are connected.

Glossary

This glossary is a comprehensive source of definitions of terms associated with the VNswitch 900 series switches.

A

address

The logical location of a terminal, node, peripheral device, or byte in memory in a network.

address filtering

Filtering of packets by switches based on source or destination addresses of the packets.

address forwarding database

See **forwarding database**.

Aging Time

A spanning tree parameter that controls how long a bridge keeps each learned entry in the forwarding database. If an entry is stored longer than the aging time, the bridge marks that entry as inactive and allows it to be overwritten.

algorithm

A computational function that determines how values for a particular object are obtained.

ANSI

American National Standards Institute. An organization that coordinates and publishes standards for use in the United States.

ARP

Address Resolution Protocol. An Internet protocol that dynamically binds a high-level internet address to a low-level physical hardware address.

asynchronous transmission

Data transmission in which each information character, or sometimes each word or small block, is individually synchronized, usually with start or stop elements (for

example, bits). The gap between each character or word is not a fixed length.

ATM

Asynchronous Transfer Mode. A high-speed networking technology that switches
network traffic between nodes on dedicated point-to-point network paths.

attenuation

The amount of optical power (or light) that is lost as the light travels from the
transmitter through the medium to the receiver. The difference between transmitted
and received power, expressed in decibels (dB).

AUI

Attachment unit interface. A connector for the Ethernet interface that attaches a
workstation to a transceiver.

B

backbone

A core network (usually high-speed) to which multiple local area networks (LANs) are
often connected by means of switches and bridges, and over which traffic can pass.

backplane LAN

An internal LAN segment configured from one or more flexible channels on the
DEChub 900 backplane.

bandwidth

A measure of the amount of traffic the media can handle at one time. In digital
communications, bandwidth describes the amount of data, in bits per second, that can
be transmitted over the line.

bits/s

Bits per second. The rate at which data is transmitted over a communications channel.

BootP

A protocol that is used by a network node to determine the Internet Protocol (IP)
address of its Ethernet interfaces used for network booting.

bridge (or MAC bridge)

An intelligent, protocol-independent, store-and-forward device that operates as a data
link layer relay. Used to connect similar or dissimilar local area networks (LANs). A
collection of LANs connected by bridges is referred to as an extended LAN.

bridge address

Used by the spanning tree algorithm, the least significant 6-octet part of the bridge
identifier.

bridge identifier

Used in a spanning tree network to uniquely identify each bridge on the network.

bridge number

The number that identifies the specific bridge to a segment or ring.

bridge port

An abstract connection point (within the VNswitch module's software) to a transparent bridge. The VNswitch module (functioning as a transparent bridge) forwards data to, or receives data from, bridge ports, based on the MAC address associated with the data.

bridge priority

The most significant 2-octet part of the bridge identifier that is used to determine which bridge on the network becomes the root bridge.

broadcast network

A network that transmits a packet of data that any computer on the network can pick up and use.

C

cell switch

A switch that processes and forwards fixed-length cells.

clearVISN

An SNMP-based graphical user interface (GUI) management tool that provides access to various statistics associated with the VNswitch modules and other network modules.

CLI

Command Line Interface. A command line driven management tool for managing a VNswitch module.

coaxial cable

A cable consisting of a central conductor surrounded by an insulator and then by another grounded conductor. The outer cable acts as a shield to prevent interference from reaching the inner conductor.

collision

An event that occurs when two computers attached to a network attempt to send a packet simultaneously.

collision domain

A shared physical medium with which systems must contend for access.

community

A set of attributes that are managed as a group. Community names are used to control access. Each software request contains a community name that the agent uses as a password to verify that the requester is authorized to access the agent's management information base (MIB) or a subset of that MIB.

community strings

A text string, also referred to as a community name, that is used as a password by an SNMP agent to authenticate or verify that the network management station (NMS) making the request is allowed access to the contents of the agent's Management Information Base (MIB).

concentrator

An FDDI physical layer repeater that allows the attachment of multiple single attachment stations, dual attachment stations, or other concentrators to the FDDI network.

configuration management

The portion of connection management that provides for configuration of PHY and MAC entities within a station. It includes the physical connection management, configuration management, and entity coordination management components.

connection

The path between two protocol modules that provides reliable stream delivery service.

cross-connect

Patch cable and passive hardware that is used to administer the connection of cables at a central or remote location.

CSMA/CD

Carrier Sensed Multiple Access/Collision Detection. A method of transmitting information in the LAN environment where only one transmitter is on the line at any one time. If two devices transmit simultaneously, the signals collide and transmissions temporarily cease. The Ethernet IEEE-802.3 standard uses CSMA/CD technology.

cyclic redundancy check (CRC)

A method of detecting errors in a frame by performing a mathematical calculation of the number of bits in the frame and appending the result to the end of the frame. The receiving station performs the same calculation on the frame and then checks that the result matches the CRC at the end of the frame.

D

DAC

Dual attachment concentrator. A concentrator that offers two connections to the FDDI network capable of accommodating the FDDI dual (counter-rotating) ring, and additional ports for the connection of other concentrators or FDDI stations.

DAS

Dual attachment station. An FDDI station that offers two connections to the FDDI dual counter-rotating ring.

designated bridge

The bridge that claims to be the closest to the root bridge in the accumulated path cost.

designate port

The port ID of the designated bridge attached to the LAN.

destination

A node designated as the intended receiver of data.

destination address (DA)

The field in a frame that contains the address of the station to which the frame is being sent.

Destination Service Access Point (DSAP)

The destination and source address fields of an LLC PDU.

distance-vector algorithm

A class of routing algorithms that emphasizes the number of hops to find the shortest path to a destination.

DMA

Direct memory access. A method to transfer data between peripheral devices and internal memory without intervention by the central processing unit (CPU).

DRAM

Dynamic random access memory. A type of storage the computer accesses at frequent intervals.

dual homing

A method of cabling concentrators and stations in a tree configuration that permits an alternate or backup path to the FDDI network in case the primary connection path fails. Can be used in a tree or in a dual ring of trees configuration.

dual ring of trees

A topology of concentrators and stations that cascade from concentrators on a dual ring.

dynamic routing

Routing that adjusts automatically to network topology or traffic changes.

E

EIA

Electronic Industries Association. A standards organization specializing in the electrical and functional characteristics of interface equipment.

encapsulation

A method by which a protocol appends required information for a peer protocol.

encoding

The act of changing data into a series of electrical or optical signals that can travel efficiently over a medium.

end-to-end

Services referred to by the lower four layers of the OSI reference model. A packet forwarded from one host to another host over the network uses end-to-end services.

entity name

In AppleTalk, an entity name is an 8-bit ASCII character string that has three fields: object, type, and zone. Each of these fields is a string of not more than 32 characters.

entity

An active element within an OSI layer or sublayer.

ES

End system. A host system in the OSI protocol that performs the functions of all of the layers of the OSI reference model.

Ethernet

A baseband LAN technology that uses the physical and data link layers of the OSI model. Software protocols, such as TCP/IP, provide network layer functions. Ethernet includes three standards: IEEE-802.3, Version 2.0, and Version 1.0.

Ethernet station

An addressable node on an Ethernet network capable of transmitting and receiving data. An Ethernet station must contain a MAC entity.

event

A network message that indicates some irregularity in the physical and software elements of a network. An event may be informational or it may require the user to perform a specific task.

F

Fast Ethernet

A LAN technology that allows standard Ethernet frames to be transmitted at 100 Mb/s. The Fast Ethernet technology implemented by the VNswitch modules is specified in the IEEE 802.3u standard.

FDDI

Fiber Distributed Data Interface. A set of ANSI/ISO standards for a 100 Mb/s token-passing ring that uses multimode fiber, single-mode fiber, unshielded twisted-pair, screened/shielded twisted-pair, or any combination of the four as the transmission medium.

FDDI connector type

A variable that specifies the type of port being managed. The four different types of ports are:

- A — The port in a dual attachment station or concentrator that attaches to the primary in and the secondary out when attaching to the dual ring.

- B — The port in a dual attachment station or concentrator that attaches to the secondary in and the primary out when attaching to the dual ring.

- S — Either the port in a single attachment station, or one port in a single attachment concentrator.

- M — A port in a concentrator that serves as a master to a connected station or concentrator.

FDDI station

A node on an FDDI ring capable of transmitting, receiving, and repeating data. A station has one instance of SMT, at least one instance of PHY and PMD, and an optional MAC entity.

fiber optics

A transmission medium used to transmit data in the form of optical signals.

filtering

Preventing (on a per-port basis) a bridge from forwarding certain frames. Filtering can be specified by port, address, or protocol.

flexible channels

High-bandwidth channels on the DEChub 900 MultiSwitch that provide for the interconnection of multiple DEChub 900 modules.

flood

Transmitting a packet across each interface on a bridge or router.

forwarding database

A table maintained by a bridge that contains station addresses, the port on which the addresses are located, and the age of these entries. A bridge forwards or filters frames based on the information in the forwarding database. A bridge creates its forwarding database by learning the source port and source address of each incoming frame.

fragment

The process of identifying an IP or ISO CLNP datagram into smaller pieces for transmission across a network that cannot handle the original datagram size.

frame

Informal name for a data-link PDU. Control information in the frame provides addressing, sequencing, and flow and error control to the respective protocol levels.

frame switch

A switch that processes and forwards variable-length LAN packets (or frames).

FTP

File Transfer Protocol. An IP application protocol that provides reliable file transfers from one network device to another.

G

Graphical User Interface (GUI)

A display format that enables the user to choose commands, start programs, and see lists of files and other options by pointing to icons and menu items on the screen. Choices can be activated either with the keyboard or with a mouse.

H

hello interval

A spanning tree parameter that controls how often a bridge sends a hello message.

hello messages

A hello message determines which bridges are elected as designated bridges and which bridge becomes the root bridge. When the spanning tree computation is complete, the root bridge originates the hello message and the other bridges propagate it down the spanning tree.

hop

The next route a packet must travel to arrive at its destination. A hop is represented by an address or a decimal character (how many).

host

A PC, workstation, or mainframe connected to a network.

hot swap

To 'hot swap' a component means to remove and replace it while the system is in operation.

hub

A central device, usually in a star topology local area network (LAN), to which each network module is attached.

I

IEEE

Institute of Electrical and Electronics Engineers.

in-band management

Managing a device over a network.

interconnect

A panel-mounted fiber-optic coupler or wallbox-mounted fiber-optic coupler used to join two cables with a single pair of connectors.

internet

A collection of packet-switching networks interconnected by gateways with protocols. This enables networks to function as a single, cooperative virtual network.

Internet Protocol (IP)

The Department of Defense (DoD) Internet standard protocol that defines the Internet datagram as the unit of information passed across the Internet. IP corresponds to the OSI reference model layer 3 and provides connectionless datagram service.

Internet Protocol (IP) address

The IP address is a series of numbers that identifies a device's location on the Internet.

Internet Protocol (IP) name

A unique alphanumeric string that identifies a device on the Internet.

interoperability

The ability of all system elements to exchange information between single vendor and multimedia equipment. Also called *open communications*.

IP

See **Internet Protocol**

L

LAN

Local area network. A data communications network that spans a limited geographical area. The network provides high-bandwidth communication over coaxial cable, twisted-pair, fiber, or microwave media and is usually owned by the user.

LAN segment

A portion of a LAN. A LAN segment may be organized in a bus, a ring, or a point-to-point configuration, or it may be formed from a virtual LAN (VLAN) or a backplane LAN.

LED

See **light-emitting diode**.

legacy LAN

A local area network that predates Fast Ethernet and ATM. Legacy LANs include standard Ethernet and FDDI local area networks.

light-emitting diode (LED)

A semiconductor light source used as an indicator of status on a network module (or other device).

listening

A bridge state, in which a bridge identifies the addresses of messages received from the network.

LLC

Logical Link Control. IEEE defined sublayer of the OSI link layer. The LLC handles error control, flow control, and framing.

local session

A session established by connecting a terminal (for example, a workstation or PC running terminal emulation software) directly to the setup port of a DEChub 900 MultiSwitch or a DEChub ONE docking station. Local sessions can be used only to

configure and manage the modules installed in the hub to which the PC or workstation is attached. Modules installed in a remote hub can be configured and managed via a remote session.

local topology

See **logical ring.**

logical interface

An abstract connection point (within the VNswitch module's software) between a physical interface and a bridge port. Ethernet, Fast Ethernet, and FDDI physical interfaces are each associated with one logical interface. An ATM physical interface is associated with 1 to 16 logical interfaces, each of which is the connection point to either an ATM emulated LAN (ELAN) or an ATM bridge tunnel.

logical ring

The path a token follows in an FDDI network made up of all the connected MAC sublayers. The accompanying physical topology can be a dual ring of trees, a tree, or a ring.

loopback

Directing signals back toward a source along a communications path.

M

MAC

Medium access control. A medium-specific access control protocol within IEEE-802. MAC provides a set of services to ensure proper operation of the LAN, including detection of, and recovery from, error conditions.

MAC address

A unique 48-bit binary number (usually represented as a 12-digit hexadecimal number) encoded in a device's circuitry to identify it on a local area network.

MAC layer

The lower portion of the datalink layer. In FDDI (and other types of shared-medium LANs) the data link layer sublayer responsible for scheduling, transmitting, and receiving data. The MAC differs for various physical media.

management agent

See **SNMP agent**.

manual mode

The condition in which a bridge does not learn addresses but only uses addresses you specify, because you have created a filter for all unspecified addresses.

MAU (Ethernet)

When used in the context of network modules such as switches, this acronym stands for media access unit.

MIB

Management information base. A dynamic, virtual collection of data about a managed object. The managed object provides this data to the network management station (NMS), which gathers the values from the managed object and loads them into the MIB representing the object.

MIC

Media interface connector. An optical fiber connector pair that links the fiber media to the FDDI station or another cable. The MIC consists of two halves. The MIC plug terminates an optical fiber cable. The MIC receptacle is attached to an FDDI station.

ModPMD

Module Physical Media Dependent.

module

An interchangeable unit that provides functionality and interconnection to a local area network (LAN) through the hub.

See also **network modules**.

multicast

A technique that allows copies of a single packet to be broadcast to a specified number of hosts.

multicast address

A type of network addressing that enables a node to send messages or data packets to an address that represents a group of stations rather than a single station.

multilayer switching

The ability to perform switching at both the Layer 2 and Layer 3 OSI layers. Typical switches perform switching only at Layer 2.

multimode

A type of fiber optic cable in which more than one transmission mode is supported.

multiswitch backplane

A backplane that allows flexible allocation of its signals so that multiple local area network (LAN) segments can be managed. The DEChub 900 MultiSwitch has a multiswitch backplane.

N

network

A collection of computers, terminals, and other devices together with the hardware and software that enables them to exchange data and share resources over either short or long distances.

network layer

Layer 3 of the OSI reference model. This layer is where all routing occurs.

network modules

Any modular device that provides network connectivity or services and can be installed in a DEChub backplane or used standalone. These include, but are not limited to, repeaters, concentrators, bridges, brouters, terminal servers, switches, and SNMP agents.

node

Any intelligent device that communicates with other devices in the network. A node is often referred to as a station.

NSAP

Network Service Access Point. The point where the communications capability of the network layer is made available at the layer boundary to its users.   An OSI network address.

O

optical bypass relay (OBR)

Used to maintain connectivity in the absence of power or during fault conditions in a station. OBR allows light to bypass the optical receiver in the faulty station.

OSI

Open Systems Interconnection. The ISO architecture for internetworking.

OSI reference model

The seven layer model specified by ISO, which specifies particular network functions.

out-of-band management (OBM)

Management of a network module (such as the VNswitch) or the DEChub 900 MultiSwitch backplane over a telephone line or direct line to a console port rather than over the network.

P

packet

A self-contained block of information containing control and user information that is transmitted across a network

packet switching

A data transfer scheme in which information is broken into individual packets, transferred across a communications link, and reassembled at the receiving end. In a packet-switching system, the route between the sender and receiver is determined by each node through which the packet travels.

PDU

See **Protocol Data Unit**.

personal Ethernet

An individual LAN segment configured for a single station.

PHY

Physical Layer Protocol. A Fiber Distributed Data Interface (FDDI) standard that defines symbols, line states, clocking requirements, and the encoding of data for transmission.

physical address

The address of the interface between the MAC interface and a LAN.

physical interface

The physical connection point on the VNswitch module to which a network transmission medium (cable or fiber, for example) is connected. Physical interfaces on VNswitch modules include Ethernet, Fast Ethernet, FDDI, ATM, and the VNbus.

physical topology

The arrangement of cables and hardware that make up the network.

PMD

Physical Media Dependent. A function of the FDDI interface card that provides the power levels and characteristics of the optical transmitter and receiver, optical signal requirements, and bit error rates.

port connector type

A variable that specifies the type of port being managed. The four different types of ports are:

- A — The port is a dual attachment station or concentrator that attaches to the primary in and the secondary out when attaching to the dual ring.

- B — The port is a dual attachment station or concentrator that attaches to the secondary in and the primary out when attaching to the dual ring.

- S — Either the port in a single attachment station, or one port in a single attachment concentrator.

- M — A port in a concentrator that serves as a master to a connected station or concentrator.

port priority

In the spanning tree algorithm, the port priority is the second 1-octet part of the port ID.

PPP

Point-to-Point Protocol. A channel or link with only two terminals whose purpose is to transmit protocol datagrams at the data link layer over serial point-to-point links.

protocol

A formal set of rules governing the format, timing, sequencing, and error control of exchanged messages on a data network.

A protocol can also include facilities for managing a communications link or contention resolution.

A protocol can relate to data transfer over an interface, between two logical units directly connected, or on an end-to-end basis between two end users over a large and complex network. There are hardware protocols and software protocols.

Protocol Data Unit (PDU)

A data object exchanged by protocol layers that contains both protocol control information and user data.

PVC

Permanent virtual circuit. A Frame Relay and X.25 feature in which data traveling between end points uses a preestablished path. A PVC gives the appearance of a permanent point-to-point connection.

R

rate limit

The total number of frames on all ports allowed by a bridge per second.

remote session

A communications session established by running Telnet on a remote system and entering the VNswitch module's IP address or the IP address of the module's out-of-band management (OBM) port. Remote sessions are established to remotely configure and manage hub modules.

ring

Connection of two or more stations in a circular logical topology. Information is passed sequentially between active stations, where each one, in turn, examines or copies the data, and returns it to the originating station, which removes the data from the network.

FDDI is an example of an implementation of a ring topology.

RISC

Reduced Instruction Set Computer. A type of processor architecture that minimizes the number of instructions performed by the processor to increase processing speed.

root bridge

The bridge with the highest priority bridge ID that is selected as the root of the spanning tree. In an extended local area network (LAN), the root bridge controls the spanning tree configuration by originating hello messages. The root bridge is responsible for keeping the spanning tree intact.

The spanning tree algorithm determines the root bridge by comparing the bridge IDs for all bridges in the extended LAN and by selecting the bridge with the lowest bridge ID (root priority and hardware MAC address).

root priority

A spanning tree parameter that determines a bridge's priority for becoming the root of the logical spanning tree. The Root Priority parameter value is used as a prefix to the bridge's address to form the bridge's identification, for example, 128/08-00-2B-2C-08-21.

route

A path through a series of LANs and bridges.

router

A device with the ability to route packets from one end station to another with multiple paths between them.

routing domain

In OSI, a set of ESs and ISs that share routing information, operate according to the same routing protocol, and are contained within a single administrative domain.

S

SAC

Single attachment concentrator. A concentrator that offers one S port for attachment to the FDDI network, and M ports for the attachment of stations or other concentrators.

SAS

Single attachment station. An FDDI station that offers one S port for attachment to the FDDI ring.

server

A node or host that provides services to a client.

shielded twisted-pair

A twisted-pair cable whose resistance to electromagnetic interference is improved by use of a flexible metallic sheath surrounding the pairs.

single mode

A type of fiber optic cable in which only one transmission mode is supported.

SLIP

Serial Line Internet Protocol (SLIP). A protocol used for transmitting Internet Protocol (IP) packets across serial lines.

SMT

See **station management**.

SNMP

Simple Network Management Protocol (SNMP). A high-level, standards-based protocol for network management, usually used in TCP/IP networks.

SNMP agent

An entity in a hardware device that executes SNMP agents.

SNMP MIB-II

The Management Information Base (MIB-II) for use with network management protocols in TCP/IP-based internets.

source address (SA)

The address of the station that originated the data transmitted on a network.

source routing

A bridging mechanism that routes frames through a multi-LAN network by specifying in the frame which route it travels.

spanning tree

A method of creating a loop-free, logical topology on any mesh topology in an extended local area network. Formation of a spanning tree topology for transmission of messages across bridges is based on the industry-standard spanning tree algorithm defined in IEEE 802.1d.

SRAM

Static Random Access Memory. A type of random-access memory that holds its contents without constant refreshing from the CPU.

standalone

A network module in a single configuration.

standby port

A redundant port that is configured as a backup port to an active port. A standby port is disabled and is held in readiness in case the active port fails.

static route

A route that is manually entered into the routing table.

station (FDDI)

A node on an FDDI ring capable of transmitting, receiving, and repeating data. A station has one SMT, at least one instance of PHY and PMD, and an optional MAC entity.

station management (FDDI)

The entity within a station on the FDDI ring that monitors and exercises overall control of station activity.

STP

See **shielded twisted-pair**.

subnet

In IP, a distinct network within a network. In OSI, subnet is the connection from the IS to the subnetwork.

subnet address

An extension of the IP addressing scheme that allows a site to use a single IP address for multiple physical networks.

subnet mask

A 32-bit address mask that is used to specify a particular subnet.

subnetwork

In OSI, a collection of ESs and ISs under the control of a single administrative domain and using a single network access protocol. In IP, the sharing of a particular subnet address.

switch

A network device that switches data traffic between two or more networks or network segments.

switch console session

A local or remote session established to configure, monitor, or manage a VNswitch module.

synchronous transmission

A form of data transmission in which data is sent continuously against a precise time base that is shared by transmitting and receiving terminals.

T

T1

A long-haul transmission medium capable of transmitting information at 1.544 Mb/s.

TCP

Transmission Control Protocol (TCP). The transport protocol offering a connection-oriented transport service in the Internet suite of protocols.

TCP/IP

Transmission Control Protocol/Internet Protocol. *See* **TCP** *and* **Internet Protocol (IP)**.

TFTP

Trivial File Transfer Protocol. A simplified version of FTP.

topology

The logical or physical arrangement of nodes on a network.

transparent bridging

A type of bridging mechanism that is invisible to each end station. The VNswitch implements transparent bridging.

traps

Messages generated in SNMP agents. The firmware monitors the device for faults and sends messages to monitoring software.

For the clearVISN application, the Alarms Poller software communicates with the SNMP agents in a hub or a community and arranges for specific traps to be sent to the

network management station running the clearVISN application.

The trap table in the Alarms Definition file determines which traps to monitor. Cold start, warm start, and authentication failure are examples of traps monitored by the alarms software.

The Alarms Poller software uses Internet Protocol (IP) addresses and community names to communicate with specific SNMP agents, such as the DECagent 90, the MultiChassis Manager, and modules with built-in SNMP agents (such as the VN-switch 900 series switches).

twisted-pair

A communications cable constructed of two helically wrapped conductors.

See also **shielded twisted-pair**.

U

unshielded twisted-pair (UTP)

A cable constructed of two helically wrapped conductors that are called a pair. In the context of FDDI, it usually refers to 100-ohm cables in a 4-pair bundle.

upstream neighbor address (UNA)

An FDDI MAC address that identifies the most recently known upstream neighbor. Stations determine the upstream neighbor by exchanging neighborhood information frames (NIFs) as part of the neighbor notification protocol. Stations also use the protocol to determine the existence of duplicate address conditions.

User Datagram Protocol (UDP)

The transport protocol offering a connectionless-mode transport service in the Internet suite of protocols.

V

virtual switched network

A network that implements switching between multiple LANs configured by means of software into logical network segments (called *virtual LANs*).

VLAN

Virtual LAN. A group of bridge ports logically linked to define a LAN on one or more hubs. This network configuration scheme enables you to configure a set of devices so that they logically appear to be on the same LAN segment, although they may be physically located on different segments.

VLAN logical interface

An abstract connection point between a VLAN and a router, providing for the connection to multiple VLANs through the router.

VNbus

Virtual network bus. A high-speed bus implemented on the DEChub 900 MultiSwitch backplane that provides for the interconnection of VNswitch modules without disabling the modules' front panel ports.

VSD

VLAN Secure Domain. A logical set of one or more VLANs that operates with a single spanning tree. VNswitch modules currently support one VLAN per VSD, but the VSD concept provides for expanded support of multiple VLANs within a single VSD.

W

wide area network (WAN)

A network spanning a large geographical area that provides communications among devices on a regional, national, or international basis.

window

A portion of the screen used for displaying information.

workgroup

A network configuration characterized by a small number of attached devices spread over a limited geographical area.

Index