

LAN Bridge 100

Technical Manual

Order No. EK-DEBET-TM-003

This manual describes hardware logic, diagnostic firmware and software, and general operating procedures. The manual is intended for use in training, in field service, and in manufacturing. The level of technical information assumes previous training or experience with Ethernet networks.

digitalTM

1st Edition, August 1986
2nd Edition, July 1987
3rd Edition, March 1989

© Digital Equipment Corporation 1986, 1987, 1989
All Rights Reserved

The information in this document is subject to change without notice and should not be construed as a commitment by Digital Equipment Corporation. Digital Equipment Corporation assumes no responsibility for any errors that may appear in this document.

Printed in U.S.A.

This manual was produced originally by Networks and Communications Publications and revised by Educational Services Development and Publishing in Merrimack, N.H.

The following are trademarks of Digital Equipment Corporation:

digital [™]	DECUS	RSTS
DEC	DECwriter	RSX
DECmate	DIBOL	Scholar
DECset	MASSBUS	ULTRIX
DECspell	PDP	UNIBUS
DECsystem-10	P/OS	VAX
DECSYSTEM-20	Professional	VMS
	Rainbow	VT
		Work Processor

Amphenol 906 is a trademark of Amphenol.
AT&T is a trademark of American Telephone and Telegraph Co.
Corning is a trademark of Corning Glass Works.
FOTEC T302D is a trademark of FOTEC, Inc.
Motorola 68000 is a trademark of Motorola, Inc.
Xerox and XNS are registered trademarks of Xerox Corporation

Contents

Preface

1 Introduction

1.1	Introduction to the LAN Bridge 100 Unit	1-1
1.1.1	Extended Networks	1-2
1.1.2	Transparent Operation	1-3
1.1.3	LAN Bridge 100 to Network Connections	1-4
1.2	LAN Bridge 100 Functional Description	1-4
1.2.1	Product Versions and Designations	1-5
1.2.2	LAN Bridge 100 Packaging	1-7
1.2.3	LAN Bridge 100 Controls, Status LEDs, and Connectors	1-8
1.2.4	LAN Traffic Monitor (LTM) Option	1-14
1.2.5	LAN Traffic Monitor Software	1-15
1.2.6	Remote Bridge Management Software (RBMS)	1-15
1.3	Configuration Considerations	1-16
1.3.1	Performance Considerations	1-18
1.3.2	Loop Considerations	1-18
1.3.3	Local LAN Bridge 100 Considerations	1-18
1.3.4	Remote LAN Bridge 100 Considerations	1-20
1.3.4.1	Fiber-Optic Cable Between LAN Bridge 100 Units	1-21
1.3.4.2	LAN Bridge 100 to Repeater Considerations	1-22
1.4	LAN Traffic Monitor Configurations	1-24
1.4.1	Single Port Configuration with Loopback Connector Installed	1-24
1.4.2	Dual Port Connections Between Two Ethernets	1-25
1.4.3	Dual Port Connections with Bridged Ethernets	1-26
1.5	Bridge/Router Considerations	1-27
1.5.1	LAN Bridge 100 Unit in Parallel with a Router	1-28
1.5.2	Modifying Router Parameters for Extended LANs	1-29
1.5.3	Setting Up Multiple Areas	1-30
1.6	Specifications	1-31
1.6.1	Physical Specifications	1-31
1.6.2	Environmental Specifications	1-32
1.6.3	Electrical Specifications	1-33

2 Installation and Verification

2.1	Overview	2-1
2.2	Unpacking	2-2
2.3	Arranging for Software Installation	2-4
2.4	Installation	2-6
2.4.1	Setting the Voltage Switch	2-6

2.4.2	Verifying Configuration Switch Settings	2-7
2.4.3	Rack-Mount Installation (Optional)	2-9
2.4.4	Wall-Mount Installation (Optional)	2-10
2.4.5	Connecting the Transceiver Cables	2-12
2.4.6	Connecting the Fiber-Optic Cable	2-13
2.4.7	Connecting the Power Cord	2-15
2.5	Verifying the Installation	2-15
2.5.1	Verifying the Bridge Installation	2-16
2.5.2	Verifying LAN Traffic Monitor (LTM) Installation	2-18

3 Operation

3.1	Chapter Overview	3-1
3.2	Operational States	3-1
3.2.1	SELF-TEST State	3-5
3.2.2	BROKEN State	3-5
3.2.3	INITIALIZATION State	3-5
3.2.3.1	Bridge Initialization	3-5
3.2.3.2	Down-Line Loading	3-6
3.2.4	PREFORWARDING State	3-6
3.2.5	FORWARDING State	3-7
3.2.6	BACKUP State	3-7
3.3	Packet Memory Data Structures	3-7
3.3.1	Initialization Blocks	3-7
3.3.2	Descriptor Rings	3-9
3.3.2.1	Receive Descriptor Ring Entry	3-11
3.3.2.2	Transmit Descriptor Ring Entry	3-14
3.4	Learning and Forwarding	3-18
3.4.1	Table Initialization	3-20
3.4.2	Address Table Entries	3-21
3.4.3	Binary Search	3-21
3.4.4	Forwarding	3-23
3.4.5	Writing the Ethernet Address Table	3-24
3.4.6	Special Cases	3-24
3.4.6.1	Stale Packets	3-24
3.4.6.2	Swapped Sides	3-24
3.4.6.3	Address Aging	3-25
3.4.7	Ethernet Address Table Maintenance	3-25
3.4.7.1	Inserting New Addresses	3-25
3.4.7.2	Single-Entry Compare	3-26
3.4.7.3	64-Bit Move	3-27
3.5	Loop Detection	3-27
3.5.1	The Spanning Tree Principle	3-28
3.5.2	The Spanning Tree Computation Process	3-28
3.5.3	Spanning Tree Parameters	3-30
3.5.4	Examples of the Spanning Tree Algorithm	3-31
3.6	Foreground and Background Operations	3-38
3.6.1	Foreground Operations	3-39
3.6.2	Background Operations	3-40
3.7	Remote Software and Bridge Access	3-41

3.8	Maintenance Operation Protocol (MOP)	3-42
3.9	Self-Test	3-42
3.9.1	Basic Tests	3-43
3.9.1.1	Program ROM Test	3-44
3.9.1.2	Program RAM Test	3-44
3.9.1.3	NVRAM Checksum Test	3-44
3.9.1.4	NVRAM Write Test	3-44
3.9.1.5	Ethernet Address ROM Checksum Test	3-44
3.9.1.6	Ethernet Address RAM Test	3-44
3.9.1.7	Timer Test	3-44
3.9.1.8	Packet Memory Test	3-44
3.9.1.9	Packet Memory Refresh Test	3-44
3.9.2	LANCE Tests	3-45
3.9.2.1	LANCE Reset Test	3-45
3.9.2.2	Internal Loop Tests	3-45
3.9.2.3	External Loop Tests	3-46
3.9.3	Table Lookup (TLU) Tests	3-47
3.9.3.1	Status RAM Test	3-47
3.9.3.2	Basic Binary Search Test	3-47
3.9.3.3	Binary Search Engine Test	3-47

4 Technical Description

4.1	Introduction	4-1
4.2	Hardware Overview	4-1
4.2.1	Overview of LAN Bridge 100 Functional Blocks	4-2
4.2.2	The Data Path	4-3
4.3	Processor Subsystem	4-5
4.3.1	Processor Circuit Descriptions	4-5
4.3.2	The Microprocessor	4-7
4.3.3	Interrupt Controller	4-7
4.3.4	Processor Memory	4-9
4.3.4.1	PROM	4-9
4.3.4.2	Program RAM	4-9
4.3.4.3	Address Table Status RAM	4-9
4.3.4.4	Physical Address PROM	4-9
4.3.4.5	NVRAM	4-9
4.3.4.6	Miscellaneous Control Registers	4-10
4.3.5	Processor Memory Map	4-10
4.3.5.1	Program ROM	4-12
4.3.5.2	Miscellaneous Control	4-12
4.3.5.3	Nonvolatile Memory	4-13
4.3.5.4	LANCE Descriptor Rings	4-13
4.3.5.5	LANCE CSR Registers	4-13
4.3.5.6	Program RAM	4-13
4.3.5.7	Packet Memory	4-13
4.3.5.8	Ethernet Address Table	4-13
4.3.5.9	Ethernet Address Table Status RAM	4-13
4.3.5.10	Compare-and-Move Registers	4-13
4.3.5.11	Ethernet Address Table 64-Bit Move	4-14

4.3.6	System Timers	4-14
4.3.7	Reset Circuitry	4-14
4.4	Network Interconnect Subsystem	4-15
4.4.1	NI Subsystem Circuits	4-15
4.4.2	NI Ports	4-16
4.4.2.1	SIA Chip	4-17
4.4.2.2	LANCE Chip	4-17
4.4.2.3	LANCE CSR Bus Registers	4-19
4.4.2.4	LANCE CSR Control Circuit	4-19
4.4.2.5	Fiber-Optic Module	4-19
4.4.3	LANCE Address Latch	4-19
4.4.4	OWN Interrupt Generator	4-20
4.4.5	Packet Memory	4-20
4.4.5.1	Address Multiplexer	4-20
4.4.5.2	Memory Control Circuit	4-21
4.4.5.3	Dynamic Packet Memory	4-22
4.4.5.4	Refresh Counter	4-23
4.4.5.5	Data Bus Latches	4-23
4.5	Table Lookup Subsystem	4-23
4.5.1	Overview of the TLU	4-23
4.5.2	Circuit Descriptions	4-26
4.5.3	Search Address Multiplexer	4-26
4.5.4	48-Bit Address Comparator	4-27
4.5.4.1	Functional Description	4-27
4.5.4.2	Address Comparator Circuit Descriptions	4-30
4.5.4.3	Ethernet Address Table Memory	4-30
4.5.4.4	Bus Transceivers	4-30
4.5.4.5	Ethernet Search Address Comparators	4-30
4.5.4.6	Search Control and Status PALS	4-31
4.5.4.7	Search Results Register	4-31
4.6	Power Supply	4-32

5 Maintenance

5.1	Scope	5-1
5.2	Maintenance Philosophy	5-1
5.2.1	Required Equipment	5-1
5.2.2	Optional Equipment	5-2
5.3	Preventive Maintenance	5-3
5.4	Corrective Maintenance	5-4
5.4.1	Troubleshooting Tips	5-4
5.4.2	Fault Diagnosis	5-5
5.5	LAN Bridge 100 Replacement Procedures	5-14
5.6	Bridge Disassembly	5-15
5.6.1	Plastic Enclosure Removal	5-16
5.6.2	Removing the Chassis Cover	5-18
5.6.3	Removing the Fiber-Optic Interface	5-21
5.6.4	Removing the Power Supply	5-23
5.6.5	Removing the Logic Module	5-29

A Glossary

B Fiber-Optic Link Analysis

B.1	Introduction	B-1
B.2	System Conventions	B-2
B.2.1	Absolute and Relative Measurements	B-2
B.2.2	Worst-Case Definition	B-3
B.3	Coupling Considerations	B-3
B.3.1	LED (Emitter)-to-Fiber	B-4
B.3.1.1	LED-to-Fiber Coupling Derating Factors	B-5
B.3.2	Fiber-to-Fiber Coupling	B-6
B.3.2.1	Splice and Barrel Connector Losses	B-7
B.3.2.2	Guidelines for Connectors and Splices on Ethernet Products	B-8
B.3.2.3	Typical Cable Attenuation Values	B-10
B.3.2.4	Single-Window and Dual-Window Fiber	B-10
B.3.2.5	Single-Mode Fiber	B-10
B.3.3	Fiber-to-Pin Diode (Receiver)	B-11
B.4	Numerical Aperture and Coupling Diameters	B-11
B.4.1	Bending Losses	B-13
B.5	System Losses and Variability	B-14
B.6	Attenuation Measurement Procedures	B-14
B.6.1	Field Measurement Procedure for Relative Power Loss	B-15
B.6.1.1	Equipment Required	B-15
B.6.1.2	Preparation	B-15
B.6.1.3	Procedure	B-15
B.6.1.4	Test Source Adjustment Procedure	B-16
B.6.1.5	Measurement Correction Calculation Procedure	B-17
B.6.1.6	System Certification Procedure	B-19
B.6.2	Optical Test Set (Single-Cable) Method	B-20
B.6.3	Optical Time Domain Reflectometry	B-22
B.6.3.1	Equipment Setup	B-23
B.6.3.2	Link Measurement	B-23
B.7	Bandwidth Analysis and Derating Factors	B-25
B.7.1	Fiber Bandwidth Rating System	B-25
B.7.2	Distance Derating Factors for Ethernet Products	B-25
B.8	Suggested System Design Methodology	B-26
B.8.1	Example Worksheets	B-27
B.8.1.1	Sample Loss Calculation	B-28
B.8.2	Worst-Case and Statistical Results	B-30
B.9	Attenuator Installation	B-30
B.9.1	When to Install the Attenuator	B-31
B.9.2	Where to Install the Attenuator	B-33
B.9.3	How to Install the Attenuator	B-34
B.10	Reference Material	B-36

Index

Figures

1-1	LAN Extended by Means of a LAN Bridge 100 Unit	1-2
1-2	Several LANs Connected with LAN Bridge 100 Units	1-3
1-3	Local and Remote Versions of the LAN Bridge 100 Unit	1-6
1-4	LAN Bridge 100 Unit with Plastic Enclosure	1-7
1-5	LAN Bridge 100 Unit with Plastic Enclosure Removed	1-8
1-6	Local LAN Bridge 100 Controls, Status LEDs, and Connectors (Switches Shown in the Default Position)	1-9
1-7	Remote LAN Bridge 100 Controls, Status LEDs, and Connectors	1-10
1-8	LAN Bridge 100 Configurations	1-17
1-9	Typical Extended LAN with a Local LAN Bridge 100 Unit	1-19
1-10	Example Showing a Fiber-Optic Link as a LAN	1-20
1-11	Extended LANs Using Remote Bridges	1-21
1-12	Remote Bridge to Repeater Configuration	1-23
1-13	LTM Single Port Configuration	1-24
1-14	LTM Connected to Two Separate Ethernets	1-25
1-15	LTM on Two Connected Ethernets	1-26
1-16	Bridge in Parallel with a Router	1-28
1-17	An Extended LAN with Routers	1-29
2-1	Unpacking the LAN Bridge 100 Unit	2-2
2-2	Checking the LAN Bridge 100 Contents	2-3
2-3	Location of Serial Number and Ethernet Address	2-5
2-4	Setting the Voltage Select Switch	2-6
2-5	Verifying Configuration Switch Settings	2-7
2-6	Removing the Plastic Enclosure from the LAN Bridge 100 Unit	2-8
2-7	Installing the LAN Bridge 100 Unit in a Rack	2-9
2-8	Fastening the Wall-Mounting Brackets to the LAN Bridge 100 Unit ...	2-10
2-9	Positioning the LAN Bridge 100 Unit for Wall Mounting	2-11
2-10	Connecting the Transceiver Cable	2-12
2-11	Removing the Protective Caps from the Fiber-Optic Connectors	2-13
2-12	Connecting the Fiber-Optic Cable	2-14
2-13	Connecting the Power Cord	2-15
2-14	Bridge Hardware Verification	2-16
2-15	LTM Hardware Verification	2-18
3-1	Operating State Transition	3-2
3-2	LANCE Initialization Block Field Functions	3-8
3-3	Descriptor Ring Structures	3-10
3-4	Bit Format of Word 0 in the Receive Descriptor Ring Entry	3-11
3-5	Bit Format of Word 1 in the Receive Descriptor Ring Entry	3-12
3-6	Bit Format of Word 2 in the Receive Descriptor Ring Entry	3-13
3-7	Bit Format of Word 3 in the Receive Descriptor Ring Entry	3-13
3-8	Bit Format of Word 0 in the Transmit Descriptor Ring Entry	3-14
3-9	Bit Format of Word 1 in the Transmit Descriptor Ring Entry	3-15
3-10	Bit Format of Word 2 in the Transmit Descriptor Ring Entry	3-16
3-11	Bit Format of Word 3 in the Transmit Descriptor Ring Entry	3-17
3-12	Table Lookup Process	3-19
3-13	Address Table Entry	3-21

3-14	Binary Decision Tree	3-22
3-15	Binary Search Process	3-23
3-16	Inserting a New Address into the Ordered Table	3-26
3-17	Result of the Spanning Tree Computation Algorithm	3-29
3-18	Hello Message Propagation	3-30
3-19	Hypothetical Extended LAN Using Bridges	3-32
3-20	Logical Extended LAN Formed by Learning Algorithm	3-33
3-21	Reconfigured Logical Extended Network	3-34
3-22	Reconfigured Logical LAN with a New Root Bridge	3-35
3-23	Extended LAN Showing Physical Loops	3-36
3-24	Logical Extended LAN with Backup Bridges	3-37
3-25	Logical Network Topology for New Root Bridge	3-38
4-1	LAN Bridge 100 Subsystems	4-2
4-2	LAN Bridge 100 Functional Block Diagram	4-3
4-3	Processor Subsystem	4-6
4-4	Hardware Interrupt Configuration	4-8
4-5	Address Mapping Used During Power Up	4-11
4-6	Address Mapping Used for Normal LAN Bridge 100 Operation	4-12
4-7	NI Subsystem	4-16
4-8	CSR Control	4-18
4-9	Memory Control Circuit	4-22
4-10	Table Lookup Subsystem	4-24
4-11	48-Bit Comparator	4-28
4-12	16-Bit Comparator	4-29
5-1	LAN Bridge 100 Status LEDs	5-3
5-2	Troubleshooting Flowchart	5-6
5-3	Removing the Top Plastic Enclosure and Side Pieces	5-16
5-4	Removing the Bottom Plastic Enclosure	5-17
5-5	Removing the Chassis Cover for Revision F08 and Below	5-19
5-6	Various Removal Procedures for Revision F09 and Above	5-20
5-7	Fiber-Optic Interface Retaining Hardware for Revision F08 and Below	5-21
5-8	Removing the Fiber-Optic Interface for Revision F08 and Below	5-22
5-9	Removing the Power Supply Connectors and Retaining Screws for Revision F08 and Below	5-24
5-10	Removing the Power Supply for Revision F08 and Below	5-25
5-11	Removing the Power Supply for Revision F09 and Above	5-26
5-12	Power Supply Assembly Removed from the Chassis (Revision F08 and Below)	5-27
5-13	Power Supply Assembly Removed from the Chassis (Revision F09 and Above)	5-28
5-14	Locations of Logic Module Retaining Hardware	5-30
B-1	Fiber-Optic System Responsibilities	B-2
B-2	Area of Network Consultant's Concern-Allowable System Loss	B-3
B-3	LED Coupling into a Fiber	B-4
B-4	Types of Fiber Misalignment	B-9
B-5	Fiber-to-Receiver Coupling	B-11
B-6	Near-End and Far-End Coupling Losses	B-12
B-7	Minimum Bend Radius	B-13
B-8	End-to-End Cable Loss Calculation Worksheet	B-17

B-9	Calibrating the Optical Test Set	B-21
B-10	Measuring the Cable Under Test	B-22
B-11	OTDR Test Configuration	B-23
B-12	OTDR Trace	B-24
B-13	Sample Worksheet for Budget Loss Calculation	B-26
B-14	Sample Worksheet for DEBET-RH Loss Calculation	B-28
B-15	Sample Worksheet for DEBET-RC Loss Calculation	B-29
B-16	Attenuator	B-30
B-17	Installing the Attenuator on Both Sides of the Link	B-32
B-18	Installing the Attenuator on -RH/-RJ Side Only	B-32
B-19	Removing the Protective Caps from the LAN Bridge 100	B-34
B-20	Connecting the Fiber-Optic Cables to the LAN Bridge 100	B-35
B-21	Fiber-Optic Cable Installed to the LAN Bridge 100	B-36

Tables

1-1	Versions of the LAN Bridge 100 Unit	1-5
1-2	LAN Bridge 100 Product Designations	1-7
1-3	LAN Bridge 100 Controls	1-11
1-4	LAN Bridge 100 Switches	1-11
1-5	LAN Bridge 100 Status LEDs	1-13
1-6	LAN Bridge 100 Connectors	1-14
3-1	Bit Descriptions of Word 0 in the Receive Descriptor Ring Entry	3-11
3-2	Bit Descriptions of Word 1 in the Receive Descriptor Ring Entry	3-12
3-3	Bit Descriptions of Word 2 in the Receive Descriptor Ring Entry	3-13
3-4	Bit Descriptions of Word 3 in the Receive Descriptor Ring Entry	3-13
3-5	Bit Descriptions of Word 0 in the Transmit Descriptor Ring Entry	3-14
3-6	Bit Descriptions of Word 1 in the Transmit Descriptor Ring Entry	3-15
3-7	Bit Descriptions of Word 2 in the Transmit Descriptor Ring Entry	3-16
3-8	Bit Descriptions of Word 3 in the Transmit Descriptor Ring Entry	3-17
3-9	Background Operations	3-40
3-10	Bridge Access Switches	3-41
3-11	MOP Request and Response Messages	3-42
4-1	Hardware Interrupts	4-8
4-2	Device Priorities	4-21
4-3	Search Results Register Bit Descriptions	4-31
5-1	Controlled Distribution Kits for the LAN Bridge 100 Unit	5-2
B-1	Types of Optical Coupling	B-4
B-2	LED-to-Fiber Coupling Derating Factors for -RC/-RD Variations	B-5
B-3	LED-to-Fiber Coupling Derating Factors for -RH/-RJ Variations	B-6
B-4	Splice and Barrel Connector Losses	B-7
B-5	Typical Cable Attenuation	B-10
B-6	Fiber Size Correction Settings for the Power Meter Reading	B-16
B-7	Wavelength Correction Values	B-18
B-8	Cable Certification Values	B-19

Preface

The *LAN Bridge 100 Technical Manual* is a reference document that provides technical information on the LAN Bridge 100 network bridge. This document provides an overview, installation and operational information, block-diagram level functional descriptions, and maintenance information on the LAN Bridge 100.

This manual is organized as follows:

- Chapter 1 Provides a general overview of the LAN Bridge 100. Configuration considerations and LAN Bridge 100 specifications are also included.
- Chapter 2 Provides installation information for the LAN Bridge 100.
- Chapter 3 Discusses LAN Bridge 100 operating states and data structures. Also described are various features provided by the LAN Bridge 100 such as loop detection and message filtering.
- Chapter 4 Provides block-level functional descriptions of the LAN Bridge 100 hardware circuits.
- Chapter 5 Describes the maintenance procedures for the LAN Bridge 100.
- Appendix A Defines some terms used in the *LAN Bridge 100 Technical Manual*.
- Appendix B Describes differences in transmission characteristics among the various types of fiber used in fiber-optic cables.

Related Documents

Additional information about the LAN Bridge 100 can be found in the following documents:

- *Communications Options Minireference Manual, Vol. 4* (Order no. EK-CMIV4-RM)
- *DECconnect System General Description* (Order no. EK-DECSY-GD)
- *DECconnect System Planning and Configuration Guide* (Order no. EK-DECSY-CG)
- *DNA Ethernet Node Product Architectural Specification* (Order no. AA-X440A-TK)

- *DNA Maintenance Operations Functional Specification* (Order no. AA-X436A-TK)
- *DECnet Router Installation/Operations Manual* (Order no. AA-X019-TK)
- *DNA Routing Layer Functional Specification* (Order no. AA-X435A-TK)
- *Ethernet Communications Server DECnet Router Software Installation Guide* (Order no. AA-X019B-TK)
- *Ethernet Transceiver Tester User's Guide* (Order no. EK-ETHTT-UG)
- *Field Maintenance Print Set* (Order no. MP-01785-01)
- *Guide to Networking on VAX/VMS* (Order no. AA-Y512A-TE)
- *LAN Bridge 100 Installation/User's Guide* (Order no. EK-DEBET-UG)
- *LAN Traffic Monitor Installation Guide* (Order no. AA-JP15A-TE)
- *LAN Traffic Monitor User's Guide* (Order no. AA-JP16A-TE)
- *LAN Traffic Monitor Identification Card* (Order no. EK-LANTM-IC)
- *Remote Bridge Management Software Guide* (Order no. AA-FY93A-TE)
- *Attenuator Installation and Configuration Reference Card* (Order no. EK-DEFOE-RC)

1

Introduction

1.1 Introduction to the LAN Bridge 100 Unit

The introduction of Ethernet and IEEE 802.3 local area networks (LANs) has reduced the cost and increased the capability of networking. This improvement has resulted in an increased demand for networking. To accommodate the increased demand, larger and denser networks are being created. However, as networks approach the design limitations of single-LAN technology, their performance may be degraded by the limitations.

The LAN Bridge 100 unit minimizes many single-LAN limitations by creating a high-speed, logical link between two LANs. Networks that are joined by bridges are called extended LANs. Note that each individual LAN can be of maximal configuration in terms of length, number of stations, and other specifications.


The LANs that make up this extended network can be either IEEE 802.3-type or Ethernet-type baseband or broadband. The LAN Bridge 100 unit provides this logical extension without creating a bottleneck in the network.

The LAN Bridge 100 unit dynamically learns the locations and station addresses of nodes for each of the networks that are connected to the bridge. This learning capability enables the LAN Bridge 100 unit to forward packets selectively, based on their destination addresses. In this way, the bridge can create an extended LAN that has the following advantages:

- More stations — Each LAN can still support its maximum number of stations including the bridge. However, the bridge is transparent to the stations; so the two LANs appear as one larger, extended LAN to all users.
- Larger network — Each LAN can still support the maximum length for a LAN, but the extended LAN can be much longer.
- Reduced traffic — The bridge forwards only nonlocal traffic. Thus, if a large LAN is broken into several smaller LANs, the traffic on any one of these smaller LANs may be greatly reduced.

1.1.1 Extended Networks


Figure 1-1 shows an example of an extended network that uses a LAN Bridge 100 unit. If these two networks were joined with a repeater instead of a bridge, all packets that originate on one of the networks would also appear on the other network. The LAN Bridge 100 unit filters packets so that only those packets destined for a station on the opposite network are forwarded. Packets destined for stations on the same network are filtered. This does not mean that all repeaters should or can be replaced by bridges. The two devices are functionally different, and a careful evaluation of the network requirements is needed to determine which device is more appropriate in a given situation.



LKG-0726

Figure 1-1: LAN Extended by Means of a LAN Bridge 100 Unit

The Ethernet-type and/or IEEE 802.3-type networks that a bridge can connect can be two maximum-length broadband, two baseband (each with the maximum number of stations), or one of each. The LAN Bridge 100 unit can connect the two Ethernet-type LANs so that they appear as a single extended LAN. This allows a logical extension of an Ethernet LAN beyond the normal limitations of 2800 meters for baseband or 3800 meters for broadband Ethernet (see Figure 1-2). Broadband and baseband Ethernet networks are each limited to 1024 stations; however, LAN Bridge 100 units can overcome this limitation by joining two or more LANs together, thereby creating an extended LAN.



*The 2800-meter (9194-foot) maximum distance between any two nodes is the sum of two 50-meter (164-foot) transceiver cables, three 500-meter (1640-foot) coaxial cable segments, four 50-meter (164-foot) transceiver cables connected to repeaters, and 1000 meters (3280-foot) of point-to-point link.

LKG-0727

Figure 1-2: Several LANs Connected with LAN Bridge 100 Units

1.1.2 Transparent Operation

Bridge operation is transparent to other stations on the LAN, and no special software is required on any station unless the LAN Bridge 100 unit is configured to operate as a LAN Traffic Monitor (more information on the LTM feature is provided in the following sections). Remote Bridge Management Software (RBMS) is available for VMS hosts. RBMS allows you to observe and control any LAN Bridge 100 unit in the network.

1.1.3 LAN Bridge 100 to Network Connections

The LAN Bridge 100 unit connects to the network through a transceiver cable and one of the following devices:

- H4000 Ethernet transceiver (baseband)
- DELNI local network interconnect
- DECOM broadband modem
- DESTA Ethernet/IEEE 802.3 transceiver

The LAN Bridge 100 unit can be either a local bridge or a remote bridge. The local bridge is connected to two LANs through transceiver cables. The remote bridge is connected to one LAN through a transceiver cable and to either another remote bridge or a remote repeater over a fiber-optic link. Remote bridges can be used when the distance between LANs is more than 100 meters or where adverse environmental conditions exist.

The LAN Bridge 100 unit operates at the Data Link layer of the ISO model and therefore is transparent to protocols above this layer. There are no jumper or switch settings required to make the bridge compatible with either 802.3-type or Ethernet-type LANs.

1.2 LAN Bridge 100 Functional Description

The LAN Bridge 100 unit provides a logical link between two LANs and extends the range of the LAN. The LAN Bridge 100 unit actually minimizes many of the limitations of conventional LAN systems and provides the following advantages:

- The LAN Bridge 100 unit selectively forwards or filters (disregards) packets based on the destination address contained in each packet.

A dynamic address-learning capability enables the LAN Bridge 100 unit to acquire a working knowledge of the network configuration. By storing source addresses from received packets, the bridge learns which port (port A or B of the bridge) is associated with each active station on the extended network.

Once the LAN Bridge 100 unit has built a database of station locations, the bridge selectively forwards packets based on the destination address. This minimizes network congestion by keeping local traffic local and allows literally thousands of stations to be connected to the extended network.

- The LAN Bridge 100 unit functions at the Data Link layer and is protocol independent. This allows non-Digital Ethernet LANs to be included in the extended LANs. Typical protocols include DECnet, Xerox Network System (XNS), Transmission Control Protocol (TCP/Internet Protocol), Local Area Transport (LAT), or any protocols based on Ethernet or on IEEE 802.3 standards.
- The LAN Bridge 100 unit has an automatic backup feature that is based on its ability to learn the locations of other bridges in the extended network. When bridges are configured in a loop, one of the bridges automatically enters a BACKUP state. Thus the data link loop is broken and the BACKUP LAN Bridge 100 unit serves as a warm standby. This enhances network availability.


- The LAN Traffic Monitor function is optional software available for use with LAN Bridge 100 units at ROM ECO Revision Level E, or later. This optional feature allows the LAN Bridge 100 unit to be used as a base from which to gather traffic data in the form of counters, and periodically forward them to a VAX/VMS system for compilation and analysis. More information on this optional feature is provided in Section 1.2.4.
- Optional Remote Bridge Management Software (RBMS) allows network operators to monitor and control individual bridges in the extended network. By using RBMS to change bridge parameters or to load a forwarding database over the network, operators can control the extended network better.

1.2.1 Product Versions and Designations

There are two versions of the LAN Bridge 100 unit: the local bridge version and the remote bridge version. Table 1-1 describes the two versions. Both versions are shown in Figure 1-3.

Table 1-1: Versions of the LAN Bridge 100 Unit

Version	Description
Local Bridge	The local bridge connects LANs separated by 100 meters (328 feet) or less. The distance from the bridge to either LAN cannot exceed the maximum allowable transceiver cable length of 50 meters (164 feet).
Remote Bridge	<p>The remote bridge connects LANs separated by more than 100 meters (328 feet) or where fiber-optic cable capabilities are needed.</p> <p>A fiber-optic cable is used either to connect two remote bridges or to connect a remote bridge and a remote repeater.</p> <p>The fiber-optic cable can be up to 3000 meters (9840 feet) in length when connecting two remote bridges or up to 1500 meters (4921 feet) when connecting a remote bridge and a remote repeater.</p>



LKG-0950

Figure 1-3: Local and Remote Versions of the LAN Bridge 100 Unit

There are four model designations of the LAN Bridge 100 unit, as shown in Table 1–2. Note that the only difference between U.S. and European versions is the product labeling and the voltage selection switch setting.


Table 1–2: LAN Bridge 100 Product Designations

Model	Version	Power Requirements	Product Labeling
DEBET-AA	Local Bridge	120 Vac Nominal (88 Vac to 132 Vac)	U.S.
DEBET-AB	Local Bridge	240 Vac Nominal (176 Vac to 264 Vac)	Non-U.S.
DEBET-RC	Remote Bridge	120 Vac Nominal (88 Vac to 132 Vac)	U.S.
DEBET-RH	Extended Remote Bridge	120 Vac Nominal (88 Vac to 132 Vac)	U.S.
DEBET-RD	Remote Bridge	240 Vac Nominal (176 Vac to 264 Vac)	Non-U.S.
DEBET-RJ	Extended Remote Bridge	240 Vac Nominal (176 Vac to 256 Vac)	Non-U.S.

1.2.2 LAN Bridge 100 Packaging

The LAN Bridge 100 unit is housed in a plastic enclosure and is ready for table-top installation. The plastic enclosure is easily removed to allow for rack or wall mounting. An optional wall-mounting kit (part number H039) is available that allows mounting the bridge to a wall or partition while it is still in its plastic enclosure. The LAN Bridge 100 unit contains its own power supply and cooling fans.


Figures 1–4 and 1–5 show the LAN Bridge 100 unit with and without the plastic enclosure.



**LAN BRIDGE 100
(LOCAL BRIDGE)**

LKG-0467

Figure 1–4: LAN Bridge 100 Unit with Plastic Enclosure




**LAN BRIDGE 100
(REMOTE BRIDGE)**

LKG-0468

Figure 1-5: LAN Bridge 100 Unit with Plastic Enclosure Removed



1.2.3 LAN Bridge 100 Controls, Status LEDs, and Connectors

All the controls, status LEDs, and connectors are located on the I/O panel of the LAN Bridge 100 unit. The I/O panel of a local bridge is shown in Figure 1-6, the I/O panel of a remote bridge, in Figure 1-7. Tables 1-3 through 1-6 describe the controls, status LEDs, and connectors available on the local and remote bridges.



LKG-0960

Figure 1-6: Local LAN Bridge 100 Controls, Status LEDs, and Connectors (Switches Shown in the Default Position)



CONTROLS, CONNECTORS, AND INDICATORS ARE SIMILAR TO THOSE OF THE LOCAL BRIDGE (SEE FIGURE A-1), EXCEPT THAT PORT A IS A FIBER-OPTIC LINK INTERFACE.

LKG-0937

Figure 1-7: Remote LAN Bridge 100 Controls, Status LEDs, and Connectors

Table 1–3: LAN Bridge 100 Controls

Control	Description
Voltage Select Switch	The voltage select switch is used to set the bridge input voltage to the range required for operation in your country. This switch was factory set for the correct power source for your country. Do not change this switch setting unless you are sure that the switch setting is incorrect (see your electrician if you are not sure). Chapter 2, Section 2.4 provides information for changing the bridge input voltage setting, if necessary.
Circuit Breaker	The circuit breaker provides overcurrent protection for the bridge. If an overcurrent condition causes the circuit breaker to trip, the white center portion of the circuit breaker pops out as a visual indication, and the power is cut off from the bridge. The circuit breaker can be reset by pressing in the white center portion of the circuit breaker.
Bridge Switches	These switches control the LAN Bridge 100 functions. Each switch is described in Table 1–4.

Table 1–4: LAN Bridge 100 Switches

Switch Number	Name	ON (Down)	OFF (Up)
1	Loop Self-Test	The bridge loops self-test continuously after powerup. Loopback terminators must be installed for this test. This switch position is for manufacturing and Field Service use only.	Bridge runs self-test once on powerup or reset.
2	NVRAM Reset	NVRAM resets to factory default settings when the bridge is powered up. NVRAM Reset removes all bridge management configuration changes.	Prevents NVRAM from resetting to factory default settings when the bridge is powered up. This setting should be used to prevent the loss of parameters stored by RBMS, during a power failure.
3	Port A Access	Stations on the LAN connected to port A that have bridge management capabilities are allowed to read and write (modify) bridge management parameters. If a load host resides on port A, the switch must be on for down-line loading software.	Stations on the LAN connected to port A that have bridge management capabilities can read but cannot write bridge management parameters.

Table 1–4 (Cont.): LAN Bridge 100 Switches

Switch Number	Name	ON (Down)	OFF (Up)
4	Port B Access	Stations on the LAN connected to port B that have bridge management capabilities are allowed to read and write (modify) bridge management parameters. If a load host resides on port B, the switch must be on for down-line loading software.	Stations on the LAN connected to port B that have bridge management capabilities are allowed to read but cannot write bridge management parameters.
5	Down-Line Load Enable	Configures the unit to operate as a LAN Traffic Monitor. Enables the unit to down-line load software (such as LTM Listener software) from a load host. When the switch is on, the bridge does not forward packets. Applicable Port A Access or Port B Access switch must be on so that the load host can write to bridge memory.	Configures unit to operate as a bridge. RBMS can override the switch causing the bridge to operate as a LAN Traffic Monitor.
6	Not Used		

NOTES

1. You can change the switch settings for port access while the bridge is operating. However, the switch settings for Loop Self-Test, NVRAM Reset and Down-Line Load Enable are read only during power up. Changing either of these switches while the bridge is operating has no effect on bridge operation. To change either of these switches, unplug the unit, change the setting, then plug the unit back in.
2. Port A and Port B Access switches can prevent bridge management software from changing any of the bridge's internal parameters. If security is a concern at the site, set the bridge's parameters with RBMS and then disable one or both ports by putting one or both switches in the up position. Bridge management software can still read the bridge's counters and other parameters. Placing either switch in the down position enables bridge management software write access from stations on the LAN connected to that port of the bridge. Normally, both switches are placed in the down position to enable bridge management software write access from stations on either LAN.

Table 1-5: LAN Bridge 100 Status LEDs

Name	ON Steady	OFF	Blinking
Port A Activity	A message is being received or transmitted on port A.	No message traffic on port A.	Infrequent messages are being received or transmitted on port A or the bridge is checking for loops (sending Hello messages) about once a second.
On-line	Unit is configured as a bridge, is fully operational, and is forwarding messages.	When configured as a bridge, the unit is in the INITIALIZE, PREFORWARDING, BACKUP, or BROKEN state. When configured as an LTM Listener, the load host has not completed down-line loading the LTM Listener software image. (Down-line loading typically takes up to two minutes depending on network traffic.)	Flashing twice every 2 seconds indicates that the load host successfully down-line loaded the LTM Listener software image. Flashing once each second indicates that the load host has started the LTM Listener function.
DC OK	Internal power supply is functioning properly.	Internal power supply is not functioning properly.	Internal power supply is not functioning properly.
Self-Test OK	Passed self-test.	Running self-test.	NVRAM failed, which requires replacement. This failure does not affect normal operation.*
Port B Activity	A message is being received or transmitted on port B.	No message traffic on port B.	Infrequent messages are being received or transmitted on port B or the bridge is checking for loops (sending Hello messages) about once a second.

* NVRAM stores network pointers, parameters, and addresses set by RBMS so that they are not lost in the event of a power failure. If the Self-test OK LED is blinking, the fault may be bypassed by setting the NVRAM Reset switch (switch 2 on the I/O panel) to the down (ON) position, and then turning the bridge power off and on. (See Chapter 5 for more information on troubleshooting.) Note that this causes the bridge to use default parameters.

Table 1-6: LAN Bridge 100 Connectors

Connector	Description
AC Input	This connector accepts ac input voltages of 120 or 240 Vac, depending on the setting of the voltage selection switch (refer to Table A-1).
Port A	<p>For local bridges, this 15-pin, female, D-type connector accepts a transceiver cable. A slide latch is provided for locking the transceiver cable in place. The pins have the following definitions:</p> <ol style="list-style-type: none">1. Chassis ground2. Collision presence +3. Transmit +4. Ground5. Receive +6. + 12 volt return7. No connection8. Ground9. Collision presence -10. Transmit -11. Ground12. Receive -13. + 12 volts14. Ground15. Ground <p>For remote bridges, port A has two fiber-optic connectors. The left-hand connector (marked ) is for receiving optical data. The right-hand connector (marked ) is for transmitting optical data.</p>
Port B	For local and remote bridges, this 15-pin, female, D-type connector accepts a transceiver cable. A slide latch is provided for locking the transceiver cable in place. The pins have the same definitions as for port A (described in this table).

1.2.4 LAN Traffic Monitor (LTM) Option

The LAN Traffic Monitor (LTM) is an Ethernet/IEEE 802.3 LAN monitor that uses the LAN Bridge 100 unit as a hardware base. Either version of the LAN Bridge 100 unit (local or remote) supports the LTM option. When the LAN Bridge 100 unit is configured to operate as a LAN Traffic Monitor, bridge operations are suspended until the unit is reconfigured for bridge operation. The LAN Bridge 100 hardware unit processes 48-bit Ethernet addresses, and the LTM software calculates the statistics. The statistics are periodically reported to a host system that performs additional data reduction, such as averaging and peak traffic analysis. LTM has two components:

- The LTM Listener — a LAN Bridge 100 hardware unit that is down-line loaded with LTM monitoring software. The LAN Bridge 100 firmware must be ROM ECO Revision Level E or later.
- The LTM User Interface (UI) — remote application software that is installed on any DECnet VAX/VMS system with an Ethernet controller and associated driver.

For more information about the LAN Traffic Monitor, refer to the *LAN Traffic Monitor User's Guide*.

1.2.5 LAN Traffic Monitor Software

The basic software required for installing and operating the LTM follows:

- LAN Traffic Monitor distribution software — installed on each LTM load host.
- DECnet Phase IV software, running on VMS Version 4.4 or later — installed on each LTM load host.

The distribution software must be installed on a load host that runs DECnet Phase IV software and that is connected to the same extended LAN as the LTM Listener. Digital recommends installing the LTM Listener software on a load host that is on the same LAN as the LTM Listener. Doing so avoids the possibility of segmenting the load host from the LTM Listener due to a bridge failure. The distribution software includes an LTM Listener software image file that is down-line loaded to the LTM Listener. All software must be installed and verified before operating the LTM.

1.2.6 Remote Bridge Management Software (RBMS)

Remote Bridge Management Software (RBMS) is an optional product available for VMS hosts. RBMS significantly enhances the network's operation by allowing you to observe and control bridges in the network. RBMS allows you to:

- Understand and modify your network topology by displaying and modifying the bridge forwarding database.
- Evaluate network performance by displaying bridge counters, status, and characteristics.
- Troubleshoot network problems by understanding your network topology, disabling selected bridges to segment your network, and signaling selected bridges to run their built-in self-test diagnostics.
- Save your configuration data in the bridge's nonvolatile RAM (NVRAM) so that it is not lost during a power failure. Changes are saved by bridge hardware.
- Remotely switch the LAN Bridge 100 unit between bridge usage and LTM Listener.
- Remotely determine whether the LAN Bridge 100 unit is operating as an LTM Listener or a bridge.

For more information on RBMS, refer to the *Remote Bridge Management Software Guide*.

1.3 Configuration Considerations

This section describes some configuration considerations that apply to implementing extended LANs with LAN Bridge 100 units. Refer to the *DECconnect Planning and Configuration Guide* (EK-DECSY-CG) for additional information on LAN configuration.


LANs connected by bridges appear as one extended LAN as far as data traffic is concerned.

Individual LANs that are connected by bridges can each be configured for up to the normal maximum for length, number of stations, and other specifications. For example, each Ethernet baseband LAN can be up to 2800 meters (9194 feet) in length and have 1024 stations.

Extended LANs may consist of combinations of any of the following LANs joined by bridges:

- Ethernet baseband
- Ethernet broadband
- IEEE 802.3 baseband (10base5)
- ThinWire Ethernet (10base2)

Within these configurations, bridges may be connected to the network through transceivers. Alternatively, remote bridges may be connected directly to other remote bridges or remote repeaters. Also, a bridge may be connected to a DELNI that may or may not be connected to a remote network. Figure 1-8 shows some of the possible configurations involving bridges.



LKG-0857-87

Figure 1-8: LAN Bridge 100 Configurations

1.3.1 Performance Considerations

A packet may have to travel through a number of bridges before reaching its destination. Note that increasing the number of bridges in the data path causes corresponding increases in the data path delay. This delay could have a negative impact on network performance, especially with time-critical protocols or with interactive tasks such as character echoing for users on terminal servers. A general rule for networks with typical traffic loading is that performance may start to degrade if a packet must travel through more than seven bridges to get from its source station to its destination station.

1.3.2 Loop Considerations

When the LAN Bridge 100 unit is turned on, it executes an internal self-test. This test takes about 15 seconds. The bridge then spends about 30 seconds learning station addresses and communicating with other bridges in the network to determine whether there are any loops (multiple paths between two or more LANs).


When bridges in an extended LAN form a loop, a loop detection process determines that one or more of the looped bridges enters the BACKUP state, so that only one path exists between any two LANs (the loop detection process is described in Chapter 3).

If an on-line bridge fails, a backup bridge takes over and begins forwarding packets. With RBMS software, bridges can be selectively placed in the BACKUP state. In this way the most direct path with the fewest number of bridges can be provided for the heaviest network traffic.

If the bridge is in a loop with a repeater, the bridge enters the BACKUP state. The bridge continues to check the loop through the repeater about once a second. If the repeater fails, the bridge automatically takes over and begins forwarding packets.

1.3.3 Local LAN Bridge 100 Considerations

The local LAN Bridge 100 unit (DEBET-AA or DEBET-AB) connects two LANs that are separated by less than 100 meters (328 feet). This distance is made up of the combined length of two transceiver cables, each 50 meters (164 feet) in length. Figure 1-9 shows a typical extended LAN configuration using a local LAN Bridge 100 unit.




LKG-0855-87

Figure 1-9: Typical Extended LAN with a Local LAN Bridge 100 Unit

1.3.4 Remote LAN Bridge 100 Considerations

Two remote LAN Bridge 100 units (DEBET-RC/RD) connect two LANs that are separated by up to 2000 meters (6560 feet) or two Extended Remote LAN Bridge 100 units (DEBET-RH/RJ) connect two LANs separated by up to 3000 meters (9840 feet) in length. The length of the fiber-optic link joining the two remote bridges can be up to 3000 meters (9840 feet) in length.


The fiber-optic link has all of the characteristics of a LAN except that it has no stations and its length cannot exceed 3000 meters (9840 feet). As such, the fiber-optic link must be included in the path cost of the extended LAN (more information on path cost calculation is provided in Chapter 3). The extended LAN example shown in Figure 1-10 is drawn to show that the fiber-optic link between the remote bridges is equivalent to a LAN.



LKG-0731-87

Figure 1-10: Example Showing a Fiber-Optic Link as a LAN

Figure 1-11 shows some of the possible extended LAN configurations using remote bridges.



LKG-0856-87

Figure 1-11: Extended LANs Using Remote Bridges

1.3.4.1 Fiber-Optic Cable Between LAN Bridge 100 Units — In a bridge-to-bridge configuration, the dual-cable fiber-optic link that connects the bridges does not affect the cable configuration guidelines of either of the LANs connected to the bridges. The end-to-end light loss of the cable must not exceed 12.5 decibels (dB) for the DEBET-RC/RD and 16 decibels (dB) for the DEBET-RH/RJ versions, regardless of the length. The type of optical fiber used affects the length of the cable. The loss budget is based on Corning 1508, 100/140 optical fiber with a bandwidth of 300 MHz/km measured at 820 nanometers.

CAUTION

Exceeding the loss limit may cause the bridge configuration to fail (see Appendix B for information on attenuation budget allocation).

To achieve long distances, particularly those more than 1000 meters (3280 feet), installation of the fiber-optic cable must be carefully planned. The type and quality of the cable's optical fiber, the cable repair strategy, and the cable's total end-to-end light loss are of great importance to successful bridge installation.

The end-to-end cable light loss depends on the quality of the fiber, the number and quality of splices required for installation, and the number and quality of the connectors used.


The cable repair strategy affects the budget in that cable repair typically consists of replacing a section of cable. This requires two splices. The repaired link must remain under the end-to-end light loss budget. If the initial installation uses the entire loss budget, a repair would not be possible. Therefore, plan for a minimum of two splices (about 0.5 dB for each splice).

For longer cable runs or for installations requiring more splices, request a lower-loss fiber-optic cable from your cable vendor. For information on measuring light loss through a fiber-optic link, refer to Appendix B.

1.3.4.2 LAN Bridge 100 to Repeater Considerations — When a remote bridge is connected to a remote repeater, the fiber-optic cable can be up to 1500 meters (4920 feet) long. When a remote bridge is connected to a remote repeater, the length of the fiber-optic cable is considered part of the length of that LAN (see Figure 1–12). Therefore, the distance between the bridge and the farthest end station on the LAN cannot exceed 2800 meters (9184 feet).

The 2800-meter (9184-foot) distance between any two stations on a typical Ethernet LAN is made up of the following:

- Two 50-meter (164-foot) transceiver cables (connecting the farthest end stations)
- Three 500-meter (1640 foot) coaxial cable segments
- Four 50-meter (164 foot) transceiver cables (connected to repeaters)
- 1000 meters (3280 feet) of point-to-point fiber-optic link



LKQ-0857-87
PRELIMINARY

Figure 1-12: Remote Bridge to Repeater Configuration

1.4 LAN Traffic Monitor Configurations

The following sections describe several ways for configuring the LAN Bridge 100 hardware unit as an LTM Listener unit. For more information about the LAN Traffic Monitor, refer to the *LAN Traffic Monitor User's Guide*.

1.4.1 Single Port Configuration with Loopback Connector Installed

As shown in Figure 1-13, the LTM Listener always monitors Ethernet 2 and sends statistics to the user interface on Ethernet 2. The LTM Listener can send statistics to a user interface on Ethernet 1 also, as long as the LAN Bridge 100 unit connects the two Ethernets. Note that port B has a loopback connector installed and is not in operation.

NOTE

The LAN Bridge 100 unit fails self-test if an unused port is left disconnected unless the unused port is the fiber-optic port used with the remote version (DEBET-RC/RD). An Ethernet loopback connector (shipped with the unit) must be connected to the unused transceiver port. Do NOT install a fiber-optic loopback connector to an unused fiber-optic port. A fiber-optic loopback causes the bridge to fail the power-up self-test.




Figure 1-13: LTM Single Port Configuration

1.4.2 Dual Port Connections Between Two Ethernets

As shown in Figure 1-14, the LTM Listener is connected to two completely separate Ethernets. In this case, the LAN Traffic Monitor can monitor either Ethernet 1 or 2 but must report to the LTM user interface on Ethernet 1.




Figure 1-14: LTM Connected to Two Separate Ethernets

1.4.3 Dual Port Connections With Bridged Ethernets

The configuration shown in Figure 1-15 describes two Ethernet LANs bridged together, forming a single extended LAN. The LTM Listener can monitor either Ethernet 1 or 2 and can report to either port.




Figure 1-15: LTM on Two Connected Ethernets

NOTE

Digital recommends that you configure the LTM Listener to report on the port that has the least number of intervening bridges between it and the LTM user interface host(s). Doing so minimizes the impact of a possible bridge failure.

1.5 Bridge/Router Considerations

This section discusses some considerations for configuring extended LANs that include LAN Bridge 100 units and routers.

NOTE

The information in this section applies specifically to equipment manufactured by Digital Equipment Corporation. However, all routers and bridges designed for layered architecture networks adhere to similar operating principles and may have similar parameters that should be considered when implementing networks with bridges and routers.

The configuration shown in Figure 1-16 involves a LAN Bridge 100 unit that is installed in parallel with a router. This configuration creates many problems and therefore is not recommended (see Section 1.5.1). The configuration shown in Figure 1-17 is a general case in which two local area networks are connected by a LAN Bridge 100 unit. This configuration does not have associated problems, provided that routing parameters are set properly in all routers.

The information contained in this section assumes that the reader is familiar with the concept of routing and how routers operate on LANs. For more information on the subject of routing, refer to the following documents:

- *DECnet Digital Network Architecture Phase IV: Routing Layer Functional Specification* (Order no. AA-X435A-TK)
- *Guide to Networking on VAX/VMS* (Order no. AA-Y512A-TE)
- *Ethernet Communications Server DECnet Router Software Installation Guide* (Order no. AA-X019B-TK)

The LAN Bridge 100 unit operates at the Data Link layer and is transparent to higher protocol layers in the Digital Network Architecture (DNA) model (such as Routing, End Communication, Session, and so on).

Since the LAN Bridge 100 unit connects two local area networks to form an extended LAN, higher protocol layers effectively see a single LAN. For the Routing layer, the creation of an extended LAN means either the addition of nodes (stations) and/or routers to the existing network or the actual merging of two distinct networks.


NOTE

An Ethernet repeater and LAN Bridge 100 unit are similar since both devices provide a channel through which data can pass (a repeater connects segments together to form a LAN, while a bridge connects LANs together to form an extended LAN). In addition, both devices are transparent to the higher protocol layers in the DNA model. Thus both devices introduce similar symptoms if implemented incorrectly.

1.5.1 LAN Bridge 100 Unit in Parallel with a Router

The configuration in which a LAN Bridge 100 unit is installed in parallel with a router (see Figure 1-16) does not enhance network performance. Since the LAN Bridge 100 unit can handle all of the traffic on both LANs, there is no need to have a router in parallel with the bridge to “assist” it.

In fact, in this configuration, both router ports are connected to the same extended LAN. This causes problems for the router since the router may receive two copies of the packet. Consequently, if the router is a designated router, it will hear its own Hello message on the other link. In response, the router defers to the “echo” Hello message and ceases to be the designated router for the LAN. This has a major impact on the network because end nodes on the LAN will not be able to communicate with nodes that are not on the LAN.



LKG-0734

Figure 1-16: Bridge in Parallel with a Router

There are additional problems associated with having routers receive multiple copies of packets on a LAN. However, it is not necessary to describe them all here. The network manager should adhere to the following guideline:


- A router should never have two ports connected to the same LAN or extended LAN. That is, a router should never be in parallel with a bridge or a repeater.

1.5.2 Modifying Router Parameters for Extended LANs

This section identifies router parameters that must be set to reflect the Extended LAN. The parameters may need to be modified when two LANs are joined with a LAN Bridge 100 unit (see Figure 1-17).

Each of the two LANs shown in Figure 1-17 has many nodes, some of which are broadcast routers (levels 1 and 2), while the remaining nodes are broadcast nonrouters (end nodes). (For a brief description of level 1 and level 2 routing, see Section 1.5.3.) When the bridge is installed, the network grows larger. Consequently, the following parameters may have to be modified in all routers in the extended LAN:

- MAXIMUM BROADCAST NONROUTERS
- MAXIMUM BROADCAST ROUTERS



LKG-0735

Figure 1-17: An Extended LAN with Routers

The **MAXIMUM BROADCAST NONROUTERS** parameter specifies the maximum number of Ethernet nonrouting nodes (end nodes) that the router database can contain at any one time. This parameter can range from 0 to 1022 and has a typical default value of 64.

The **MAXIMUM BROADCAST NONROUTERS** parameter must be greater than or equal to the actual number of nonrouting nodes on the LAN. The larger this number is, the greater the memory overhead for the router. However, if this number is less than the number of active end nodes, some end nodes will be unreachable because of limits on the size of the router's database.

NOTE

When nodes on a LAN are configured into multiple DECnet areas, this parameter is set on a per-area basis. For more information about areas, see Section 1.5.3. Also refer to the *Ethernet Communications Server DECnet Router Software Installation Guide*.

If the **MAXIMUM BROADCAST NONROUTERS** parameter is set too low, it is difficult to diagnose on networks with many nodes (both routing nodes and end nodes). One of the symptoms of this parameter being set too low is partial partitioning of the network. Under this condition, some nodes may lose their ability to initiate conversations with other nodes even though they are on the same LAN.

Therefore, whenever network changes are made, the network manager should take steps to verify (and, if necessary, modify) this parameter in all routers so that it is either equal to or greater than the number of nonrouting nodes in the extended LAN.

The **MAXIMUM BROADCAST ROUTERS** parameter specifies the maximum number of Ethernet routing nodes that the router database can contain at any one time (this applies to routing nodes on the same Ethernet). Routing nodes include level 1 and level 2 routers (see Section 1.5.3). In a multiple area network, this parameter must include all the level 1 routers in a particular area plus all the level 2 routers in the entire network. This parameter can range from 0 to 32 and has a typical default value of 10.

The larger the value for this parameter, the greater the memory overhead for the router. Also the more routers there are on the Ethernet, the greater is the control traffic associated with the temporary looping property of the routing algorithm.

If this parameter value is too small, some routers will be unreachable because of limits on the size of the router's database. That is, when the **MAXIMUM BROADCAST ROUTERS** parameter is exceeded, routers with the lowest priority are dropped from the router's database.

Therefore, whenever network changes are made, the network manager should take steps to verify (and, if necessary, modify) this parameter so that it is either equal to or greater than the number of level 1 routers in the DECnet area plus the number of level 2 routers in the extended LAN.

1.5.3 Setting Up Multiple Areas

When the total number of routers and/or nonrouters on the extended LAN approach the maximum values, the network manager can set up multiple areas on the Ethernet.

In general, each area is a group of nodes. Nodes are grouped together in areas for purposes of hierarchical routing. Hierarchical routing involves the addition of a second level of routing to the network. Routing within an area is referred to as level 1 routing; routing between areas is called level 2 routing. When creating multiple areas on the Ethernet, the number of routers and nonrouters are considered separately for each area.

For more information on configuring multiple area LANs with routers, refer to *DECnet Digital Network Architecture Phase IV: Routing Layer Functional Specification*.

1.6 Specifications

Specifications for the LAN Bridge 100 unit are divided into the following categories:

- Physical specifications
- Environmental specifications
- Electrical specifications

1.6.1 Physical Specifications

The LAN Bridge 100 unit is housed in a plastic table-top enclosure that has the following dimensions and weight.

With plastic enclosure:

Revision F08 and below	Revision F09 and above
Height 16.2 cm (6.4 in)	Height 16.2 cm (6.4 in)
Width 49.4 cm (19.4 in)	Width 49.4 cm (19.4 in)
Depth 31.3 cm (12.3 in)	Depth 31.3 cm (12.3 in)
Weight 9.5 kg (21 lb)	Weight 6.7 kg (15 lb)

Without plastic enclosure

Revision F08 and below	Revision F09 and above
Height 13.3 cm (5.3 in)	Height 13.3 cm (5.3 in)
Width 43.6 cm (17.2 in)	Width 43.6 cm (17.2 in)
Depth 29.8 cm (11.7 in)	Depth 29.8 cm (11.7 in)
Weight 7.3 kg (16 lb)	Weight 4.5 kg (10 lb)

AC power-cord length	U.S. 1.83 m (6 ft) Others 2.5 m (8.2 ft)
----------------------	---

Transceiver cable length	BNE4 12 m (39 ft) maximum BNE3 50 m (164 ft) maximum
--------------------------	---

Fiber-optic cable

Type	Corning 1508 (100/140)
Minimum bandwidth	300 MHz measured at a wavelength of 820 nm

Fiber-optic connectors

Type	Stainless steel, Amphenol type 906, SMA style, or equivalent
Maximum attenuation	Less than 1.5 dB

The plastic enclosure is easily removed, and brackets are provided to allow mounting of the unit on a wall or in a cabinet. An optional kit (part number H039) is available for mounting the bridge on a wall or partition without removing the plastic enclosure.

1.6.2 Environmental Specifications

The LAN Bridge 100 unit is designed to operate in a non-air-conditioned environment or in an exposed area of an industrial site. However, 50°C is the maximum ambient temperature which must not be exceeded at the air intake of the bridge. This applies even when the LAN Bridge 100 unit is mounted in a cabinet. The bridge is not intended to operate in an air plenum.

Operating Environment:

Temperature	5°C to 50°C (41°F to 122°F)
Maximum rate of change	20°C/hr (36°F/hr)
Relative humidity	10% to 95% (noncondensing)
Wet-bulb temperature	32°C (90°F) maximum
Dew point	2°C (36°F) minimum
Altitude	Sea level to 2.4 km (8000 ft)
Air flow	37.5 CFM. About 10 to 15 cm (4 to 6 in) of space must be provided on both ends of the unit for adequate air flow.

Shipping Environment:

Temperature	-40°C to 66°C (-40°F to 151°F)
Relative humidity	0% to 95% (noncondensing)
Altitude	Sea level to 9.1 km (30,000 ft)

1.6.3 Electrical Specifications

The LAN Bridge 100 unit requires the following power for proper operation:

AC input power

	Revision F08 and below	Revision F09 and above
Voltage	Range	Range
120 V	88 V to 132 V	88 V to 132 V
240 V	176 V to 264 V	176 V to 264 V
Current		
At 120 V	1.6 amps	1.1 amps
At 240 V	0.9 amps	0.7 amps
Frequency	47 to 63 Hz	47 to 63 Hz
Power Consumption	160 W	168 W

The LAN Bridge 100 unit provides the following power features:

- Self-contained power supply
- Power-cord options for all major DIGITAL markets
- Adequate power to drive two external transceivers

2

Installation and Verification


2.1 Overview

This chapter provides all the information necessary for successful installation and subsequent checkout of the LAN Bridge 100 unit. Included are instructions for:

- Unpacking and inspection
- Step-by-step installation
- Verification of operation

2.2 Unpacking

Unpack the LAN Bridge 100 unit (Figure 2-1).



*BUBBLE PACKING CONTAINS LOOPBACK CONNECTORS, MOUNTING BRACKETS, AND SCREWS

LKG-0953

Figure 2-1: Unpacking the LAN Bridge 100 Unit

Check the contents of the carton for damage and missing parts (see Figure 2-2). In case of damage, contact your delivery agent and your Digital sales representative. In case of missing parts, contact your Digital sales representative.





Figure 2-2: Checking the LAN Bridge 100 Contents

2.3 Arranging for Software Installation

If you are installing the LAN Bridge 100 unit as a bridge, software installation is not required. Go to Section 2.4.

If you are installing the LAN Bridge 100 unit to operate as an LTM Listener, arrange to have the LTM distribution software installed on a load host while you install the hardware.

The identification card (ID card) that you received with the unit provides information that is required by the person installing the software. This information includes the LAN Bridge 100 serial number and Ethernet address (see Figure 2-3), where the device is to be located, and your name. Fill in this information at the top portion of the ID card, and then give the card and any software cartons you received with the shipment to the system/network manager. The system/network manager can then fill in the required information as described on the ID card and arrange for the software installation. Ask to be notified when the LTM distribution software is installed on a load host and when the LAN Bridge 100 unit is configured in the load host database. You can continue with the hardware portion of the installation, but do NOT power up the unit until notified that software installation is complete and that the LAN Bridge 100 unit is configured in the load host database.



LKG-0961

Figure 2-3: Location of Serial Number and Ethernet Address

2.4 Installation


Installation consists of setting switches, placing the unit and connecting cables.

2.4.1 Setting the Voltage Switch

Determine the operating voltage for the LAN Bridge 100 unit by comparing your power cord to those shown in Figure 2-4. If necessary, set the voltage switch to match the power-cord voltage. Do not connect the power cord at this time.

CAUTION

An incorrect voltage setting can damage the LAN Bridge 100 unit.



LKG-0931


Figure 2-4: Setting the Voltage Select Switch

2.4.2 Verifying Configuration Switch Settings

Check to see that your configuration switch settings are properly set, as shown in Figure 2-5.

NOTE

If you are installing the LAN Bridge 100 unit as only an LTM Listener, set Switch 5, Down-Line Load Enable, to ON (down = ON). You can use RBMS to switch between LTM functions and bridge operation if Switch 5 is OFF (up = OFF).



LKG-0959

Figure 2-5: Verifying Configuration Switch Settings


The LAN Bridge 100 unit is housed in a plastic, table-top enclosure. This plastic enclosure can be removed, and the sheet metal unit inside it can be either wall mounted or mounted in a standard 48-cm (19-inch) rack. A set of mounting brackets for both rack mounting and wall mounting is included.

NOTE

Whichever installation you choose, leave enough clearance around the bridge's air inlets and outlets to ensure optimal air flow.

Rack mounting requires removal of the bridge's plastic enclosure. For wall mounting, the bridge can be installed with or without its plastic enclosure. To mount the LAN Bridge 100 unit with the enclosure on a wall or partition, order the special wall-mounting kit (part no. H039). To remove the plastic enclosure, take the following steps (see Figure 2-6).

1. Turn the bridge upside down.
2. Remove the 8 screws from the bottom of the unit.
3. Remove the plastic enclosure.



LKG-0952

Figure 2-6: Removing the Plastic Enclosure from the LAN Bridge 100 Unit


CAUTION

Do not reinstall screws in the bridge's metal casing. Doing so could damage the bridge.

2.4.3 Rack-Mount Installation (Optional)

To rack mount the LAN Bridge 100 unit, take the following steps (see Figure 2–7).

1. Fasten the mounting brackets to the bridge with the screws provided.
2. Fasten the bridge to the rack (screws not provided).



LKG-0932

Figure 2–7: Installing the LAN Bridge 100 Unit in a Rack

NOTE

The air inside the rack may be hotter than the ambient room temperature. Therefore, be sure that the air entering the bridge's air inlet does not exceed the maximum air-inlet temperature of 50°C (122°F).

In addition, be sure that all cables connecting to a rack-mounted bridge are dressed to the rack to relieve strain. If fiber-optic cables are used, the cables must not exceed their 15-cm (6-inch) bend radius.


2.4.4 Wall-Mount Installation (Optional)

To wall mount the LAN Bridge 100 unit without the plastic enclosure, take the following steps:

1. Fasten the mounting brackets to the bridge with the screws provided (Figure 2–8).

NOTE

To meet federal safety codes, never install a bridge with the I/O panel facing downward.




LKG-0938

Figure 2–8: Fastening the Wall-Mounting Brackets to the LAN Bridge 100 Unit

2. Fasten the brackets to the wall (screws not provided) (Figure 2-9).

NOTE

The air entering the bridge's air inlet must not exceed the bridge's maximum temperature of 50°C (122°F). If two bridges are mounted side by side, one bridge receives the heated exhaust air from the other. This usually does not cause a problem unless the ambient air temperature is close to 50°C (122°F). Do not mount several bridges side by side.




LKG-0954

Figure 2-9: Positioning the LAN Bridge 100 Unit for Wall Mounting

2.4.5 Connecting the Transceiver Cables

If you are installing a DEBET-RC/RH or DEBET-RD/RJ bridge with a fiber-optic cable, there will be only one transceiver cable (see Figure 2-10).

1. Push the slide latch to the left to the unlocked position.
2. Plug the transceiver cable into the jack.
3. Push the slide latch to the right until it snaps into the locking position.
4. Gently pull on each connector to make sure that the latch is secure.



LKG-0933

Figure 2-10: Connecting the Transceiver Cable


2.4.6 Connecting the Fiber-Optic Cable

If you are installing a DEBET-RC/RH or DEBET-RD/RJ bridge with a fiber-optic cable, use the following procedure to connect it (see Figure 2-11):

WARNING

Never look into a fiber-optic connector or cable. High-intensity light can damage your eyes.

1. Pull the protective caps off the fiber-optic connectors (Figure 2-11).




LKG-0934

Figure 2-11: Removing the Protective Caps from the Fiber-Optic Connectors

2. Note the polarity of the cable connections. The transmit cable (identified on the cable by an arrow pointing away from the connector) must be connected to the fiber-optic output of the bridge. The receive cable (identified on the cable by an arrow pointing toward the connector) must be connected to the fiber-optic input of the bridge.
3. Connect the fiber-optic cable. Tighten the fiber-optic connectors so they are only finger tight (see Figure 2-12).

CAUTION

The fiber-optic cable will be damaged by sharp bends. Do not allow the bend radius to be less than 15 cm (6 inches).



LKG-0955

Figure 2-12: Connecting the Fiber-Optic Cable

2.4.7 Connecting the Power Cord

To connect the power cord, take the following steps (Figure 2–13):

1. Read and remove the yellow CAUTION label.
2. Plug in both ends of the power cord.




Figure 2–13: Connecting the Power Cord

2.5 Verifying the Installation

Proper installation of the LAN Bridge 100 unit is verified by the state of the status LEDs on the I/O panel. The LED states vary depending whether the LAN Bridge 100 hardware unit is configured to operate as a bridge or as a LAN Traffic Monitor.


If you are installing the LAN Bridge 100 unit to operate as a bridge, go to Section 2.5.1 to verify correct installation.

If you are installing the LAN Bridge 100 unit to operate as a LAN Traffic Monitor, go to Section 2.5.2 to verify correct installation.

2.5.1 Verifying the Bridge Installation

Whenever power is applied to the unit, the bridge performs its diagnostic self-test. The bridge's self-test normally takes about 15 seconds to complete. The bridge then spends about 30 seconds communicating with other bridges in the network to determine if there are loops and build its initial database of forwarding addresses.

Allow up to 45 seconds for the bridge's self-test and communications tasks to complete, then compare the state of the five status LEDs on the bridge with those shown in Figure 2-14.



If:
The status LEDs match those shown here.

The status LEDs **DO NOT** match those shown here.

Then:
The LAN Bridge 100 hardware is functional.

Go to Chapter 5.

LKG-0957

Figure 2-14: Bridge Hardware Verification

NOTE

For definitions of the status LEDs on the bridge, refer to Chapter 1, Table 1-5.


Check the logical link by sending a message from a node on one side of the bridge to a node on the other side. On a VAX/VMS host running DECnet, you can send a message by using the SET HOST command.

Installation of the LAN Bridge 100 unit is now complete.

2.5.2 Verifying LAN Traffic Monitor (LTM) Installation

Whenever power is applied to the unit, the bridge performs its diagnostic self-test and, if successful, initiates a request for a down-line load of the LTM Listener software image from a load host. The bridge's self-test normally takes about 15 seconds to complete, but the down-line loading of the software image could take longer if the network is busy.

Allow up to 2 minutes for the bridge's self-test and down-line loading of the LTM Listener software image to complete; then compare the state of the five status LEDs on the unit with those shown in Figure 2-15.



If:

The status LEDs match those shown here.

The status LEDs **DO NOT** match those shown here.

Then:

The LTM Listener is operational. Notify the system manager.

Go to Chapter 5.

LKG-0956

Figure 2-15: LTM Hardware Verification

NOTE

For definitions of the status LEDs on the bridge, refer to Chapter 1, Table 1-5.

Installation of the LAN Bridge 100 unit as a LAN Traffic Monitor is now complete.

3 Operation


3.1 Chapter Overview

This chapter describes the operational states and processes control operation of the LAN Bridge 100 unit. This chapter also describes the initialization block and descriptor ring data structures.

3.2 Operational States

The LAN Bridge 100 unit has six states that control its functions. An operating state transition diagram is shown in Figure 3–1. The LAN Bridge 100 unit has the following operational states:


- SELF-TEST
- BROKEN
- INITIALIZATION
- PREFORWARDING
- FORWARDING
- BACKUP



LKG-0737-87

Figure 3-1


Figure 3-1: Operating State Transition (Sheet 1 of 3)



* DLL STANDS FOR DOWN-LINE LOAD

LKG-0737-87

Figure 3-1: Operating State Transition (Sheet 2 of 3)



LKG-0737

Figure 3-1: Operating State Transition (Sheet 3 of 3)

3.2.1 SELF-TEST State

The SELF-TEST state is entered on power up, on a soft reset command received from bridge management software, or from the BROKEN state.

The SELF-TEST state sets the LAN Bridge 100 unit to a known condition and tests the hardware and firmware by performing read/write operations to various hardware locations.

The SELF-TEST state loops (repeats) if the Loop Self-test switch (located on the LAN Bridge 100 connector panel) is set to the on position. A description of the self-test procedure is provided in Section 3.9.

If the self-test fails (except for nonvolatile RAM [NVRAM] failures), the LAN Bridge 100 unit enters the BROKEN state.

If NVRAM fails and the NVRAM Reset switch (located on the I/O panel) is enabled, the LAN Bridge 100 unit will enter the INITIALIZATION state using default parameters. If NVRAM fails and the NVRAM Reset switch is disabled, the LAN Bridge 100 unit will enter the BROKEN state.

If the self-test passes, the LAN Bridge 100 unit enters the INITIALIZATION state.

3.2.2 BROKEN State

The BROKEN state is entered if errors occur in self-test. In the BROKEN state, the LAN Bridge 100 unit reenters the SELF-TEST state every 15 seconds until the LAN Bridge 100 unit is powered off.

3.2.3 INITIALIZATION State

The INITIALIZATION state is entered on completion of the self-test.

3.2.3.1 Bridge Initialization — The microprocessor uses the INITIALIZATION state to perform the following operations:

- Initialize various system counters and registers.
- Allocate necessary buffer space in packet memory.
- Write the initialization blocks and the descriptor rings into packet memory (see Section 3.3).
- Write the LAN Bridge 100 unit's physical address into the Ethernet address table and into the LANCES.
- Write the pointers to the initialization block into the LANCE control and status registers.
- Determine whether the bridge down-line loads as part of the INITIALIZATION state.
- Write the initialization and start bits into the LANCE control and status registers.

If the bridge determines that a down-line load will not occur, the INITIALIZATION state exits to the PREFORWARDING state after the initialization and start bits are written into the LANCE control and status registers.

3.2.3.2 Down-Line Loading — The LAN Bridge 100 unit may perform a down-line load as part of the INITIALIZATION state.

The down-line load feature is controlled using a hardware “down-line load enable” switch and two software parameters: DLL-switch and DLL-info (DLL stands for down-line load). The hardware switch must be manually set or cleared. Software parameter values for DLL-switch and DLL-info are entered using RBMS.

An algorithm examines the hardware switch and software parameters to determine whether the bridge down-line loads. If the bridge is configured to down-line load, the algorithm specifies which file to load. A block diagram of the algorithm is contained in Figure 3-1.

MOP protocol messages control the down-line load operation. Once the down-line load operation begins, the bridge attempts to complete the operation until the load is successful or the bridge power is turned off.

When the down-line load is complete, the INITIALIZATION state exits to the state determined by the loaded software.

3.2.4 PREFORWARDING State

The PREFORWARDING state is entered after the initialize and start bits are written into the LANCE control and status registers. This state may also be entered from the BACKUP state when the LAN Bridge 100 unit becomes the root bridge in the extended LAN or a designated bridge on a LAN (see Section 3.5).

The LAN Bridge 100 unit uses the PREFORWARDING state to:

- Learn the location of active stations on the network.
- Perform loop detection functions to determine whether it is the root bridge in the extended LAN or a designated bridge on a LAN (see Section 3.5).

The operations performed in the PREFORWARDING and FORWARDING states (see Section 3.2.5) are identical except that packets are not forwarded during the PREFORWARDING state. This prevents the formation of transient loops.

The duration (forwarding delay) of the PREFORWARDING state has a default value of 30 seconds. The forwarding delay can be changed using RBMS (see the *Remote Bridge Management Software Guide*).

The PREFORWARDING state exits to the BACKUP state if the LAN Bridge 100 unit determines that the following loop detection conditions (see Section 3.5) are met:

- The LAN Bridge 100 unit is not a root bridge.
- The bridge is not a designated bridge on a LAN.

If the loop detection conditions are not met, the PREFORWARDING state exits to the FORWARDING state.

3.2.5 FORWARDING State

The FORWARDING state is entered from the PREFORWARDING state when loop detection conditions are not met (refer to Sections 3.2.4 and 3.5).

The FORWARDING state is similar to the PREFORWARDING state except that packets may be forwarded. In this state, the LAN Bridge 100 unit performs address filtering functions. The LANCE chips receive and store packets in memory. The packets are discarded or processed depending on a determination made by the microprocessor and table lookup subsystem.

Other tasks carried out in the FORWARDING state include address table maintenance and loop detection functions. Bridge management functions such as Remote Bridge Management Software (RBMS) and maintenance functions which use Maintenance Operation Protocol (MOP) may also be performed.

The FORWARDING state exits to the SELF-TEST state if a reset operation is performed.

The FORWARDING state exits to the BACKUP state if loop detection conditions are met (refer to Section 3.5).

3.2.6 BACKUP State

The BACKUP state is entered when loop detection conditions are met (see Section 3.5). Remote Bridge Management Software (RBMS) commands may be used to force the LAN Bridge 100 unit into the BACKUP state by “disabling” a link in the bridge.

During the BACKUP state, the LAN Bridge 100 unit “listens” to network traffic. If the LAN Bridge 100 unit determines that loop detection conditions are not met, it enters the PREFORWARDING state.

3.3 Packet Memory Data Structures

The LANCE initialization blocks and the descriptor rings are the data structures that manage packet memory.

The LANCE chips use these structures to locate available buffer space for storing incoming packets and for finding packets that are to be transmitted.

The microprocessor writes the initialization blocks and descriptor rings into packet memory during the initialization state. The microprocessor uses the descriptor rings to find packets in memory and to identify packets that are ready for transmission by a LANCE.

3.3.1 Initialization Blocks

The microprocessor writes two initialization blocks into packet memory to initialize the two LANCE chips. Each block uses 12 words of contiguous memory.

Each LANCE uses its own initialization block to configure its operating parameters. The fields in an initialization block are identified in Figure 3-2.

15	00
MODE OF LANCE OPERATION	
ETHERNET PHYSICAL ADDRESS	
ETHERNET PHYSICAL ADDRESS	
ETHERNET PHYSICAL ADDRESS	
LOGICAL ADDRESS FILTER	
RECEIVE DESCRIPTOR RING POINTER	
RECEIVE DESCRIPTOR RING POINTER	
TRANSMIT DESCRIPTOR RING POINTER	
TRANSMIT DESCRIPTOR RING POINTER	

LKG-1224-87

Figure 3-2: LANCE Initialization Block Field Functions

The functions of each field in the initialization block are defined as follows:

- Mode of LANCE operation — Promiscuous mode is selected, signifying that all addresses are received.
- Ethernet physical address — This field contains the 48-bit Ethernet address assigned to the LAN Bridge 100 unit. The LANCE uses this address to detect packets (such as RBMS or MOP commands) addressed to the LAN Bridge 100 unit. This address is also transmitted by the LAN Bridge 100 unit in loop detection (Hello) messages.
- Logical address filter — The 64-bit filter helps to detect multicast addresses. This filter is not used by the LAN Bridge 100 unit since the LANCE operates in promiscuous mode and receives all packets.
- Receive descriptor ring pointer — This field points to the base (lowest) address of the receive descriptor ring. Another part of this field defines the number of entries in the receive descriptor ring (1, 2, 4, 16, 32, 64, or 128 entries).

- **Transmit descriptor ring pointer** — This field points to the base (lowest) address of the transmit descriptor ring. Another part of this field defines the number of entries in the transmit descriptor ring (1, 2, 4, 16, 32, 64, or 128 entries).

3.3.2 Descriptor Rings


Four descriptor rings (one transmit and one receive ring for each LANCE) are written into packet memory by the processor during the INITIALIZATION state.

A descriptor ring is a circular queue of tasks which the LANCE and processor use to locate buffer space or packets in packet memory. Figure 3–3 illustrates the concept of descriptor rings.

Each descriptor ring is made up of entries that define the location and status of buffer space in packet memory. Each entry is four words in length. Thus each entry starts on a quadword boundary in the descriptor ring address space. The number of entries in a descriptor ring is defined in the ring pointer of the initialization block.

Receive descriptor rings are used by the LANCE chips to find available buffer space for writing incoming packets. The processor uses the receive descriptor rings to find the destination and source addresses in packets the LANCE chips receive (the destination and source addresses are contained in the first 12 bytes of a packet).


Transmit descriptor rings are written to by the processor to identify packets awaiting transmission. The LANCE reads the transmit descriptor rings to find packets awaiting transmission.



LKG-0739

Figure 3-3: Descriptor Ring Structures

3.3.2.1 Receive Descriptor Ring Entry — Each receive descriptor ring entry is made up of four words that define the location and status of buffer space in packet memory. Figures 3–4 through 3–7 show the fields that make up each word in the receive descriptor ring entry. Tables 3–1 through 3–4 describe the bit functions of each word in the receive descriptor ring entry.




LKG-0740

Figure 3–4: Bit Format of Word 0 in the Receive Descriptor Ring Entry

Table 3–1: Bit Descriptions of Word 0 in the Receive Descriptor Ring Entry

Bit(s)	Name	Description
< 15:00 >	LADR	Low address. The low-order 16 bits of the address of the transmit data buffer corresponding to this receive descriptor ring entry.




LKG-0742

Figure 3-5: Bit Format of Word 1 in the Receive Descriptor Ring Entry

Table 3-2: Bit Descriptions of Word 1 in the Receive Descriptor Ring Entry

Bit(s)	Name	Description
<15>	OWN	Ownership bit. When set, indicates that the LANCE owns or has use of the receive descriptor ring entry and the corresponding receive data buffer. The LANCE clears the OWN bit to give the entry to the microprocessor. The microprocessor sets the bit to give the entry back to the LANCE.
<14>	ERR	Error summary. Indicates that one or more of the following error bits is set: FRAM, OFLO, CRC, or BUFF.
<13>	FRAM	Framing error. Indicates that the received packet has both a CRC error and that the number of bits in the packet is not a whole multiple of 8.
<12>	OFLO	Overflow. Indicates that the LANCE has lost part or all of the packet because the internal data silo overflowed before the LANCE could write the packet into packet memory.
<11>	CRC	CRC error. Indicates that the packet has a CRC error.
<10>	BUFF	Buffer error. Indicates that the LANCE filled the receive data buffer and was unable to chain to another receive data buffer to store the remainder of the packet. Either there was not another receive buffer available or the LANCE was unable to read the receive descriptor ring entry quickly enough to find the new buffer before the LANCE internal data silo overflowed. Data buffers must be at least 64 bytes long to allow the LANCE time to find the next data buffer for buffer chaining.
<09>	STP	Start of packet. Indicates that this receive data buffer starts a new packet.
<08>	ENP	End of packet. Indicates that this receive data buffer ends a packet.
<07:00>	HADR	High address. The high-order 8 bits of the address of the receive data buffer corresponding to this receive descriptor ring entry.




LKG-0743

Figure 3-6: Bit Format of Word 2 in the Receive Descriptor Ring Entry

Table 3-3: Bit Descriptions of Word 2 in the Receive Descriptor Ring Entry

Bit(s)	Name	Description
<15:12>	1	Must be ones.
<11:00>	BCNT	Buffer byte count. The length of the data buffer corresponding to this entry in the receive descriptor ring expressed as a twos complement.




LKG-0744

Figure 3-7: Bit Format of Word 3 in the Receive Descriptor Ring Entry

Table 3-4: Bit Descriptions of Word 3 in the Receive Descriptor Ring Entry

Bit(s)	Name	Description
<15:12>	RESERVED	Not used.
<11:00>	PCNT	Packet byte count. The number of bytes in the received packet.

3.3.2.2 Transmit Descriptor Ring Entry — Each transmit descriptor ring entry is made up of four words that define the location and status of packets stored in packet memory. Figures 3–8 through 3–11 show the fields that make up each word in the transmit descriptor ring entry. Tables 3–5 through 3–8 describe the bit functions of each word in the transmit descriptor ring entry.




LKG-0745

Figure 3–8: Bit Format of Word 0 in the Transmit Descriptor Ring Entry

Table 3–5: Bit Descriptions of Word 0 in the Transmit Descriptor Ring Entry

Bit(s)	Name	Description
<15:00>	LADR	Low address. The low-order 16 bits of the address of the transmit data buffer corresponding to this transmit descriptor ring entry.




LKG-0746

Figure 3-9: Bit Format of Word 1 in the Transmit Descriptor Ring Entry

Table 3-6: Bit Descriptions of Word 1 in the Transmit Descriptor Ring Entry

Bit(s)	Name	Description
< 15 >	OWN	Ownership bit. When set, indicates that the LANCE owns or has use of the transmit descriptor ring entry and the corresponding transmit data buffer. The microprocessor sets the OWN bit to tell the LANCE to transmit the contents of the data buffer as a packet. The LANCE clears the bit to give the entry back to the microprocessor.
< 14 >	ERR	Error summary. Indicates that one or more of the following error bits is set in the transmit descriptor ring entry word 3: UFLO, LCOL, LCAR, or RTRY.
< 13 >	RES	Not used.
< 12 >	MORE	More. Indicates that more than one retry was needed to transmit the packet.
< 11 >	ONE	One. Indicates that exactly one retry was needed to transmit the packet.
< 10 >	DEF	Deferred. Indicates that the LANCE had to defer or wait for a pause in the message traffic on the network before transmitting the packet.
< 09 >	STP	Start of packet. Indicates that this is the first buffer to be used by the LANCE for this packet. It is used for data chaining buffers.
< 08 >	ENP	End of packet. Indicates that this is the last buffer to be used by the LANCE for this packet. It is used for data chaining buffers.
< 07:00 >	HADR	High address. The high-order 8 bits of the address of the transmit data buffer corresponding to this transmit descriptor ring entry.




LKG-0747

Figure 3-10: Bit Format of Word 2 in the Transmit Descriptor Ring Entry

Table 3-7: Bit Descriptions of Word 2 in the Transmit Descriptor Ring Entry

Bit(s)	Name	Description
<15:12>	1	Must be ones.
<11:00>	BCNT	Buffer byte count. The length of the data buffer corresponding to this entry in the transmit descriptor ring expressed as a twos complement number.



LKG-0748

Figure 3-11: Bit Format of Word 3 in the Transmit Descriptor Ring Entry

Table 3-8: Bit Descriptions of Word 3 in the Transmit Descriptor Ring Entry

Bit(s)	Name	Description
< 15 >	BUFF	Buffer error. Indicates that the LANCE did not find the end of packet (ENP) bit in the current buffer and did not own the next buffer.
< 14 >	UFLO	Underflow error. Indicates that the LANCE has truncated the packet because it was unable to read data words from memory quickly enough.
< 13 >	RES	Reserved. Not used.
< 12 >	LCOL	Late collision. Indicates that a collision occurred after the LANCE transmitted the first 64 bytes of the packet. All nodes should have detected that there was activity on the network by this time. This suggests that the collision detect circuitry of some other node on the network has failed. The LANCE does not retry on late collisions.
< 11 >	LCAR	Loss of carrier. Indicates that the carrier-presence input to the LANCE went false during the transmission of this packet. The LANCE does not retry upon loss of carrier.
< 10 >	RTRY	Retry error. Indicates that the LANCE failed to transmit the packet in 16 attempts because of repeated collisions on the network.
< 09:00 >	TDR	Time domain reflectometry. An internal counter that counts system clocks (10 MHz) from the start of a transmission to the occurrence of a collision. This value is useful in determining the approximate distance to a fault in the network cable. The TDR value is valid only if RTRY is set.

3.4 Learning and Forwarding


The main function of the LAN Bridge 100 unit is filtering network traffic. This minimizes network traffic by keeping local traffic local.

The LAN Bridge 100 unit “listens” to network traffic and acquires a working knowledge of which stations are associated with which LAN (LANs are identified as being connected to port A or port B of the bridge). This knowledge is acquired by reading the source addresses of incoming packets and noting the port through which the packet entered. Thus each station address that is heard from becomes associated with a port of the bridge.

The LAN Bridge 100 unit stores the information in a station address table that holds the station address (48 bits) and the status of each station (16 bits). The status is written by the firmware and defines the port associated with an address, the age of the address, and so on.


When subsequent packets are received that are intended for a known (previously stored) address, the bridge can determine whether the packet should be forwarded or ignored. The bridge also updates any existing status for the station that sent the packet. If the sending station was previously unknown, the bridge adds the new address and its status to the table.

Figure 3–12 illustrates the table lookup (TLU) process. The Troll feature shown in the figure may be set by RBMS. This feature controls access to certain stations on the network (see Section 3.4.5).



LKG-0749

Figure 3-12: Table Lookup Process (Sheet 1 of 2)



LKG-0749

Figure 3-12: Table Lookup Process (Sheet 2 of 2)

3.4.1 Table Initialization

When the LAN Bridge 100 unit first becomes active, the network address table is empty. Because the table is empty, the processor has no means of accessing status information to perform packet filtering. The bridge must construct a list of active stations on the network. To do this, the bridge processor places the bridge in the PREFORWARDING state to monitor all traffic on the two LANs to which it is connected.

In the PREFORWARDING state, source addresses are extracted from the packets received from the LANs, but the packets are not forwarded. These source addresses and a status word for each station address are inserted in ascending order into the network address table.

The bridge exits the PREFORWARDING state after the preforwarding delay time of 30 seconds. The preforwarding delay time can be changed by use of RBMS.

3.4.2 Address Table Entries

Address table entries are used to keep track of the status of known network stations. A station becomes known when the LAN Bridge 100 unit receives a packet from the station.

An address table entry has the fields shown in Figure 3–13.



LKG-0750

Figure 3–13: Address Table Entry

The status is a 16-bit word stored in address table RAM. The status functions include:

- Identifying the port associated with the station address.
- Logging the age of the last packet received.
- Noting any special status that may be set by RBMS such as a nonaging or limited access address and so on.

The station address is the unique 48-bit station address of the originating station.

3.4.3 Binary Search

Under the heaviest network traffic conditions, the bridge must be able to filter or forward a continuous stream of minimum-size (64 bytes) incoming packets. Thus it is essential that the TLU subsystem quickly find entries in the active station list or determine that an entry is not present. The active station list is an ordered list (ascending numerical order) of 48-bit station addresses corresponding to active stations on the network.

A rapid search is accomplished using a binary decision tree search method. Figure 3–14 shows a binary decision tree for 16 stations.

The binary search is started at the middle entry of the table. A comparison is made between the address received and the middle entry and a “greater than,” “less than,” or “equal to” determination is made. This determination results in the following actions:


- If the received address is greater than the accessed entry of the table, the search continues at the center of the upper half of the tree. In Figure 3–14, the search would continue at address 12.
- If the received address is less than the accessed entry of the table, the search continues at the center of the lower half of the tree. In Figure 3–14, the search would continue at address 4.
- If the received address is equal to the accessed entry of the table, the search ends.

If the search continues, another “greater than,” “less than,” or “equal to” determination is made on the remaining part of the tree. This process continues until the address is found or is determined to be missing from the table.

When the search ends, the results are made available to the search results register. The results include a pointer to the last location probed, and bits that indicate whether the received address was less than, greater than, or equal to the address in the last location probed.


The binary search algorithm requires a maximum of 13 probes in a table containing 8000 entries. The example in Figure 3–14 requires a maximum of four probes. To optimize the LAN Bridge 100 processor performance, the results of a search are sought only after enough time has elapsed to guarantee that the required number of probes has been performed.

Figure 3–15 illustrates the decisions that control the binary search process.



LKG-0751

Figure 3–14: Binary Decision Tree



LKG-0752

Figure 3-15: Binary Search Process

3.4.4 Forwarding

When the LAN Bridge 100 unit has found an incoming packet's destination address to be stored in the active station list of the address table, it checks the corresponding status word to determine whether the packet should be forwarded or ignored.

When the bridge has not found the destination address from an incoming packet to be stored in the active station list of the address table, the bridge forwards the packet by default.

NOTE

If certain conditions are set by use of RBMS, the LAN Bridge 100 unit places additional constraints on the decision of whether to filter or forward packets. Some of these conditions are discussed in Section 3.4.5.

3.4.5 Writing the Ethernet Address Table

An optional software package called Remote Bridge Management Software (RBMS) can be used to specify additional conditions for the bridge's decision on whether to filter or forward packets. Users of RBMS can write into the active station list to regionalize multicast packets or limit the access of certain stations on opposite sides of the bridge.

More detailed information on implementing the following features is provided in the *Remote Bridge Management Software Guide* (part no. AA-FY93A-TE).

Regionalization

A number of protocols use multicast packets for initializing, locating other stations, and so on. It is easy to confine such protocols to a particular region of the network by instructing the bridge not to pass such multicast packets. This is done by entering the multicast addresses into the Ethernet address table of the bridge. In this way the bridge recognizes the addresses as being on the same side of the bridge as the originating station and does not forward them.

Limiting Access

It may be desirable from a network management standpoint to limit the ability of some stations on one side of the bridge to send packets to certain station on the opposite side of the bridge. This is referred to as the Troll feature, which can be activated by RBMS.

The appropriate bridge management protocol packet instructs the bridge to flag none, some, or all of the addresses in the active station list as controlled addresses. This causes the bridge firmware to drop packets originating from one side of the bridge destined for these controlled addresses on the opposite side. In other words, the bridge never forwards a packet to what is considered to be a controlled address.

3.4.6 Special Cases

The LAN Bridge 100 unit handles the following special cases.

3.4.6.1 Stale Packets — A packet that is to be forwarded becomes stale if the bridge is not able to forward the packet within two seconds because of outbound congestion. This stale packet time is measured from the reception of the last byte of the CRC off the inbound local area network (LAN) to the start of transmission of the destination address on the outbound LAN. The LAN Bridge 100 unit recovers from a stale packet condition by causing an unsuccessful transmission of the packet.

3.4.6.2 Swapped Sides — If a station is moved from one side of the bridge to the other, its presence on the new port supersedes its presence on the old port (the status word is updated).

3.4.6.3 Address Aging — The relative age of station addresses in the Ethernet address table is determined by the amount of time elapsed since the station was last heard from.

If two minutes have elapsed since the station was last heard from, the bridge considers the station to have ceased to exist on that side of the bridge. Any data in the status word indicating that the station exists on that side of the bridge is ignored.

If that station later transmits from its previous side of the bridge, its address is relearned.

NOTE

Two minutes is the normal amount of time allowed for aging. However, for two minutes following a topology change, only 30 seconds of aging are allowed. Topology changes can change the location (port) associated with an address and cause partitioning of the extended LAN. Shortened aging time minimizes partitioning by clearing the address table of entries whose locations were affected by the topology change.

3.4.7 Ethernet Address Table Maintenance

When the LAN Bridge 100 unit finds the source address from an incoming packet to be in the active station list of the Ethernet address table, it updates the corresponding status word. Updating the status word includes resetting the aging timer, noting the port associated with the address and so on.

When the source address from an incoming packet is not stored in the active station list, the bridge must insert the new address (and status) into its place in the ordered table.

New addresses and their search status (last location probed and the equal to or greater than bits) are first stored in the new address entry portion of the address table.

As a background operation (see Section 3.6), the new address entries are presorted in the new entry portion of the table. Then each address is inserted into the active station list.

3.4.7.1 Inserting New Addresses — When a new source address is inserted into the active station list, it must be inserted in its proper location in the ordered table. Often, this requires numerous existing entries to be shifted up so that the proper location is “vacated” to accommodate the new address.

As an example, Figure 3–16 shows addresses that must be moved to accommodate the new address 55. Address 55 must be inserted into location N to maintain the ordered nature of the table.

Since the new address 55 must be inserted into location N in the active station list, seven addresses must be transferred up one location at a time until location N is vacated.

A special location (0000) in the address table is used to initiate single-entry comparisons. When location 0000 is probed by the microprocessor, a single comparison is made and no further locations are probed.

Once the new entries are presorted, the microprocessor then uses 64-bit moves (see Section 3.4.7.3) to insert each address into its location in the active station list.

3.4.7.3 64-Bit Move — Shifting any address in the address table involves moving 48 bits of address and the corresponding 16 bits of status, a total of 64 bits. Since the microprocessor can manipulate only 16 bits of data at a time, each address shift would require four address-to-address transfers.

To accelerate the task of table maintenance, a special address range is mapped such that when the microprocessor addresses the special address range, all 64 bits (the address and status) are moved at once. The low-order 48 bits that make up the Ethernet address are read from RAM, temporarily stored in the comparator PALs and then written into the new location in RAM. The high-order 16 bits that make up the status word are written from the status RAM, temporarily stored in the microprocessor status register, and then written into the new location in the status RAM.

3.5 Loop Detection

LAN configurations cannot include any loops (redundant paths) because packets transmitted onto the LAN will circulate around any loop indefinitely.

All LAN Bridge 100 units have a loop detection feature. This is an automatic process that logically configures a loop-free extended network. This process, performed by a spanning tree computation algorithm, provides the following features to enhance extended network performance:

- Loop detection — If LAN Bridge 100 units are accidentally configured in a loop, the process will compute a loop-free portion of the topology. This prevents packets from circulating around the extended network indefinitely.
- Backup — LAN Bridge 100 units can be deliberately configured in a loop in such a way that one LAN Bridge 100 unit in the loop can serve as the backup for any one of the other LAN Bridge 100 units in the loop.
- The process is self-maintaining — The process is a continuous one that responds dynamically to changes in network topology. Changes may be caused by a malfunctioning LAN Bridge 100 unit or by using Remote Bridge Management Software to change parameters in the spanning tree computation algorithm. When changes are sensed, the algorithm recomputes the new network topology and again becomes stable and self-maintaining.
- The process is deterministic — A fixed set of rules controls the process so that when variables are changed, the results can be predicted.
- The process requires low network overhead — The messages that control the spanning tree are transmitted by participating LAN Bridge 100 units at 1-second intervals, thus using a very small percentage of the available network bandwidth. For example, if the interval between Hello messages is 1 second (the default), the overhead caused by Hello messages in a stable configuration is 64 bytes per second (0.005% on the network).

NOTE

Hello messages are 60 bytes in length. The Hello multicast address is (HEX) 09-00-2B-01-00-01, and the type field is 80-38. Because of the proprietary nature of the Hello message, other fields are not discussed in this manual.

3.5.1 The Spanning Tree Principle

Although the LAN Bridge 100 units in an extended network can be physically placed in an arbitrary fashion, the logical network that the LAN Bridge 100 units automatically create is always a **spanning tree**. The spanning tree has the following basic properties:

- There are no loops; that is, there is only one path between any two LAN Bridge 100 units.
- All LANs are connected.

The process by which the LAN Bridge 100 units construct the logical spanning tree from an arbitrary physical configuration is called the spanning tree computation process.

3.5.2 The Spanning Tree Computation Process

The spanning tree computation process consists of the following steps:

1. The extended network elects a single root bridge.
2. The extended network elects a designated bridge on each of its LANs. The designated bridge is the LAN Bridge 100 unit with the lowest path cost to the root bridge. A path cost is the sum of all the line-cost parameter values in any given path of lines. RBMS may be used to modify line cost parameter values (see Section 3.7).
3. All LAN Bridge 100 units except designated bridges turn off all lines, except for the single line that is the lowest path cost to the root bridge. Designated bridges turn off all lines except for the single line that is the lowest path cost to the root bridge and any lines attached to LANs for which they are the designated bridge.

When LAN Bridge 100 units are first activated (either by being powered on or by being reset using a bridge management command), each LAN Bridge 100 unit assumes that it is the root bridge of the spanning tree. Each LAN Bridge 100 unit sends out Hello messages on all its LANs declaring itself to be the root bridge of the extended network. These Hello messages are multicast to all other bridges that are directly connected to the same LANs. For example, in Figure 3-17, bridge B10's Hello message is received by bridges B3, B4, B9, B11, and B14.


Because the LAN Bridge 100 unit is declaring itself the root bridge, it also declares itself the designated bridge on all its LANs.

As LAN Bridge 100 units receive Hello messages from other LAN Bridge 100 units, each bridge compares the root bridge and designated bridge information in the received Hello messages to its own information. When performing the comparison, the bridge compares the new root bridge, root path cost, and designated bridge to the current root bridge, root path cost, and designated bridge. The root bridge and designated bridge fields are simply the network address of the root bridge and designated bridge with a root priority prefix. This root priority prefix can be set by bridge management and defaults to a fixed value that is equal for all bridges. If a bridge finds that a Hello message

contains better root information (lower values), it ceases to declare itself as the root, stores the new root information just received, and begins to send out the new information.

In a stable environment (all bridges and lines remain active), this process eventually leads to a single root bridge being elected as well as a single designated bridge per LAN. If no further events occur, the network will remain configured in the original spanning tree created by the spanning tree computation algorithm. Note that given the same bridges, the same set of conditions, and the same set of parameters, the spanning tree computation algorithm will always derive the same configuration.

Figure 3-17 shows the end result of the algorithm for a 14-bridge extended network. In this network, the bridges have the addresses B1 through B14. All root priority fields are assumed to be the default (128). If address B1 is assumed to be arithmetically less than address B2, then B1 will become the root bridge. Figure 3-17 also shows the lines that are put into the BACKUP state. In this state the bridges continue to receive Hello messages over these lines; however, the lines are not used for forwarding. This BACKUP state is what allows the bridges to form a loop-free tree from an arbitrary physical connection and yet have the flexibility to reconfigure themselves if a line goes down. Figure 3-17 also shows the designated bridges that result from the algorithm: bridges B2, B3, B5, B6, B8, B10, B12, B13, and B14.




- Broken lines indicate lines that are in the BACKUP state.
- Bridge B1 is the root bridge.
- Bridges B2, B3, B5, B6, B8, B10, B12, B13, B14 are the designated bridges.

LKG-0754

Figure 3-17: Result of the Spanning Tree Computation Algorithm

When a stable spanning tree configuration has been established and a bridge determines that it is not the root bridge, it stops sending out Hello messages. Instead, it sends out a Hello message only when it receives a Hello message from the designated bridge on the line to the root bridge. Other than topology change notification messages, bridges do not send a Hello message on the line over which they received the Hello message (the line to the designated bridge).

Therefore, in a stable network, each LAN has only one bridge (the designated bridge) that sends out Hello messages. Figure 3-18 shows the Hello message propagation in the stable network shown in Figure 3-17.



- ↑ and ↓ indicate direction of Hello message propagation.
- Broken lines indicate lines that are in the BACKUP state.
- Bridge B1 is the root bridge.
- Bridges B2, B3, B5, B6, B8, B10, B12, B13, B14 are the designated bridges.

LKG-0755

Figure 3-18: Hello Message Propagation

3.5.3 Spanning Tree Parameters

Certain parameters used in the spanning tree computation algorithm can be set using bridge management commands, and almost all parameters and variables involved in the algorithm can be read using management commands. This capability allows network managers to monitor and influence decisions made by individual bridges. For information on using this capability, refer to the *Remote Bridge Management Software Guide* (part no. AA-FY93A-TE).


3.5.4 Examples of the Spanning Tree Algorithm

Figure 3–19 represents a possible physical network in which seven networks are tied together with bridges to form a large extended LAN. This particular configuration may not be practical, but it serves to explain the function of the loop detection process. When the bridges in the network of Figure 3–19 are powered up, the following events occur as part of the spanning tree algorithm.

1. All bridges send out Hello messages which contain the bridge ID.
2. Since bridge 0 has the lowest ID, it is recognized as the root of the spanning tree.
3. The root bridge establishes itself as the designated bridge on each of the LANs it connects.
4. Each of the other bridges determines its distance from the root bridge in terms of the number of bridges between it and the root bridge.
5. The algorithm elects a designated bridge on each LAN. The designated bridge is the one with the lowest path cost to the root bridge.

Designated Bridge	LAN
1	D
2	C
3	F
4	E
5	G


6. In the case of LAN G, the distance to the root bridge is the same through bridge 5 or bridge 6. Bridge 5 becomes the designated bridge because its ID is lower than that of bridge 6. Bridge 6 goes into the BACKUP state because it is neither a root bridge nor a designated bridge.



LKG-0756

Figure 3-19: Hypothetical Extended LAN Using Bridges


The logical network formed as a result of the learning algorithm is shown in Figure 3-20.



LKG-0757

Figure 3-20: Logical Extended LAN Formed by Learning Algorithm

If a bridge goes from FORWARDING state to a BROKEN state, the spanning tree algorithm reconfigures the network topology. For example, if bridge 4 malfunctions, bridge 6 goes to the FORWARDING state and becomes the designated bridge for LAN E. Bridge 0 is still the root bridge. The new logical network topology is shown in Figure 3–21.




LKG-0758

Figure 3–21: Reconfigured Logical Extended Network

If the root bridge of Figure 3–21 malfunctions:

- Bridge 1 becomes both the root bridge and designated bridge for LANs A and D.
- Bridge 6 goes into the FORWARDING state and becomes the designated bridge for LAN E.
- Bridge 4 becomes the designated bridge for LAN C.
- Bridge 2 becomes the designated bridge for LAN B.


The new logical network topology is shown in Figure 3-22.



LKG-0759

Figure 3-22: Reconfigured Logical LAN with a New Root Bridge

Another example of an extended network is shown in Figure 3–23. Notice that bridges are connected in such a way that two physical loops exist. The loop detection function of the bridge helps to configure the network logically so that no loops exist.



LKG-0760

Figure 3–23: Extended LAN Showing Physical Loops

After the network is configured by the spanning tree algorithm:

- Bridge 0 is established as the root bridge of the spanning tree and is the designated bridge for LANs A and B.
- Bridge 1 is established as the designated bridge for LAN C.
- Bridge 3 is established as the designated bridge for LAN D.
- Bridges 2 and 4 enter the BACKUP state.

The logical configuration is shown in Figure 3–24.





Figure 3-24: Logical Extended LAN with Backup Bridges

In the network of Figure 3-24:

- If bridge 1 malfunctions, bridge 4 will become operational and the designated bridge for LAN C.
- If bridge 3 malfunctions, bridge 2 will become the designated bridge for LAN D.
- If bridge 0 (the root bridge) malfunctions, bridge 1 will become the new root bridge and designated bridge for LANs B and C, bridge 2 will become operational and the designated bridge for LAN D, and bridge 3 will become the designated bridge for LAN A.

The new logical network topology formed as a result of the root bridge malfunction is shown in Figure 3–25.



LKG-0762

Figure 3–25: Logical Network Topology for New Root Bridge

3.6 Foreground and Background Operations

The main function of the LAN Bridge 100 unit is filtering and forwarding packets. To operate efficiently under the heaviest network traffic conditions, the LAN Bridge 100 unit must be able to filter or forward a packet every 32 μ s.

During normal operation, the LAN Bridge 100 unit performs a variety of operations that includes:

- Performing loop detection functions.
- Receiving packets and writing them into packet memory.
- Using the descriptor rings to allocate buffer space and to find packets in packet memory.
- Transmitting packets from packet memory.

- Performing a binary search of the addresses in the Ethernet address table.
- Presorting new address entries.
- Storing new address entries in the Ethernet address table.
- Responding to bridge management instructions.
- Performing MOP operations.

Though all of the operations listed here support the main function of the LAN Bridge 100 unit, some operations may be less critical; thus their execution may be delayed until an operation of higher priority is completed.

To prioritize the various operations, those in the critical path of filtering or forwarding packets are considered foreground operations. Less critical operations are considered background operations.

3.6.1 Foreground Operations

Foreground operations are those operations that must occur so that the LAN Bridge 100 unit can filter or forward packets. The LAN Bridge 100 unit must:

- Receive packets from a LAN.
- Scan and write the receive descriptor ring.
- Write the packets into packet memory.
- Compare the destination address to the Ethernet address table.
- Decide whether to forward or filter the packets.
- Scan and write the transmit descriptor ring.
- Transmit (forward) the packets to the other LAN.
- Compare the source address to the Ethernet address table.
- Perform loop detection functions.

Most of the foreground operations are in the critical path of filtering or forwarding packets and are called critical path operations.

Loop detection functions, though not in the critical path, have the highest priority and are performed as foreground operations. Loops in an extended network cause unnecessary packet traffic that can degrade performance of critical path operations. This condition can ultimately prevent background operations from executing, since background operations are deferred until foreground operations are completed.

Loop detection functions maintain LAN Bridge 100 unit performance levels by detecting and eliminating loops. This optimizes usage of the critical path by minimizing unnecessary traffic, thus allowing time for execution of background operations.

3.6.2 Background Operations

Background operations support the main function of the LAN Bridge 100 unit but are not in the critical path of filtering or forwarding packets. Table 3–9 identifies major background operations in order of their priority.

Table 3–9: Background Operations

Operation	Description
Address table maintenance	Typical address table maintenance operations include: <ul style="list-style-type: none">• Updating the address table status word• Presorting new address entries• Storing new addresses in the address table
Bridge management operations	Responding to bridge management instructions.
Maintenance Operation Protocol (MOP) functions	MOP functions include responding to requests for system identification, loop data messages, or counter messages.
Timer functions	Timer functions include resetting various timers.
Counter maintenance functions	Counter functions include maintaining pointers for various queues such as the transmit and receive descriptor rings.
Control of state changes	State changes are dictated by timers and counters that are controlled in the background.

3.7 Remote Software and Bridge Access

Remote software such as Remote Bridge Management Software and LAN Traffic Monitor software use the network to communicate with the bridge. Bridge switches 3, 4, and 5 control communications between the bridge and the node running remote software. Table 3–10 describes the switch positions that control bridge access.

Table 3–10: Bridge Access Switches

Switch Number	Name	ON (Down)	OFF (Up)
3	Port A Access	<p>Stations on the LAN connected to port A that have bridge management capabilities are allowed to read and write (modify) bridge management parameters.</p> <p>If a load host resides on port A, the switch must be on for down-line loading software.</p>	<p>Stations on the LAN connected to port A that have bridge management capabilities can read but cannot write bridge management parameters.</p>
4	Port B Access	<p>Stations on the LAN connected to port B that have bridge management capabilities are allowed to read and write (modify) bridge management parameters.</p> <p>If a load host resides on port B, the switch must be on for down-line loading software.</p>	<p>Stations on the LAN connected to port B that have bridge management capabilities are allowed to read but cannot write bridge management parameters.</p>
5	Down-Line Load Enable	<p>Configures unit to operate as a LAN Traffic Monitor. Enables the unit to down-line load software (such as LTM Listener software) from a load host. When the switch is on, the bridge does not forward packets.</p> <p>Applicable Port A Access or Port B Access switch must be on so that the load host can write to bridge memory.</p>	<p>Configures unit to operate as a bridge. Disables the down-line function.</p> <p>RBMS can override the switch causing the bridge to operate as a LAN Traffic Monitor.</p>

3.8 Maintenance Operation Protocol (MOP)

The LAN Bridge 100 unit implements network interconnect (NI) maintenance protocols as specified in the *Digital Network Architecture Maintenance Operations Functional Specification* (part no. AA-X436A-TK). The operational description and message formats of these protocols are specified in that document and are not repeated here. The minimum functional requirements for an NI port are outlined in the *Digital Network Architecture Ethernet Node Product Architecture Specification* (part no. AA-X440A-TK).

The bridge supports the following minimum set of maintenance functions:

- Request identification and system identification packets
- Loop data and loop data packets
- Request counters and counters messages

These maintenance functions are further defined in Table 3-11.

Table 3-11: MOP Request and Response Messages

Request Message

Request Identification — A message from network management requesting a system identification message from the LAN Bridge 100

Loop Data — A message from network management instructing the LAN Bridge 100 to loop data.

Response Message

System Identification — A response from the bridge containing:

- MOP version number
- Console state
- Hardware address
- Functions
- Device

Looped Data — A message from the LAN Bridge 100 in response to a loop data message. Loop messages are forwarded in the direction indicated by the Ethernet address table.

3.9 Self-test

The self-test is resident in the LAN Bridge 100 firmware and is started on power up or soft reset. The self-test program is designed to:

- Test 95% of all stuck-at (SA1 or SA0) faults.
- Provide a maximum amount of fault isolation to the failing functional block.

The self-test is run in one of two switch-selectable modes:

- Normal
- Loop-on-self-test

In normal mode, the self-test starts from the power-up vector and begins testing at the lowest level with the CPU and ends testing at the more functional network tests. If no errors are found, self-test illuminates the Self-test OK status LED and jumps to the LAN Bridge 100 firmware. If an error is found, a code corresponding to the failing test is written to NVRAM, and the Self-test OK status LED remains off. The LAN Bridge 100 unit enters the BROKEN state, waits 15 seconds, and then reenters the SELF-TEST state. If an error is again detected, the error is handled in the same manner as before. The error code previously stored in NVRAM is overwritten with the most recent error code.

When the loop-on-self-test switch is turned on, the self-test operates the same as it does in normal mode except that it repeats the entire self-test until either an error is detected or the LAN Bridge 100 unit is powered off. The NVRAM write test is disabled after the first pass to avoid exhausting its write capability. The Self-test OK status LED is turned on at the end of the first pass and remains on unless an error is found. If an error is found, the Self-test OK status LED is turned off, and a code corresponding to the failing test is written to the NVRAM and the manufacturing register. The self-test then loops on the failing test and generates a trigger pulse (once each pass) at one of the external pins. This trigger pulse is an aid to troubleshooting the LAN Bridge 100 module.

The self-test contains three main modules:

1. Basic tests
2. LANCE tests
3. TLU tests

3.9.1 Basic Tests

Basic tests are run from ROM and include testing of the:

- Program ROM
- Program RAM
- NVRAM
- Ethernet address ROM
- Ethernet address RAM
- Timer check
- Packet memory
- Packet memory refresh

Sections 3.9.1.1 through 3.9.1.9 provide a brief description of each of these basic tests.

3.9.1.1 Program ROM Test — This test does the following:

- Performs a CRC32 calculation on the program ROM excluding the last longword location in the ROM
- Compares the results of the calculated CRC with the CRC character stored in the last word location of the ROM

3.9.1.2 Program RAM Test — Program RAM is tested for all stuck-at (SA1 or SA0) faults and coupling faults using a modified version of the Nair, Thatte, and Abraham's testing procedure (refer to *Functional Testing of Semiconductor Random Access Memories – Computing Surveys*, Vol 15, No. 3, Sept. 1983).

3.9.1.3 NVRAM Checksum Test — This test performs a checksum calculation on the NVRAM and compares the calculated checksum with the value stored in the NVRAM. The checksum is calculated using the 16-bit ones complement binary arithmetic, shifting before adding.

3.9.1.4 NVRAM Write Test — This test verifies that the NVRAM can be written into. The NVRAM write test is disabled after the first pass of the self-test in manufacturing mode.

3.9.1.5 Ethernet Address ROM Checksum Test — This test performs a checksum calculation on the Ethernet address ROM and compares the result of the calculation to the checksum value previously stored in the ROM. The checksum is calculated using the 16-bit ones complement binary arithmetic, shifting before adding.

3.9.1.6 Ethernet Address RAM Test — The Ethernet address RAM is tested for stuck-at faults (SA1 or SA0) and coupling faults using a modified version of the Nair, Thatte, and Abraham's testing procedure.

3.9.1.7 Timer Test — The timer test does the following:

- Starts the software clock and checks that the interrupt occurs at the correct vector and the correct IPL.
- Checks that the interrupt can be acknowledged or reset by reading address location 4000.
- Checks that the interrupt occurs at the correct interval (498 milliseconds).

3.9.1.8 Packet Memory Test — Packet memory is tested for stuck-at faults (SA1 or SA0) and coupling faults using a modified version of the Nair, Thatte, and Abraham's testing procedure.

3.9.1.9 Packet Memory Refresh Test — This test verifies that packet memory refresh is working. A data pattern is written into the RAM, the test waits 1 second, then reads the location to verify that the pattern read agrees with the pattern written. This is repeated until all bits are checked for SA1 and SA0 faults.

3.9.2 LANCE Tests

There are three categories of LANCE tests:

- Reset
- Internal loop
- External loop

3.9.2.1 LANCE Reset Test — This test verifies that LANCE chips can be reset to a known state. Bits in both LANCE chips are set to a predetermined state. LANCE A is reset and LANCE B is checked to verify that the bits are still set, then LANCE A is checked to verify that it is cleared. Next, bits in LANCE A are again set to the predetermined state. This time, LANCE B is reset and LANCE A is checked to verify that the bits are still set. Finally, LANCE B is checked to verify that it is cleared.

3.9.2.2 Internal Loop Tests — The internal loop tests include:

- Transmit CRC Logic Test
- Receiver CRC Logic Test
- Receive Bad CRC Test
- Collision Test
- Multicast Address Test
- Reject Physical Address Test
- Byte Swap and Broadcast Address Test

Transmit CRC Logic Test

This test clears the disable transmit CRC (DTCR) bit in the mode register to enable CRC generation on the packet transmission. The test then transmits a packet and compares the received CRC character with a precalculated value.

Receive CRC Logic Test

This test sets the disable transmit CRC (DTCR) bit in the mode register to disable CRC generation on packet transmission. The test then transmits a packet with a precalculated CRC character and verifies that the packet was received correctly with no CRC errors.

Receive Bad CRC Test

This test sets the disable transmit CRC (DTCR) bit in the mode register to disable CRC generation on packet transmission. The test then transmits a packet with a bad CRC character appended and verifies that the receiver flags a CRC error.

Collision Test

In this test, a packet is transmitted in internal loop mode with the mode register collision bit set. The test then verifies that the LANCE detects an error and retries the transmission 16 times. After the sixteenth retry failure, a retry error is indicated.

Multicast Address Test

This test checks the ability of the LANCE to accept or reject a packet with the multicast bit set in the destination address of the transmitted packet.

The test uses the following procedure:

1. Transmit a packet with an address that the logical address filter should accept.
2. Verify that the packet is transmitted and received correctly.
3. Transmit a packet with an address that the logical address filter should reject.
4. Verify that the packet is transmitted but not received.

Reject Physical Address

This test transmits a packet with the destination address not equal to the LANCE address. The test then verifies that the packet is transmitted correctly but is not accepted by the LANCE receiver. The test also checks the transmit status and looks for unexpected interrupts.

Byte Swap and Broadcast Address Check

The LANCE has the BYTE SWAP bit set in CSR3 and the destination address of the packet is a broadcast address. When an internal loop packet is transmitted, the receiver unit should accept the packet. Transmit and receive status is checked along with the data.

3.9.2.3 External Loop Tests — These tests are used to determine the port-to-port loop configuration. A flag is set to enable the network port to network port testing if there is a valid configuration available. In port-to-port testing, packet sizes vary from 64 to 1500 bytes. There are two external loop tests:

- External loop test
- Port-to-port loop test

External Loop Test

This test takes the following steps to determine what network loopback configuration is connected to the LAN Bridge 100 unit:

1. An external loop packet is sent on port A.
2. If the packet is not received, the self-test considers it a hard error and halts.
3. If the packet is received back on port A, the self-test determines that there must be transceivers connected to both ports.
4. An external loop packet is sent on port B.
5. Status and data are checked at both ports.
6. If the external loop packet transmitted on port A is received on port B, the self-test determines that there must be a port-to-port loop configuration.

7. If the external loop packet transmitted on port B is received on port A, the self-test sets a flag to indicate that a port-to-port configuration exists.
8. If any of the status/data checks or transmit/receptions fail, the self-test indicates the hard error and halts.

Port-to-Port Loop Test

The port-to-port test requires that both ports are connected together through a manufacturing loopback cable. The test also requires that the port-to-port flag be set.

NOTE

This test will fail unless the loop-on-self-test switch located in the LAN Bridge 100 I/O panel is set to the Enable position.

The port-to-port loop test creates a maximum amount of activity on the LAN Bridge 100 hardware by starting operations with the binary search hardware and timer hardware at the same time as it does port-to-port looping through the manufacturing loopback cable. The size of the packets transmitted varies from 64 bytes to 1518 bytes. Both transmit and receive data chaining is tested.

3.9.3 Table Lookup (TLU) Tests

The TLU tests test the following:

- Status RAM
- Basic binary search
- Binary search engine

3.9.3.1 Status RAM Test — Status RAM is tested for stuck-at faults (SA1 or SA0) and coupling faults using a modified version of the Nair, Thatte, and Abraham's testing procedure.

3.9.3.2 Basic Binary Search Test — This test verifies that the search registers can be read from and written into by the processor. The test also verifies that data from the search registers can be written into the Ethernet address table RAM in 64-bit increments by writing to addresses FC000–FFFFFF (this tests the 64-bit move function).

3.9.3.3 Binary Search Engine Test — This test verifies that the binary search engine can search the network address RAM for entries scattered throughout its address range. A check is made to ensure that a binary search is actually done during the search for stored network addresses.

4

Technical Description


4.1 Introduction

This chapter contains a technical description of the hardware and hardware subsystems that make up the LAN Bridge 100 unit. The discussion takes a top-down approach, beginning with the system level description (LAN Bridge 100 system) and ending with discussions at the circuit level.

4.2 Hardware Overview

The LAN Bridge 100 unit is made up of four subsystems, as shown in Figure 4-1. Each subsystem provides the following functions:

- The network interconnect (NI) subsystem receives packets from the LAN on one side of the LAN Bridge 100 unit and stores them in memory. The NI may forward the packets to the LAN on the other side depending on a determination made by the LAN Bridge 100 unit firmware.
- The table lookup (TLU) subsystem performs searches of 48-bit Ethernet addresses stored in an 8K-by-48-bit address table. The results of the search are used by the LAN Bridge 100 firmware to determine whether a packet should be forwarded to the LAN on the other side of the LAN Bridge 100 unit.
- The processor subsystem keeps track of packets stored in packet memory and determines whether those packets will be filtered or forwarded. The processor maintains the Ethernet address table in the TLU. The processor also responds to loop detection messages, Maintenance Operation Protocol functions, and bridge management or Remote Bridge Management Software (RBMS).
- The power supply subsystem provides all of the voltages required to support the NI, TLU, and processor subsystems as well as the transceivers connected to port A and port B.




LKG-0763

Figure 4-1: LAN Bridge 100 Subsystems

4.2.1 Overview of LAN Bridge 100 Functional Blocks

The LAN Bridge 100 subsystems interact through a system of bus structures and control lines. The block diagram in Figure 4-2 shows relationships among the major circuits of the LAN Bridge 100 unit.

The design incorporates a system of bus switches (latches) that isolate memory structures in each of the subsystems. This important design element enhances the operating speed of the LAN Bridge 100 unit by allowing concurrent operation of the NI, TLU, and processor subsystems.



LKG-0764

Figure 4-2: LAN Bridge 100 Functional Block Diagram

4.2.2 The Data Path

This section provides a general discussion of one typical LAN Bridge 100 operation to illustrate how the hardware is used to forward packets. Each of the circuits mentioned in the discussion here is shown in Figure 4-2.

As an example of LAN Bridge 100 operation, assume that a minimum size packet is present at port A. The LAN Bridge 100 unit performs the following steps:

1. The serial interface adapter (SIA) receives the packet from port A transceiver and sends it to the Local Area Network Controller for Ethernet (LANCE).

The LANCE stores the packet in packet memory and clears an OWN bit in the LANCE A receive descriptor ring. This causes the OWN interrupt generator to send an OWN interrupt to the processor.

2. The processor responds to the OWN interrupt by reading the LANCE A receive descriptor ring that points to the buffer containing the packet. The processor then reads the 48-bit destination address (located in the first 6 bytes of the packet) and copies it into the TLU.
3. The TLU performs a binary search, which compares the destination address against addresses already stored (in ascending numerical order) in the address table. The processor uses the results of the search to determine whether the packet must be forwarded.

One of the following will occur:

- If no match is found (that is, if the destination address does not equal a known source address), the processor, by default, forwards the packet to the LAN at port B. Go to step 4.
- If a match is found (that is, if the destination address equals a known source address), the processor checks the corresponding 16-bit status word in the processor RAM to determine the location (port A or port B) of the station.

If the destination is on side A, the processor disregards the packet. Go to step 5.

If the destination is on side B, the processor determines, by default, that the packet is to be forwarded to the LAN at port B.

4. The LAN Bridge 100 unit transmits the packet to port B as follows:
 - The processor writes the data from the LANCE A receive descriptor ring into the LANCE B transmit descriptor ring, resets the OWN bit in the descriptor ring, and sets the transmit demand bit in LANCE B CSR.
 - When transmit conditions at port B are met, LANCE B reads the packet from packet memory and transmits it to the LAN at port B using the serial interface adapter (SIA).
5. The LAN Bridge 100 unit compares the source address of the incoming packet to addresses stored in the Ethernet address table. This keeps the address table and status word up to date.

4.3 Processor Subsystem

The processor subsystem is responsible for overall operation of the LAN Bridge 100 unit. The processor allocates buffers to the LANCE network controllers, initiates binary searches, and maintains the address table of the TLU.

4.3.1 Processor Circuit Descriptions

The processor consists of the following circuitry (see Figure 4-3):

- Motorola 68000 microprocessor
- Interrupt controller
- Processor memory
- System timers
- Reset circuitry

To minimize capacitive loading of the processor data bus (D00–D15), a secondary data bus (UD00–UD15) is incorporated into the architecture (see Figure 4-3). The secondary data bus is joined to the primary data bus by a pair of octal bus transceivers. The secondary data bus is enabled by the address decoding circuitry.




Figure 4-3: Processor Subsystem

4.3.2 The Microprocessor

The heart of the processor subsystem is the microprocessor. The microprocessor used in the LAN Bridge 100 unit is the Motorola 68000, a single-chip device with address and data on separate buses. The 68000 has built-in mechanisms for handling incoming interrupt information and for controlling bus arbitration. The 68000 also provides status information that allows decoding of the current bus operation.

The 68000 microprocessor has a 32-bit internal architecture and a 16-bit external data path. The address bus is 23 bits wide with separate control lines for accessing the upper and lower bits either as separate bytes or together as words. The processor is driven by a 10-MHz (100-ns) system clock that results in a processor cycle time of 400 ns without WAIT states.


4.3.3 Interrupt Controller

The interrupt controller (see Figure 4-4) monitors the interrupt priority level (IPL) lines. The controller responds to the interrupt request if the priority of the request is higher than the priority of the processor. If two or more interrupt requests are presented to the interrupt controller at the same time, the highest level is passed on to the processor and the other levels remain pending.

The microprocessor acknowledges the interrupt by placing the interrupt level on address lines A1 through A3. When the interrupt vector generator is enabled by the function code lines (FC0-FC2), it reads the address lines A1 through A3 and places a corresponding interrupt vector onto the data bus. Table 4-1 identifies each interrupt, its priority, and its vector.

Table 4-1: Hardware Interrupts

Interrupt Level	Vector Address	Interrupt
6	46	Watchdog timer.
5	45	LANCE B interrupt. Generated when a CSRO status flag is set for LANCE B.
4	44	LANCE A interrupt. Generated when a CSRO status flag is set for LANCE A.
3	43	OWN bit interrupt B. Generated when an OWN bit (set by LANCE B) is detected.
2	42	OWN bit interrupt A. Generated when an OWN bit (set by LANCE A) is detected.
1	41	Fiber-optic self-test. Generated during self-test if a fiber-optic module is installed. This causes self-test to verify operation of the fiber-optic circuitry.



LKG-0766

Figure 4-4: Hardware Interrupt Configuration

4.3.4 Processor Memory

The processor memory contains the following components:

- PROM (stores the program code)
- RAM (holds a copy of the program code)
- RAM (holds the address table status)
- PROM (holds the physical address of the LAN Bridge 100 unit)
- NVRAM (records diagnostic and bridge management parameters)
- Miscellaneous control registers

4.3.4.1 PROM — The PROM (programmable, read-only memory) section of processor memory stores 16K words and contains all native LAN Bridge 100 firmware, including self-test.

4.3.4.2 Program RAM — Program RAM contains 8K words and is made up of two 8K-by-8-bit static RAM devices. The program RAM contains a copy of the LAN Bridge 100 program code and data structures. The program executes from program RAM after the LAN Bridge 100 unit has powered up and self-test has passed.

4.3.4.3 Address Table Status RAM — The address table status RAM contains 8K words and is made up of two 8K-by-8-bit static RAM devices. The address table status RAM contains status words that correspond to each entry in the Ethernet address table of the TLU. The status includes information such as the following:

- The port (port A or port B) on which the address resides
- The age of the address
- Whether the address is used by the LAN Bridge 100 unit for special purposes
- Whether the address is learned through normal LAN Bridge 100 operation or is fixed (written into memory by bridge management so that it cannot age)

4.3.4.4 Physical Address PROM — The physical address PROM is a 32-bit-by-8-bit PROM device. This PROM contains the unique 48-bit (6 bytes) station address of the LAN Bridge 100 unit and a 2-byte checksum of the address.

The physical address is used for bridge management functions, Maintenance Operation Protocol (MOP) functions and for loop detection purposes.

4.3.4.5 NVRAM — Nonvolatile RAM (NVRAM) contains 2K bytes and is a single 2K-by-8-bit memory device. NVRAM stores diagnostic information and LAN Bridge 100 parameters during a power-down period.

NVRAM is write protected by the reset circuitry during all power-up and power-down cycles.

4.3.4.6 Miscellaneous Control Registers — The miscellaneous control registers monitor and control operation of the LAN Bridge 100 unit. The registers and their functions are as follows:

- Miscellaneous control registers enable interrupts and loopback functions. A memory map control bit causes RAM address space to be remapped after power up is completed (see Section 4.3.5.6).
- A self-test error latch holds error status if errors occur during self-test.
- A programming selection buffer holds LAN Bridge 100 configuration information such as switch settings or whether a fiber-optic module is present. The processor uses this buffer to set program functions.

4.3.5 Processor Memory Map

The LAN Bridge 100 unit has approximately 1 megabyte of address space. When the LAN Bridge 100 unit is powered on, the program executes from program ROM. Program RAM is hidden and is not accessible. The memory map shown in Figure 4–5 illustrates the mapping used for powering up the LAN Bridge 100 unit.

PROGRAM ROM	(0000) (03FFF)
MISCELLANEOUS CONTROL	(04000) (04FFF)
NONVOLATILE MEMORY	(05000) (05FFF)
LANCE DESCRIPTOR RINGS	(06000) (06FFF)
LANCE CSR REGISTERS	(07000) (07FFF)
PROGRAM ROM	(10000) (1FFFF)
PACKET MEMORY	(20000) (3FFFF)
ETHERNET ADDRESS TABLE	(E0000) (EBFFF)
ETHERNET ADDRESS TABLE STATUS RAM	(EC000) (EFFFF)
COMPARE-AND-MOVE REGISTERS	(F4000) (F4FFF)
ETHERNET ADDRESS TABLE STATUS RAM	(FB000) (FBFFF)
ETHERNET ADDRESS TABLE 64-BIT MOVE	(FC000) (FCFFF)

LKG-0767

Figure 4-5: Address Mapping Used During Power Up

After the bridge completes self-test, the address space for program RAM is remapped to the range of 00000–03FFF and the program firmware is copied from ROM into program RAM. Thereafter, the program executes from program RAM. Executing the program from RAM provides more flexibility by allowing the use of soft interrupt vectors.

The memory map shown in Figure 4-6 illustrates the mapping used for operation of the LAN Bridge 100 unit after power up and self-test have completed.

PROGRAM RAM	(0000) (03FFF)
MISCELLANEOUS CONTROL	(04000) (04FFF)
NONVOLATILE MEMORY	(05000) (05FFF)
LANCE DESCRIPTOR RINGS	(06000) (06FFF)
LANCE CSR REGISTERS	(07000) (07FFF)
PROGRAM ROM	(10000) (1FFFF)
PACKET MEMORY	(20000) (3FFFF)
ETHERNET ADDRESS TABLE	(E0000) (EBFFF)
ETHERNET ADDRESS TABLE STATUS RAM	(EC000) (EFFFF)
COMPARE-AND-MOVE REGISTERS	(F4000) (F4FFF)
ETHERNET ADDRESS TABLE STATUS RAM	(FB000) (FBFFF)
ETHERNET ADDRESS TABLE 64-BIT MOVE	(FC000) (FCFFF)

LKG-0768

Figure 4-6: Address Mapping Used for Normal LAN Bridge 100 Operation

4.3.5.1 Program ROM — Program ROM contains the firmware that is used to control the operation of the LAN Bridge 100 unit.

During power up, program ROM is mapped into the ranges of 00000–03FFF and 10000–1FFFF. After power up and self-test are complete, RAM is mapped into the range of 00000–03FFF and the firmware (except for self-test) is copied into RAM. Subsequent program execution is from RAM.

4.3.5.2 Miscellaneous Control — The miscellaneous control region (04000–04FFF) contains various registers used to monitor and control operation of the LAN Bridge 100 unit (see Section 4.3.4.6).

4.3.5.3 Nonvolatile Memory — Nonvolatile memory (NVRAM) resides in memory address space 05000–05FFF. NVRAM stores bridge management parameters when the LAN Bridge 100 unit is powered off. If power up self-test fails, NVRAM stores a diagnostic error code.

4.3.5.4 LANCE Descriptor Rings — The LANCE A and LANCE B transmit and receive descriptor rings reside in address space 06000–06FFF. To clear the OWN interrupt, the OWN bit in the descriptor rings are cleared using the range of 06000–06FFF.

The LANCE descriptor rings are used to allocate and find buffer space in packet memory.

Note that the descriptor rings also reside within the packet memory address space and range from 26000 through 26FFF. This is the range used by the LANCE processors to access the descriptor rings.

4.3.5.5 LANCE CSR Registers — The LANCE A and LANCE B CSR registers (a total of eight registers) are used to monitor and control the operation of the LANCE processors. The registers are mapped into the range of 07000–07FFF.

4.3.5.6 Program RAM — Program RAM is used for program execution after the LAN Bridge 100 unit powers up and self-test is complete.

During power up of the LAN Bridge 100 unit, program RAM is hidden and is not accessible. After power up and self-test are complete, RAM is remapped into the range of 00000–03FFF and the firmware (except for self-test) is copied into RAM. The interrupt vector address space is included in RAM after RAM is remapped.

4.3.5.7 Packet Memory — Packet memory (20000–3FFFF) is used to store packets while the LAN Bridge 100 unit is determining whether to forward or filter the packets. The LANCE A and LANCE B descriptor rings and the initialization blocks also reside in packet memory.

4.3.5.8 Ethernet Address Table — The Ethernet address table (EC000–EBFFF) is 8K by 48 bits and contains 48-bit Ethernet station addresses. A station address (and a corresponding status word) is stored when a packet is received from that station.

4.3.5.9 Ethernet Address Table Status RAM — The Ethernet address table status RAM is 8K by 16 bits and contains 16-bit status words that correspond to each station address stored in the Ethernet address table (see Section 4.3.4.3).

The Ethernet address table status RAM is mapped into two regions of address space: EC000–EFFFF, and F8000–FBFFF. Both regions contain the same data.

4.3.5.10 Compare-and-Move Registers — Three 16-bit compare-and-move registers (F4000–F4FFF) are used to store search arguments (48-bit Ethernet addresses), while magnitude comparisons are performed by the TLU. These registers are also used in the read/write operation associated with a 64-bit move.

Each register and address is identified as follows:

F4xx0	Search register with address bits < 47–32 >
F4xx2	Search register with address bits < 31–16 >
F4xx4	Search register with address bits < 15–00 >
F4xx6	Results register

Writing to address F4xx4 triggers the binary search of the Ethernet address table. The binary search starts at a midpoint in the address table defined by address lines BA13–BA03. If the address lines BA13–BA03 are all zeros, then a single entry comparison is performed using location 0 in the Ethernet address table.

Writing to address F4xx6 also triggers a binary search of the Ethernet address table. In this case, the TLU searches the address table using the same search argument that was used for the previous search.

4.3.5.11 Ethernet Address Table 64-Bit Move — Address space FC000–FCFFF accesses the Ethernet address table and the Ethernet address table status RAM. The memory configuration addressed in this range is 8K by 64 bits (64 bits include a 48-bit Ethernet address and a 16-bit status word). This capability is necessary for performing a 64-bit move (see Chapter 3 for more information on 64-bit moves).

4.3.6 System Timers

The system timer is driven by the 10-MHz system clock from the NI subsystem and provides the following functions:

- Clock interrupt to the processor
- Hardware watchdog timer
- Clock to the refresh counter

The clock interrupt to the processor occurs about every 0.5 seconds and drives a real-time clock. This clock is used to keep track of how long each network address has remained unused in the address table. If a station has been inactive for 2 minutes, the LAN Bridge 100 unit considers the corresponding address to be invalid.

The watchdog timer is used as a hardware check on the firmware. It is the responsibility of the firmware to reset the watchdog timer once every 15 seconds. If the timer is not reset, it will expire and cause a system reset. If the firmware is operating properly, this situation will not occur.

The clock to the refresh counter occurs every 12.5 μ s and is used to increment the refresh counter.

4.3.7 Reset Circuitry

The reset circuitry is responsible for the following functions:

- Resetting the hardware on power up
- Resetting the hardware upon a watchdog timer reset
- Write protecting the NVRAM during power up, power down, and reset sequences

The reset circuitry monitors VCC (5 volts) to sense when power is stable. During power up, the reset circuitry holds the LAN Bridge 100 unit in reset mode and keeps the NVRAM write protected until all state-saving devices are stable.

4.4 Network Interconnect Subsystem


The network interconnect (NI) subsystem provides the interface between the other LAN Bridge 100 subsystems and the physical channel of the extended LAN by performing the following functions:

- Receiving and decapsulating packets from the LAN connected to port A or port B of the LAN Bridge 100 unit.
- Writing packets into packet memory in locations determined by the processor.
- Writing packet status information into the receive descriptor rings.
- Reading the transmit descriptor rings to determine the location(s) of any packet that is ready for transmission.
- Reading packets from packet memory and encapsulating packets for transmission to the LAN connected to side A or side B.
- Transmitting encapsulated packets to the specified LAN according to the Ethernet channel access method.

4.4.1 NI Subsystem Circuits

The NI subsystem is made up of two physical channel interface circuits (NI ports), a LANCE CSR control circuit, and a packet memory. Figure 4-7 is a block diagram of the NI subsystem.

The NI subsystem bus structures and packet memory are isolated from processor and TLU bus structures and memory. This allows concurrent operation of the NI, processor, and TLU and helps minimize the time required to perform LAN Bridge 100 functions. In addition, the LANCE processor and microprocessor buses have independent access to packet memory. This allows maximum utilization of the memory cycle time and guarantees minimal delay for access.



LKG-0465

Figure 4-7: NI Subsystem

4.4.2 NI Ports

The NI subsystem has two ports. Each interface circuit connects to one of the two LANs joined by the LAN Bridge 100 unit. Each NI port consists of one SIA chip and one LANCE chip. This discussion describes one NI port; the other NI port is identical.

NOTE

In a remote LAN Bridge 100 unit, a fiber-optic module is included in port A. The fiber-optic interface is described in Section 4.4.2.5.

4.4.2.1 SIA Chip — The SIA chip is a 24-pin device that interfaces the LANCE chip to the network transceiver. The SIA chip performs the following functions:

- Detection of data presence on the network
- Conversion of collision-presence signals to TTL levels
- Manchester decoding of data
- Manchester encoding of data
- Interfacing of differential signal line pairs and TTL levels

The SIA chip requires a 20-MHz crystal for its internal oscillator. The oscillator generates the clock for the Manchester-encoded data stream. The 20-MHz is halved to provide the 10-MHz transmit clock and receive clock. The LANCE uses the transmit clock to set the data rate. The LANCE uses the receive clock to gate sampling of the receive data line.

4.4.2.2 LANCE Chip — The LANCE is a 48-pin, very-large-scale-integration (VLSI) device designed to provide data link services over a CSMA/CD local area network. The LANCE chip performs the following functions:

- Direct memory access to packet memory
- Buffer management by using descriptor rings in packet memory
- Runt packet filtering
- CRC generation and checking
- Transmission backoff and retry

The LANCE operates in promiscuous mode. This mode causes the LANCE to accept all packets regardless of their destination address. The LANCE is monitored and controlled by control and status registers (CSRs), which are physically separate from the LANCE. CSR control is described in more detail in the following section.

The LAN Bridge 100 unit does not use the LANCE's ability to generate CRC. However, the bridge does preserve incoming CRC sequences for use on outgoing packets. Note that the probability of externally undetectable errors originating within the LAN Bridge is no greater than 1 in 109.

A table called the initialization block that is stored in system memory is loaded into packet memory upon initialization. The initialization block is used by the LANCE to do the following:


- Determine the size and location of the descriptor rings in packet memory
- Set operating parameters for the LANCE (such as promiscuous mode)

LANCE Control and Status Registers — There are four control and status registers (CSRs) within each LANCE that are programmed by the processor. The CSRs are accessed through a register address port (RAP) and a register data port (RDP) as shown in Figure 4–8.

Data is read from or written into a CSR in a two-step operation.

1. The CSR is addressed by writing the register address into the register address port (RAP) of the LANCE.
2. Data that is read from or written into the register data port (RDP) is also read from or written into the CSR that is selected by the RAP.

During the final stage of LAN Bridge 100 initialization, the microprocessor sets the RAP to CSR0. Once the RAP is set, it remains set until rewritten, and it is not rewritten once the LANCE starts operation. With the RAP set to CSR0, each time the RDP is accessed, CSR0 is accessed.



LKG-0464

Figure 4-8: CSR Control

4.4.2.3 LANCE CSR Bus Registers — The CSR bus registers are bidirectional latches that interface the processor data bus to the LANCE address/data bus (see Figure 4–8). The latches isolate the two buses to allow concurrent operation of the NI and processor subsystems. They also minimize processor WAIT states since the processor does not have to wait for the LANCE address/data bus to read the LANCE CSR registers.

4.4.2.4 LANCE CSR Control Circuit — The LANCE CSR control circuit (see Figure 4–8) manages the transfer of data between the LANCE internal CSRs and the CSR bus registers.

After the LANCE is initialized, the register address port (RAP) is set to CSR0 (see Section 4.4.2.3). CSR0 remains selected as long as the LANCE continues to operate.

When the processor wants to read from CSR0, it actually reads from the CSR bus register, which is part of the processor's memory map. Reading from the bus register triggers the CSR control to initiate a read from the RDP into the bus register. Therefore, each time the bus register is read, the data in the bus register is updated. This guarantees that the processor does not have to wait for the LANCE bus, but it also means that the data obtained from the bus register is old and must be treated as such. The processor must read the bus register twice to get fresh data. However, the processor will incur WAIT states if it attempts back-to-back accesses of the CSR while the LANCE is performing a DMA operation.

When the processor wants to write to CSR0, it writes to the CSR bus register. The CSR control arbitrates for the LANCE bus and writes the data into CSR0 as soon as the LANCE bus is available. The maximum delay incurred is 4.8 μ s.

4.4.2.5 Fiber-Optic Module — The fiber-optic module is included only in remote LAN Bridge 100 units and performs functions similar to those of a transceiver.

The fiber-optic module converts emitter-coupled logic (ECL) levels from the SIA to optical signals for transmission through a fiber-optic cable. Conversely, optical signals received from the fiber-optic cable are converted to ECL levels for use by the SIA circuitry.

4.4.3 LANCE Address Latch

There are two LANCE address latches, one for each LANCE (see Figure 4–7). This section describes the address latch for one LANCE. An identical circuit exists for the other LANCE.

The LANCE address latch multiplexes the LANCE address/data bus to provide a separate address bus and data bus for packet memory access.

The LANCE latches a packet memory address into the latch by asserting the address latch enable (ALE) signal to the latch and writing the address onto the LANCE address/data bus. The latched address is then placed onto the LANCE address bus where it is used by the address multiplexer to access packet memory.

When the ALE signal is deasserted, the LANCE address/data bus becomes available for data functions.

4.4.4 OWN Interrupt Generator

There are two OWN interrupt generators, one for each LANCE (see Figure 4–7). This section describes the OWN interrupt generator for one LANCE. An identical circuit exists for the other LANCE.

The OWN interrupt generator is a programmed array logic (PAL) device that generates an interrupt when it detects that an OWN bit has been set in the LANCE receive descriptor ring. The interrupt is sent to the interrupt controller in the processor subsystem.

In packet memory, every fourth address in the range of receive descriptor ring addresses contains an OWN bit. When the OWN interrupt generator detects these specific addresses on the LANCE address bus, it looks at bit < 15 > of the LANCE address/data bus (bit < 15 > carries the OWN bit). If the PAL detects that an OWN bit is cleared, it generates an interrupt signal. This signals the processor that a packet memory buffer is ready for processing by the processor and TLU subsystems. See Chapter 3 for more information on the OWN bit function.

4.4.5 Packet Memory

Packet memory (see Figure 4–7) stores packets during the forward/filtration process. In addition, this memory contains the descriptor rings (transmit and receive descriptor rings) and the initialization blocks for each LANCE.

The packet memory architecture allows access by either of the LANCES, by the processor, or by the refresh address counter.

The cycle times of the devices that access packet memory are inherently slower than the memory cycle access time, which is about 300 ns. To minimize memory circuit WAIT states, data bus latches hold the data retrieved from packet memory until it can be read by the requesting device. This improves memory performance by allowing the memory to begin its next cycle immediately without waiting for slower devices.

Packet memory is made up of the following circuits:

- Address multiplexer
- Memory control circuit
- Dynamic RAM (DRAM)
- Data latches

4.4.5.1 Address Multiplexer — The address multiplexer (see Figure 4–7) determines which device (LANCE A, LANCE B, the processor, or the DRAM refresh counter) addresses packet memory. Steering inputs to the multiplexer come from the memory control circuits (see Section 4.4.5.2).

When a 16-bit address is present at the input to the multiplexer, the steering inputs apply the high-order (row) address and then the low-order (column) address onto the DRAM address bus.


4.4.5.2 Memory Control Circuit — The memory control circuit (see Figure 4–7) coordinates memory activity and arbitrates the priority of the LANCE, processor, and refresh counter. Table 4–2 identifies the priorities assigned to the four devices. Figure 4–9 is a block diagram of the memory control circuit. The memory control circuits perform the following functions:

- The synchronization circuit synchronizes the control signals of memory accessing devices with the memory control circuit.
- The arbitration circuit selects which device (processor, LANCE A, LANCE B, or refresh counter) addresses packet memory.
- The timing circuit ensures that row and column address strobes and the write signal are properly gated to the DRAM.

Table 4–2: Device Priorities

Priority	Device
1	Refresh
2	Processor
3	LANCE A *
4	LANCE B *

* Note: LANCE A and LANCE B alternate priority.



LKG-0770

Figure 4-9: Memory Control Circuit

4.4.5.3 Dynamic Packet Memory — The packet memory stores packets while the processor and TLU perform the forward filtration process. The memory also contains the LANCE descriptor rings and the LANCE initialization blocks. Refer to Chapter 3 for a description of the initialization blocks and descriptor rings.

DRAM memory is made up of sixteen 64K-by-1-bit devices (each device is an array of 128 rows by 512 columns).

The DRAM address bus is 16 bits wide but is multiplexed into 8-bit row addresses and 8-bit column addresses. Row address strobe (RAS) and column address strobe (CAS) signals select which row and column are read from memory. RAS, CAS, and WRT (write-enable) signals are driven by the memory control circuit (see Section 4.4.5.2).

The DRAM data bus is 16 bits wide. When CAS strobes the DRAM, the data is either read from or written to the data bus depending on the state of WRT.

Each DRAM row must be refreshed at least once every 2 ms. Since each DRAM contains 128 rows, one row must be refreshed every 15 μ s (128 x 15 μ s = 1.920 ms). To allow for any delays, the refresh counter addresses a row in memory every 12.5 μ s (see Section 4.4.5.4).

4.4.5.4 Refresh Counter — The refresh counter increments and accesses a row address in packet memory every 12.5 μ s.

The refresh counter is made up of a dual 4-bit binary counter that is configured as an 8-bit binary counter. The incoming 12.5- μ s clock that increments the counter comes from the system timer (see Section 4.3.6).

For additional information on the DRAM refresh cycle, see Section 4.4.5.3.

4.4.5.5 Data Bus Latches — Packet memory data bus latches (see Figure 4–7) are controlled by the memory control circuit (see Section 4.4.5.2). The latches interface the packet memory data bus (PMD 00–15) to the following bus structures:

- ADAL 00–15 (LANCE A address/data bus)
- BDAL 00–15 (LANCE B address/data bus)
- D00–15 (processor data bus)

The data latches and buffers perform two functions:

- They isolate the memory data bus from external data buses to allow concurrent operation of the NI and other subsystems.
- They hold data retrieved from memory until the data can be read by the requesting device. This improves memory performance by allowing the memory to begin its next cycle without waiting for slower devices.

4.5 Table Lookup Subsystem


Each time a LANCE writes a new packet into packet memory, the processor reads the destination address of the packet and sends the address to the table lookup (TLU) subsystem.

The TLU compares the destination addresses from incoming packets to addresses stored in memory. The results of each comparison are made available to the processor, which determines whether the packet should be forwarded or ignored.

The TLU, with assistance from the processor subsystem, also maintains the Ethernet address table. Table maintenance functions include comparing source addresses from incoming packets to addresses stored in memory. The results of each comparison are made available to the processor, which determines whether the address should be stored in the address table or ignored.

4.5.1 Overview of the TLU

The functional blocks of the TLU are shown in Figure 4–10. This illustration can guide you through the overview of TLU operation.



LKG-0771

Figure 4-10: Table Lookup Subsystem

1. The processor initiates the TLU binary search process in the following way:
 - The processor writes the search argument (the new 48-bit Ethernet address) from the data bus into the search compare registers of the comparator PALs. The processor uses three write cycles to write the three address words into the search compare registers. The words are written in sequence of the low, middle, and high word of the address.
 - The search control and status PAL detects when the final (highest) byte is written into the search compare registers by the processor. After a delay of about 50 ns, the search control and status PAL deasserts the DONE signal. The deasserted DONE signal does two things:
 - After 50 ns, it steers the multiplexer to select subsequent addresses from the search control and status PALs.
 - After 100 ns, it prevents the processor from addressing any registers in the TLU.

Before the DONE signal is deasserted (during the 50-ns delay), the processor addresses the location in the Ethernet address table, where the binary search begins. This address is used to start the binary search. Addresses for subsequent probes of the address table are provided by the search control and status PALs.
2. Once the processor has written the first address to be probed in the Ethernet address table, the search control and status PALs in the TLU assume control of the binary search process.

NOTE

Ethernet addresses are stored in order of magnitude in the address table, higher locations having the higher addresses. This ordering is necessary for the binary search process to function. See Chapter 3 for a detailed description of the binary search process.

- The comparator PALs perform bit-by-bit comparisons of the latched search argument against addresses probed in the address table.
 - The search control and status PALs use the results of each comparison (> [greater than], < [less than], or = [equal to]) to determine the next step in the search process.
 - “Greater than” causes a search to be performed at a higher location in the address table.
 - “Less than” causes a search to be performed at a lower location in the address table.
 - “Equal to” indicates a match. A match terminates the search.
 - The binary search ends either when a match is found or when the bottom address of the binary tree (an odd address) is probed and no match has been found.

The search control and status PALs assert DONE. This enables the decoder to allow the processor to read the search results register. This register indicates whether a match was found and contains the location of the last address searched.
3. After the search, the search control and status PALs report the results of the search by setting “equal to” or “greater than” bits in the search results register. The address of the last location probed is included in the register data. These results dictate what the processor does with the location address.

- If the “equal to” bit is set, it means that a match was found.

If a destination address was compared, the processor checks a corresponding 16-bit status word in RAM to determine the location (bridge port A or port B) of the Ethernet address and whether the packet should be forwarded or ignored. The status word location in RAM corresponds to the Ethernet address location in the address table.

If a source address was compared, the processor updates the status word in RAM.

- If the “greater than” bit is set, it means the search argument was greater than the address in the last location probed. This also indicates that no match was found.

If a destination address was compared, the processor forwards the packet since it cannot determine the location (port A or B) of the station with that address.

If a source address was compared, the processor adds the new address and status word to the Ethernet address table as a background task.

- If neither the “greater than” nor the “equal to” bit is set, it means the search argument was less than the address in the last location probed. This also indicates that no match was found.

If a destination address was compared, the processor forwards the packet, since it cannot determine the location (port A or B) of the station with that address.

If a source address was compared, the processor adds the new address and status word to the Ethernet address table as a background task.

4.5.2 Circuit Descriptions

The TLU contains the following circuits:

- Search address multiplexer
- 48-bit address comparator circuitry, which includes:

Ethernet address table memory

Bus transceivers

Ethernet search address comparator PALs

Binary search control and status PALs

Search results register

4.5.3 Search Address Multiplexer

The search address multiplexer (see Figure 4–10) selects the processor bus (BA01–BA13) or the search control bus (PA01–PA13) to address the Ethernet address table memory.


Four quad 2:1 multiplexers make up the multiplexer circuit. The single steering input to the multiplexer is a signal called SMUXL. This signal is derived from the DONE signal when DONE is deasserted by the search control and status PALs. SMUXL is gated to the multiplexer by the 100-ns clock signal.

4.5.4 48-Bit Address Comparator

The 48-bit address comparator performs bit-by-bit comparisons of the 48 bits that make up an Ethernet address. The PALs contain the compare and move registers (see Section 4.3.5.10) that are used for 64-bit moves (see Chapter 3).

4.5.4.1 Functional Description — The 48-bit comparator is made up of three 16-bit comparators. Figure 4–11 shows the structure of a 48-bit comparator. This illustration shows the circuits that compare the low 16 bits, the middle 16 bits, and the high 16 bits of the 48-bit Ethernet address.

The address table memory is 8K by 48 bits. Thus one location in memory accesses the entire 48-bit Ethernet address.




LKG-0772

Figure 4-11: 48-Bit Comparator

Figure 4-12 shows the structure of a 16-bit comparator. This illustration shows the circuits that perform a bit-by-bit comparison of the four nibbles that make up a 16-bit portion of the 48-bit address.

The results of each byte comparison are sent to the search control PALs, which determine the next step.



LKG-0773

Figure 4-12: 16-Bit Comparator

4.5.4.2 Address Comparator Circuit Descriptions — The address comparator is made up of the following circuits:

- Ethernet address table memory
- Bus transceivers
- Ethernet search address comparator PALs
- Binary search control and status PALs
- Search status register

4.5.4.3 Ethernet Address Table Memory — The Ethernet address table memory stores (in ascending numerical order) Ethernet addresses. The addresses can be either or both of the following:

1. Addresses of stations from which the LAN Bridge 100 unit has received packets
2. Addresses that have been down-line loaded by Remote Bridge Management Software (RBMS)

The memory is mapped as 8K by 48 bits and is made up of six 8K-by-8-bit RAM devices. The devices are addressable by the processor or by the search control PALs.

Data (48-bit addresses) can be read out of address table memory for the following purposes:

- The comparator PALs use addresses for performing address comparisons.
- The comparator PALs temporarily store addresses during table maintenance operations.
- Memory data may be used by RBMS or for Maintenance Operation Protocol (MOP) purposes.

4.5.4.4 Bus Transceivers — The bus transceivers are bidirectional and allow data to be transferred between the data bus (D00–D15) and the internal compare bus.

In normal LAN Bridge 100 operation, Ethernet addresses are transferred from the data bus to the compare bus. The addresses are latched into the comparator PALs for comparison against addresses stored in the Ethernet address table.

Bidirectional transfer capabilities allow bridge management to read addresses from the Ethernet address table.

The transfer direction is controlled by the processor BWRT control line and the SXC0, SXC1, and SXC2 control lines from the decoder.

4.5.4.5 Ethernet Search Address Comparators — The comparator PALs perform bit-by-bit comparisons between data received through the processor's data bus and data accessed from the Ethernet address table.

When the bus transceivers place the 48-bit Ethernet address onto the compare bus, each PAL latches one nibble (4 bits) of the address. As the search control and status PALs probe the address table, each comparator PAL compares its latched nibble against a corresponding nibble accessed from the address table. A "greater than," "less than," or "equal to" result is sent to the search control and status PALs.

The comparator PALs are also used during a 64-bit move for temporarily storing bits that make up the Ethernet address (see Chapter 3 for more information on 64-bit moves).

4.5.4.6 Search Control and Status PALs — The search control and status PALs have two functions. They control the binary search process once it is initiated by the processor and they write the search status to the search results register.

When the search control PALs detect that the final (highest) byte of the 48-bit Ethernet address is latched into the comparator PALs, they read the address that is on the search address bus (SA01–SA13). This address, which is used for the first probe of the Ethernet address table, is also used by the search control PALs to calculate successive addresses to be probed during the binary search of the address table.

The search control and status PALs deassert the DONE signal. This has two effects.

- It inhibits the decoder from allowing the processor to address any of the TLU hardware.
- It steers the multiplexer to select search addresses from the search control and status PALs.

The PALs continue the binary search process by monitoring the results (“greater than” or “equal to”) from the address comparator PALs. Depending on the results, the PAL determines the next step in the search process as follows:

- When the results are “greater than,” the PAL equation calculates the next address to be probed. The calculated address is placed onto the address bus.
- When the results are “less than” (actually, not “greater than” and not “equal to”), the PAL equation calculates the next address to be probed. The calculated address is placed onto the address bus.
- When the results are “equal to,” the PALs terminate the search and reassert the DONE signal. The reasserted DONE signal enables the decoder to allow the processor to read the search results register. It also steers the multiplexer to select the search address from the processor.

4.5.4.7 Search Results Register — The search results register contains the address and status of the last location probed by the search control and status PALs. The register data is latched in the search control and status PALs and is addressed through a buffer on the processor data bus.

The 16 bits of register data are written onto the PAL address lines (PA00–PA13) and PAL control lines. The register is addressable by the processor only when the binary search is not in progress (the DONE signal is deasserted). The 16 bits have the following definitions:

Table 4–3: Search Results Register Bit Descriptions

Bit	Definition
<15>	Equal to (EQ) bit. Asserted when a search results in a found entry.
<14>	Not used.
<13:01>	Search address bits corresponding to the last address probed by the search PALs. When accompanied by the EQ bit being set, the address indicates where the entry is located. When EQ is not set, the address indicates where the entry should reside.
<00>	Greater than (GT) bit. Asserted when a search does not result in a found entry and the last comparison found the search argument to be greater in magnitude than the table entry.

4.6 Power Supply

The power supply subsystem is a switching type, regulated ac-to-dc converter that uses a transformer in a half-wave transformer-coupled mode.

Power regulation is achieved by modulating the pulse width of the inverter's primary current conduction time. An increase in the input power decreases the pulse width of the inverter's primary current conduction time to lower the output voltage. Conversely, a decrease in the input power increases the pulse width of the inverter's primary current conduction time to raise the output voltage.

The voltage input range is switch selectable and can be either of the following:

- 88 Vac to 132 Vac (115 Vac nominal)
- 176 Vac to 264 Vac (230 Vac nominal)

The power supply provides the following regulated output voltages that are used by the LAN Bridge 100 circuitry:

Revision F08 and below	Revision F09 and above
+5 volts at 16 amps	+5 volts at 20 amps
+12 volts at 3 amps	+12 volts at 2.5 amps
-12 volts at 1 amp	-12 volts at 2.0 amps
Ground	Ground

5 Maintenance

5.1 Scope

The purpose of this chapter is to provide information on maintaining the LAN Bridge 100 unit. This chapter includes sections on:

- Maintenance philosophy
- Preventive maintenance
- Corrective maintenance
- LAN Bridge 100 disassembly

5.2 Maintenance Philosophy

The maintenance philosophy for the LAN Bridge 100 unit is option swap. In other words, maintenance is concerned with determining whether or not a LAN Bridge 100 unit is faulty, and replacing the LAN Bridge 100 unit when necessary.

Maintenance of the LAN Bridge 100 unit consists of corrective procedures only. Instructions for replacing faulty LAN Bridge 100 units are provided as part of these procedures.

5.2.1 Required Equipment

The controlled distribution (CD) spares kit contains a LAN Bridge 100 unit and loopback connectors. The loopback connectors allow off-line testing of the LAN Bridge 100 unit. Table 5-1 describes the CD kits that are available for local and remote LAN Bridge 100 units.

Table 5–1: Controlled Distribution Kits for the LAN Bridge 100 Unit

Part Number	Description
A2-W0948-10	Local LAN Bridge 100 Spares Kit (120 V Version) This kit contains: <ul style="list-style-type: none">• 1 LAN Bridge 100 (part no. DEBET-AA)• 2 loopback connectors (part no. 12-22196-01)
A2-W0948-11	Local LAN Bridge 100 Spares Kit (240 V Version) This kit contains: <ul style="list-style-type: none">• 1 LAN Bridge 100 (part no. DEBET-AB)• 2 loopback connectors (part no. 12-22196-01)
A2-W1043-10	Remote LAN Bridge 100 Spares Kit (120 V Version) This kit contains: <ul style="list-style-type: none">• 1 LAN Bridge 100 (part no. DEBET-RA)• 1 loopback connector (part no. 12-22196-01)• 1 fiber-optic loopback connector (part no. 29-25037-01)
A2-W1043-11	Remote LAN Bridge 100 Spares Kit (240 V Version) This kit contains: <ul style="list-style-type: none">• 1 LAN Bridge 100 (part no. DEBET-RB)• 1 loopback connector (part no. 12-22196-01)• 1 fiber-optic loopback connector (part no. 29-25037-01)

5.2.2 Optional Equipment

The H4080 test connector may be helpful in performing some corrective maintenance procedures. Note that this test device is not supplied with the CD kit.

The H4080 test connector is a transceiver that is connected to a short length of coaxial cable. This test connector can replace an on-line transceiver for off-line self-testing of the LAN Bridge 100 unit.


NOTE

The H4000-TA (or H4000-TB) Ethernet transceiver testers cannot be used to test the data path through a LAN Bridge 100 unit because all H4000-Tx testers have the same Ethernet address. This prevents the LAN Bridge 100 unit from forwarding the test packets. However, testers may be used to test the links that normally connect to the ports of the bridge.

5.3 Preventive Maintenance

There are no preventive maintenance (PM) procedures for the LAN Bridge 100 unit. However, the LAN Bridge 100 self-test should be run when network PM is performed. To run self-test, perform the following steps:

1. Start the self-test either by interrupting power or by remote bridge management intervention.
2. Observe the status LEDs shown in Figure 5-1. The Self-test OK status LED should light within 10 seconds after self-test is started. If loop-detection conditions are not met, the On-line status LED should light within 20 seconds after the Self-test OK status LED lights. If a failure is noted, refer to Figure 5-2.



MKV85-2776

Figure 5-1: LAN Bridge 100 Status LEDs

5.4 Corrective Maintenance

Corrective maintenance should be performed when there are indications that a LAN Bridge 100 unit is faulty. Corrective maintenance comprises verification and replacement of the faulty LAN Bridge 100 unit. This is done by analyzing the results of self-test and isolating the LAN Bridge 100 unit from the network by using loopback connectors.

5.4.1 Troubleshooting Tips

The following tips on troubleshooting a LAN Bridge 100 unit are also referenced in the troubleshooting flowchart (Figure 5-2).

- Understand the symptoms of the problem before starting to troubleshoot. This can help prevent misinterpretation of the symptoms.
- Always look for the obvious problems first, such as the following:
 - Loose cable connections: Make sure that transceiver cables are locked in place with the slide latches. Make sure that fiber-optic cables are not misthreaded, that they are screwed firmly together, and that transmit and receive cables are not interchanged. (Cable connections are illustrated in Chapter 2.)
 - Bridge switches set incorrectly (see Chapter 1).
 - Power cord not plugged in.
 - Incorrect power or no power at the electrical outlet.
 - Circuit breaker (located on the bridge's I/O panel) tripped.
 - Note whether the fan is running. If it is not, power problems may exist.
 - Operator error.
- If the self-test fails, it automatically executes again after 15 seconds. It is not necessary to cycle bridge power.
- Test communication through the bridge by attempting to communicate to a node on the other side of the bridge. VMS systems use the SET HOST command to establish a logical link.
- Consider possible environmental problems, such as power fluctuations, high ambient temperature, interference from other equipment, and so on.
- When down-line load is enabled using the hardware switch or RBMS, ensure that a load host is established for down-line loading the appropriate LAN Traffic Monitor software. For software installation details, refer to the *LAN Traffic Monitor Installation Guide*.

5.4.2 Fault Diagnosis


This section includes a troubleshooting flowchart (Figure 5–2) that is designed to isolate problems and identify a faulty LAN Bridge 100 unit, associated transceivers, transceiver cables, or fiber-optic cables.

It is possible that a fault exists in the transceiver or other equipment electrically close to the LAN Bridge 100 unit. In some cases, this type of fault initially appears to be in the LAN Bridge 100 unit. However, careful execution of the troubleshooting procedures in this section will either isolate the fault to the LAN Bridge 100 unit or point to other possible sources of the malfunction.

When the troubleshooting flowchart leads to a malfunction in a transceiver or remote repeater, use the appropriate documents and tools to troubleshoot those devices. These include the *H4000 Transceiver Technical Manual* (part no. EK–H4000–TM) and the *Ethernet Repeater Technical Manual* (part no. EK–DEREP–TM).


Notes

1. After troubleshooting the LAN Bridge 100 unit, be sure to reconnect all cables and to reset any switches to their correct positions.
2. When NVRAM failure occurs, the fault can be bypassed and the bridge can continue to operate. However, the bridge should be replaced as soon as possible.
 - Even when NVRAM has failed, RBMS parameters may be set over the network if S2 is placed in the down position. However, the bridge reply to RBMS commands indicates partial success in setting parameters, since any set parameters will be lost in the event of subsequent hard or soft resets (including power-down periods). The parameters must be set each time a reset occurs. As long as this condition exists, the Self-test OK status LED blinks to indicate that NVRAM has failed.
 - Default parameters are used if S2 is placed in the down position. The Self-test OK status LED is steadily lit to indicate a successful self-test.



LKG-0774-87

Figure 5-2: Troubleshooting Flowchart (Sheet 1 of 8)




* in 15 seconds

† see section 5.4.2; Item 2 under "Notes."

LKG-0774-87


Figure 5-2: Troubleshooting Flowchart (Sheet 2 of 8)



* After 15 seconds
 † Refer to Troubleshooting Tips (Section 5.4.1).

LKG-0774-87

Figure 5-2: Troubleshooting Flowchart (Sheet 3 of 8)



¹ FROM HOST ON PORT A TO HOST ON PORT B


² DOWN-LINE SWITCH MAY BE ENABLED BY SWITCH 5 BEING SET OR VIA A REMOTE COMMAND BY RBMS. SEE LAN BRIDGE 100 HARDWARE INSTALLATION/OWNER'S GUIDE.

³ FLASHING TWICE EVERY 2 SECONDS INDICATES THAT THE LOAD HOST SUCCESSFULLY DOWN-LINE LOADED THE LTM LISTENER SOFTWARE IMAGE. FLASHING ONCE EACH SECOND INDICATES THAT THE LOAD HOST HAS STARTED THE LTM LISTENER SOFTWARE.

⁴ CHECK THAT THE DOWN-LINE LOAD HOST HAS BEEN SET UP. SEE LAN TRAFFIC MONITOR INSTALLATION GUIDE FOR DETAILS ON SETTING UP A LOAD HOST.


LKG-0774-87

Figure 5-2: Troubleshooting Flowchart (Sheet 4 of 8)




LKG-0774-87

Figure 5–2: Troubleshooting Flowchart (Sheet 5 of 8)



LKG-0774-87


Figure 5-2: Troubleshooting Flowchart (Sheet 6 of 8)



*AFTER 15 SECONDS

LKG-0774-87

Figure 5-2: Troubleshooting Flowchart (Sheet 7 of 8)



*AFTER 15 SECONDS

LKG-0774

Figure 5-2: Troubleshooting Flowchart (Sheet 8 of 8)

5.5 LAN Bridge 100 Replacement Procedures

Use the following procedures to replace a LAN Bridge 100 unit.

1. Unplug the power cord from the wall outlet and from the defective LAN Bridge 100 unit.
2. Disconnect the transceiver cables and/or the fiber-optic cable.
3. Test the replacement LAN Bridge 100 unit by installing loopback connectors and connecting power to the LAN Bridge 100 unit.
4. If the LAN Bridge 100 unit is rack mounted or wall mounted, go to step 5. If the LAN Bridge 100 unit is a desk-top unit, replace it with another LAN Bridge 100 unit and go to step 8.
5. Support the LAN Bridge 100 unit and remove the screws holding the mounting brackets to the rack or wall.
6. Remove the brackets from the LAN Bridge 100 unit.
7. Install a replacement LAN Bridge 100 unit in the rack or on the wall according to the instructions in Chapter 2.
8. Connect the transceiver cables and/or fiber-optic cable to the replacement LAN Bridge 100 unit according to the instructions in Chapter 2.

5.6 Bridge Disassembly

This section describes the disassembly procedure for the LAN Bridge 100 unit. Removal procedures for the following parts are provided:

- Plastic enclosure
- Chassis cover
- Fiber-optic interface
- Power supply
- Logic board

WARNING

To prevent electrical shock and damage to components, disconnect the power cord from the LAN Bridge 100 unit before opening the chassis.


The instructions in this section assume that:

- All external cables to the bridge have been removed. (Cables should be marked for proper replacement.)
- The LAN Bridge 100 unit has been removed from its rack or wall mounting.

5.6.1 Plastic Enclosure Removal

To remove the plastic enclosure (see Figures 5-3 and 5-4), follow these steps:

1. Remove the 4 screws (see Figure 5-3) that secure the rubber feet to the bottom of the plastic enclosure. These screws release the top plastic of the enclosure and the two side pieces.
2. Remove the 4 screws (see Figure 5-4) that hold the bottom of the enclosure to the chassis.



LKG-0775

Figure 5-3: Removing the Top Plastic Enclosure and Side Pieces




Figure 5-4: Removing the Bottom Plastic Enclosure

LKG-0776

5.6.2 Removing the Chassis Cover

To remove the chassis cover for Revision F08 and below, see Figure 5–5 and follow these steps.


NOTE

To identify the revision level on the unit, the label that contains the Part Number (PN) and Serial Number (SN) also lists the revision level as Rev. Older revision units are labeled as Rev F08 and below. The newer units are designated as Rev F09 and above.

1. Remove the top and bottom plastic enclosures and side pieces (refer to Section 5.6.1).
2. Remove 24 chassis cover screws as follows (see Figure 5–5):
 - Remove 10 screws from the top of the chassis.
 - Remove 14 screws (7 on each end of the chassis).
3. Lift off the chassis cover (see Figure 5–5).


To remove the chassis cover for Revision F09 and above, see Figure 5–6 and follow these steps:

1. Remove the top and bottom plastic enclosures and side pieces (refer to Section 5.6.1).
2. Remove 14 chassis cover screws as follows (see Figure 5–6):
 - Remove 2 screws from the top of the chassis.
 - Remove 12 screws (6 on each end of the chassis).
3. Lift off the chassis cover (see Figure 5–6).



LKG-0777

Figure 5-5: Removing the Chassis Cover for Revision F08 and Below



MKV88-1686

Figure 5-6: Various Removal Procedures for Revision F09 and Above


5.6.3 Removing the Fiber-Optic Interface

To remove the fiber-optic interface, follow these steps:

1. Remove the top and bottom plastic enclosures and side pieces (refer to Section 5.6.1).
2. Remove the chassis cover (refer to Section 5.6.2).
3. Remove the 2 screws (see Figure 5-7) that hold the interface bracket to the LAN Bridge 100 chassis.
4. While releasing the standoff clips that hold the interface board in place (see Figure 5-7), gently remove the interface board and bracket from its connector on the logic board, as shown in Figure 5-8.

NOTE

A fiber-optic interface is present only in a remote LAN Bridge 100 unit.



LKG-0779


Figure 5-7: Fiber-Optic Interface Retaining Hardware for Revision F08 and Below

NOTE

See Figure 5-6 for fiber-optic interface retaining hardware for Revision F09 and above.

NOTE

Figure 5-8 shows the angle at which the fiber-optic interface should be removed from the LAN Bridge 100 unit.



LKG-1228-87

Figure 5-8: Removing the Fiber-Optic Interface for Revision F08 and Below

NOTE

See Figure 5-6 for removing the fiber-optic interface for Revision F09 and above.

5.6.4 Removing the Power Supply


To remove the power supply, follow these steps.

Power supply removal procedure for Revision F08 and below:

1. Remove the top and bottom plastic enclosures and side pieces (refer to Section 5.6.1).
2. Remove the chassis cover (refer to Section 5.6.2).
3. Unplug the 3 connectors (A) shown in Figure 5–9.
4. Remove the 2 retaining screws (B) that secure the power supply bracket to the LAN Bridge 100 I/O panel.
5. Remove the fan power harness from the plastic retaining clips (A) shown in Figure 5–10.
6. Loosen the 2 captive screws (B) on the power supply bracket until they pop up.
7. Carefully lift the power supply assembly out of the chassis. Figure 5–12 shows the power supply assembly after it is removed from the chassis.

Power supply removal procedure for Revision F09 and above:

1. Remove the top and bottom plastic enclosures and side pieces (refer to Section 5.6.1).
2. Remove the chassis cover (refer to Section 5.6.2).
3. Remove the 6 screws that secure the power supply module to the top chassis cover. See Figure 5–11.
4. Remove the fan power harness from the power supply module.
5. Remove the Ground wire from the top chassis cover by removing the retainer nut.
6. Carefully lift the power supply assembly out of the chassis. Figure 5–13 shows the power supply assembly after it is removed from the chassis.




LKG-1227-87

Figure 5-9: Removing the Power Supply Connectors and Retaining Screws for Revision F08 and Below


NOTE

See Figure 5-6 for removing the power supply connectors and retaining screws for Revision F09 and above.




LKG-0780

Figure 5-10: Removing the Power Supply for Revision F08 and Below




MKV88-1688

Figure 5-11: Removing the Power Supply for Revision F09 and Above



LKG-0781

Figure 5-12: Power Supply Assembly Removed from the Chassis (Revision F08 and Below)



MKV88-1687

Figure 5-13: Power Supply Assembly Removed from the Chassis (Revision F09 and Above)

5.6.5 Removing the Logic Module


To remove the logic module (see Figure 5–14), follow these steps:

1. Remove the top and bottom plastic enclosures and side pieces (refer to Section 5.6.1).
2. Remove the chassis cover (refer to Section 5.6.2). Remove the power supply (refer to Section 5.6.4).
3. If necessary, remove the fiber-optic interface (refer to Section 5.6.3).
4. Unplug the remaining interface connector (A) shown in Figure 5–14. To remove the interface connectors completely, loosen the captive screw (see Figure 5–14, inset) on each connector bracket until it pops up. Lift each interface connector out of the chassis.

NOTE

For Revision F09 and above, the interface connector next to the ac power area is fix-mounted and, therefore, cannot be removed.

5. Remove the retaining screw (B) that holds the logic module to the LAN Bridge 100 chassis.
6. While releasing the standoff clips (C) that hold the logic module in place, carefully lift the logic module out of the chassis (see Figure 5–14).



LKG-0782

Figure 5-14: Locations of Logic Module Retaining Hardware

A

Glossary

Baseband Ethernet

An unmodulated, single-channel, CSMA/CD network. One station at a time can be transmitting information on the medium without disruption. (*See also* CSMA/CD.)

Bridge physical address

A unique, 48-bit address that distinguishes a bridge from all other addressable network devices (including other bridges).

Broadband Ethernet

A single CSMA/CD channel, modulated to occupy a frequency bandwidth channel on a broadband coaxial medium.

CSMA/CD

Carrier-sense, multiple-access with collision detection. This is the generic term for the class of link management procedure used by Ethernet.

Cost

See Path Cost.

Data link

A logical connection between two stations on the same circuit.

Data Link layer

The higher of the two lowest layers in the International Standards Organization model for open systems interconnection. The Data Link layer implements a medium-independent link-level communication facility on top of the physical channel provided by the Physical layer.

Descriptor ring

A circular queue of tasks that the LANCE and microprocessor use to point either to buffer space for storing received packets or to packets in packet memory that are to be transmitted.

Designated bridge

The bridge on a local area network that has the lowest path cost to the root or the lowest ID if more than one bridge has the lowest path cost. (*See also* Path cost.) The ID consists of the physical address (which is fixed) and a priority prefix (which can be set using RBMS).

Destination address

A field in an Ethernet packet that consists of the first 6 bytes (48 bits) of the packet once the preamble and start bit have been removed. This field identifies the destination of the packet.

Down-line load

The process of sending a software image from a load host to the bridge.

Extended LAN

See Extended local area network.

Extended local area network

A LAN whose capabilities, such as length and number of stations, are extended beyond normal Ethernet or IEEE 802.3 LAN limitations through the use of certain devices such as the LAN Bridge 100 unit.

Fiber-optic link

A link between devices whose transmission medium is made up of optical fibers.

Filtering

The LAN Bridge 100 decision not to forward a packet from one side of a bridge to the other side.

Forwarding

The LAN Bridge 100 decision to pass a packet from one side of a bridge to the other side.

Hello message

A special message transmitted among bridges in an extended LAN and used for loop detection purposes. The message is transmitted as part of the spanning tree algorithm.

IEEE 802.3

A standard of the Institute of Electrical and Electronic Engineers that defines the CSMA/CD Media Access Method and Physical Layer Specifications of a network.

LAN

See Local area network.

LAN backbone

The central segment in a LAN configuration to which multiple segments can be connected using bridges. A typical use for a configuration with a LAN backbone might be in a multistory office building. A vertical backbone could connect to segments that serve individual floors in the building.

LAN Traffic Monitor (LTM)

An optional mode of LAN Bridge 100 operation that provides data about network traffic. A load host down-line loads the LTM software image to the LAN Bridge 100 unit.

LANCE

Local Area Network Controller for Ethernet. A VLSI (very-large-scale-integration) device that provides data link functions for Ethernet controllers.

Local area network

A network in which communications are limited to a moderately sized geographic area, such as a single office building, a warehouse, or a campus. The maximum length of an Ethernet LAN is 2800 meters (9184 feet). A LAN depends upon a communications medium of moderate to high data rate (10 megabits/second for Ethernet) and normally operates with a consistently low error rate.

Local bridge

A bridge that can connect LANs separated by not more than 100 meters. The distance from the local bridge to either LAN cannot exceed the maximum allowable transceiver-cable length of 50 meters (164 feet).

Loop detection

An implementation of a spanning tree algorithm that ensures a single data path between any two stations on an extended LAN.

LTM

See LAN Traffic Monitor

MOP

Maintenance Operation Protocol. A maintenance protocol that is specified in the *Digital Network Architecture (DNA) Low-Level Maintenance Operations Architectural Specification*.

Multicast message

A message that is addressed to a group of logically related stations.

NI subsystem

Network interconnect subsystem. The subsystem in a LAN Bridge 100 unit that receives and stores packets from and transmits packets to LANs that are connected to the bridge.

Node

See Station.

Packet

The basic Ethernet network message unit that is made up of a preamble and data stream. The data stream is preceded by a preamble and includes the destination address and source address fields. The minimum packet size is 512 bits (64 octets). The maximum packet-size length is 12,144 bits (1518 octets). Packets must be preceded by the interpacket gap of at least 9.6 microseconds.

Path cost

The sum of the circuit costs along a path between two stations. A circuit cost is a positive integer value associated with use of a given circuit.

Physical address

The unique, 48-bit address value associated with a given station on the network. An Ethernet physical address is defined to be distinct from all other physical addresses on all Ethernet networks.

Physical layer

The lower of the two lowest layers in the International Standards Organization model for open systems interconnection. The Physical layer is implemented by the physical channel using a specified medium. The Physical layer insulates the Data Link layer from medium-dependent physical characteristics such as baseband, broadband, or fiber-optic technologies.

Processor subsystem

The subsystem in a LAN Bridge 100 unit that controls overall bridge operation.

RBMS

Remote Bridge Management Software. An optional, VMS-layered software package residing on a station that is used to monitor and control LAN Bridge 100 units in an extended LAN.

Remote bridge

A LAN Bridge 100 unit that also converts electrical signals to optical signals that are suitable for transmission over a fiber-optic cable. Two remote bridges and a fiber-optic cable can be used to connect LANs separated by more than 100 meters (328 feet). The maximum length of the fiber-optic cable is limited by its attenuation characteristics. The distance from either remote bridge to its associated LAN cannot exceed the maximum allowable transceiver-cable length of 50 meters (164 feet).

Repeater

A device used in a LAN to extend the length, topology, or interconnectivity of the transmission medium beyond that imposed by a single transmission segment. A segment can be up to 500 meters (1640 feet) in length.

Root

The bridge in an extended LAN that has the lowest ID. The ID consists of the physical address (which is fixed) and the root priority prefix (which can be set using RBMS).

SIA

Serial interface adapter. A VLSI (very-large-scale-integration) device that provides the electrical interface between a LANCE and a transceiver cable. (*See also* LANCE.)

Source address

A field in an Ethernet packet that encompasses the second 6 bytes (48 bits) of the packet once the preamble and start bit have been removed. This field identifies the source of the packet.

Spanning tree algorithm

An algorithm used to ensure a single data path between any two stations on an extended LAN.

Station

A single addressable site on an Ethernet network, generally implemented as a computer and appropriate peripherals and connected to the network through a controller and transceiver. Also called a node.

Station address

See Physical address.

TLU

See Table lookup subsystem.

Table lookup subsystem

The subsystem in a LAN Bridge 100 unit that performs address comparisons. The results of comparisons are used by the processor to help determine whether packets should be forwarded or filtered.

Transceiver

The portion of the Physical layer implementation that connects directly to the coaxial cable and provides both the electronics that send and receive the encoded signals on the cable and the required electrical isolation.

Transceiver cable

A four-pair, shielded cable used to connect a transceiver to the LAN Bridge 100 unit.

Troll feature

A feature in a LAN Bridge 100 unit that is controlled by RBMS. This feature can prevent stations on one side of the bridge from sending packets to some stations on the opposite side of the bridge. (*See also* RBMS.)

B

Fiber-Optic Link Analysis

B.1 Introduction

This appendix discusses the general tradeoffs in a fiber-optic system and identifies and defines the characteristics of the components in such a system. Since a system is composed of many independent entities, it is possible to make a system work while violating one or more of the recommendations. Digital Equipment Corporation assumes no responsibility for providing engineering assistance if failure occurs in such cases. If the guidelines are followed properly, the system will perform to the engineering specifications. If the installer or consultant wishes to ignore the published limits and use actual measurements in forming a fiber-optic system, he or she does so with the understanding that the system may work marginally, may fail prematurely, or may fail entirely.

NOTE

Guidelines should be used to aid in determining the feasibility of a system. Following the guidelines minimizes risk.

The material provided in this appendix assumes that the reader is familiar with fiber-optic technology and has been trained in proper fiber-optic installation and maintenance techniques.

A clear understanding of the guidelines included in this appendix is essential for successful link installation, especially where many variables are involved.

WARNING

Be especially careful when working with fiber-optic cables and devices.

When exposed, optical fibers (contained within the fiber-optic cable) are extremely brittle and fragments from the fiber can easily penetrate the skin or eyes. Wear protective goggles and clothing when working with the optical fibers.

Some fiber-optic equipment may emit laser light that can injure your eyes. Never look into an optical fiber.

B.2 System Conventions

This section describes the conventions used to develop the data included in this appendix.

B.2.1 Absolute and Relative Measurements


Absolute values of physical quantities in a fiber-optic system require a laboratory environment and sophisticated laboratory techniques. It is assumed that the system under consideration has been thoroughly characterized, that parameters have been identified and specified, and that the sensitivities are understood.

Relative measurements such as end-to-end cable loss or connector loss can be done in the field. Since losses are relative in nature, the decibel (dB) is the common unit of measure rather than an absolute value such as the microwatt. The system designer must use the allowable system loss (relative measurements) as a guideline for system implementation.

NOTE

The actual power levels of emitted light and detected light in Ethernet devices are not essential in the planning and prediction of a fiber-optic network interconnect. The key measurement is the allowable system loss.

You should assume that the product is a certified product that meets the minimum specifications. The equation shown in Figure B-1 is the basis for the design engineer's system limits. The key parameter is the allowable loss, not the absolute values of light intensity.




LKG-0783

Figure B-1: Fiber-Optic System Responsibilities

In fiber-optic systems characterized by Digital Equipment Corporation, the range of absolute values of light has been determined by engineering and is the amount of light in the fiber just after the launch connector specified as **the amount of light in the fiber just after the launch connector** (see Section B.3.1).

Light loss through the last connector of the link is negligible since fiber-to-receiver coupling losses are minimal; unless the connector is damaged or the fiber end is marred, smudged, or dirty (see Section B.3.3). Therefore, the connectors on both ends of the link should not be considered in any loss calculations. They have been accounted for in the issuance of the allowable system loss number. Figure B-2 identifies the portion of the fiber system where allowable loss can occur.



LKG-1226-87

Figure B-2: Area of Network Consultant's Concern — Allowable System Loss

B.2.2 Worst-Case Definition

Parameters specified in this appendix represent the worst case unless otherwise specified. That is, all parameters have been characterized and the point specified as worst case is 2-sigma (two standard deviations from the normal data distribution). 2-sigma implies that 98% of the values obtained were better than the given worst-case number.

B.3 Coupling Considerations

There are three types of optical coupling in fiber-optic systems. Each type has a different coupling mechanism; therefore, coupling losses that are specified for each type cannot be interchanged.

The three types of optical coupling are as shown in Table B-1.

Table B-1: Types of Optical Coupling

Type of Coupling	Measurement Expertise
LED (emitter)-to-fiber	Laboratory
Fiber-to-fiber	Field
Fiber-to-pin diode (receiver)	Laboratory


IMPORTANT

It cannot be overemphasized that the quality of the polished ends of the fiber is the major cause of connector loss. Connector loss is significantly increased by dirt, scratches, finger smudges, and so on. Keep the connector ends clean.

B.3.1 LED (Emitter)-to-Fiber

Emitter-to-fiber measurements are confined to the laboratory environment. There is no field method established for determining quantitative data regarding the LED-to-fiber coupling. When using a characterized fiber type, the engineering and manufacturing groups have ensured that minimum coupling specifications are met.

The LED emits a spherical shaped pattern (see Figure B-3) that may vary from the form $r(\cos\theta)^m$ to $r[(\cos\theta)^m]$. The value of m is typically from 25 to 50 in the case of a microlensed LED. In the case of a laser, m can be as high as 200. The coupling depends on mechanical alignment of the actual LED die and the core of the fiber, the emission and acceptance angle of the LED and fiber, and the differences in refractive indexes of the materials between the die and the fiber.



LKG-0785

Figure B-3: LED Coupling into a Fiber

B.3.1.1 LED-to-Fiber Coupling Derating Factors — Tables B-2 and B-3 identify the Allowable System Loss that can be assumed for those systems that use DIGITAL fiber-optic daughterboard 54-16053-01 and stainless steel Amphenol 906 type SMA connectors. The 54-16053-01 daughterboard has two revisions that give differing power budgets. Be sure to check the variation of the product and follow the appropriate table based on the variation you are configuring.

The amount of Allowable System Loss must be compared to the sum of the fiber link losses in order to determine whether or not the link will function. The losses that must be considered are listed below. Methods for calculating system budgets are outlined in Section B.8.

Cable Attenuation	dB/km × cable length
Splice Loss	dB/each × # of splices
Connector Loss	dB/connector × # of additional connectors
Bandwidth Derating	0.5 dB/km × cable length
Margin	1.5 dB for system margin

NOTE

The 1 dB for LED aging and .7 dB for temperature-caused variations are already included as part of the fixed losses and do not have to be included when making budget calculations.

Table B-2: LED-to-Fiber Coupling Derating Factors for -RC/-RD Variations

(Using 54-16053 Revision A2 or lower)

Fiber Type	Dimensions	Derating Factor	Allowable System Loss
Corning 1508	100/140 μm	0 dB	12.5 dB
Corning 1519	85/125 μm	-2.4 dB	10.1 dB
AT&T Multimode	62.5/125 μm	-4.2 dB	8.3 dB
Corning 1509	62.5/125 μm	-4.2 dB	8.3 dB
Corning 1517	50/125 μm	-9.4 dB*	3.1 dB
AT&T Multimode	50/125 μm	-9.4 dB*	3.1 dB

*The greater loss incurred by these cable types results from a smaller numerical aperture and a smaller core diameter.

Table B-3: LED-to-Fiber Coupling Derating Factors for -RH/-RJ Variations

(Using 54-16053 Revision B2 or higher)

Fiber Type	Dimensions	Derating Factor	Allowable System Loss
Corning 1508	100/140 μm	0 dB	16.0 dB**
Corning 1519	85/125 μm	-2.5 dB	13.5 dB
AT&T Multimode	62.5/125 μm	-4.0 dB	12.0 dB
Corning 1509	62.5/125 μm	-4.0 dB	12.0 dB
Corning 1517	50/125 μm	-8.0 dB*	8.0 dB
AT&T Multimode	50/125 μm	-8.0 dB*	8.0 dB

* The greater loss incurred by these cable types results from a smaller numerical aperture and a smaller core diameter.

DEBET and DEREK (-RH/-RJ versions **ONLY) links using 100/140 micron fiber, must incur a minimum of 3 dB loss. Refer to section B.9 for use of the fiber-optic attenuator if your system does not have at least 3 dB worth of loss.

B.3.2 Fiber-to-Fiber Coupling

Fiber-to-fiber coupling occurs in two instances: a splice or a bulkhead/barrel connection. Measurement of fiber attenuation, splice loss, and fiber-to-fiber barrel connections are the only measurements that can be done in the field.

NOTE

If possible, do not mix fiber types or sizes in a system. In a system with mixed fiber sizes, all fibers must be considered to act as the smallest fiber size in the system.

Losses occur in splices and barrel connections because of mechanical misalignment, mismatch of core diameters, mismatch of refractive index profile, and mismatch of numerical aperture. Statistical data is available for most of these parameters, but worst-case numbers should be used in predicting performance of a system that is going to be installed.

NOTE

It is extremely difficult to measure attenuation in a system with mixed fiber diameters. Light traveling from a small fiber to a larger one can yield erroneous test readings, such as apparent light increases instead of attenuations (decreases).

B.3.2.1 Splice and Barrel Connector Losses — In typical fiber-optic interconnect systems, there may be several kinds of cables (using the same fiber type), several sections of cables, and usually some type of a junction box strategy. This implies that there are barrel connections and splices in the fiber-optic cable.

Table B-4 shows the loss in decibels for the connectors and splices that have either been characterized by engineering or are available and can be characterized by the installer.

IMPORTANT

The values in the following table are results that can be achieved only by a person who is trained and experienced in using proper tools and correct procedures. The values were obtained under laboratory conditions.

Table B-4: Splice and Barrel Connector Losses

Connector or Splice Type	Typical Loss in dB	Worst Case Loss in dB (Note 1)
<i>Barrel Connectors</i>		
100/140 Amphenol type 906 (stainless steel)	0.8	1.5
50/125 Amphenol type 906 (stainless steel)	1.5	Note 2
62.5/125, 85/125 Amphenol type 906	1.5	Note 2
50/125 AT&T ST type	1.0	1.5
62.5/125 AT&T ST type	0.8	1.2
85/125 AT&T ST type	Note 2	Note 2
100/140 AT&T ST type	Note 2	Note 2
<i>Splices</i>		
Elastomeric	0.1	0.2 (Note 3)
Capillary tube	0.1	0.2 (Note 3)
Fusion	.05	0.1 (Note 3)

NOTES:

1. Assumes that the barrel connector is close to the LED source. This is the worst-loss condition.
2. Consult vendor information.
3. Loss assumes the same lot (reel) of fiber is used. Splicing different lots of fiber yields results based on the fiber differences rather than on the mechanics of the splice.

B.3.2.2 Guidelines for Connectors and Splices on Ethernet Products — This section provides guidelines that should be followed when connecting or splicing fiber-optic cable for use with Ethernet products.

- All fibers must be terminated in an Amphenol type 906 stainless steel connector for attachment of the cable to the Ethernet device.
- For 100/140 fiber systems, splices or barrel connections can be used. Use Amphenol type 906 stainless steel barrel connectors and follow the appropriate connection procedures for the connectors and alignment sleeves.
- For 62.5/125 or 85/125 fiber systems, splicing is preferred for intermediate connections, but Amphenol type 906 or AT&T type ST connectors may be used for these connections.
- For 50/125 fiber systems, all intermediate connections should be spliced.
- Be aware of the types of misalignment or separation that are illustrated in Figure B-4.

NOTE

No guarantee of a specified attenuation can be made for splicing or barrel connecting fibers with different core diameters.

The ideal treatment of loss from one fiber size to another is represented by the following equations. These generally do not apply to real-world situations since the loss depends on the spot size of the light in the core, not on the core diameter.

Ideal cases:

1. Loss due to diameter mismatch:


$$20 \log \left[\frac{\text{diameter 1}}{\text{diameter 2}} \right]$$

where diameter is the diameter of the fiber core.


2. Loss due to mismatch of numerical aperture:

$$20 \log \left[\frac{\text{NA 1}}{\text{NA 2}} \right]$$


Where NA is the numerical aperture of the fiber.



LATERAL MISALIGNMENT



END SEPARATION



ANGULAR MISALIGNMENT

LKG-0786

Figure B-4: Types of Fiber Misalignment

B.3.2.3 Typical Cable Attenuation Values — Table B-5 contains typical cable attenuations. Each type of cable may have a different value for attenuation based on the application, cable type, and vendor guarantee for the installation environment.

This table should be used as a reference for typical values only.

Table B-5: Typical Cable Attenuation

Fiber Type	Typical Cable Type	Typical Value
Corning 1508	100/140; DEC BN25B-xx	4.5 dB/km*
Corning 1519	85/125; Siecor loose tube**	3.5 dB/km
Corning 1519	85/125; Siecor tight buffer	4.0 dB/km
Corning 1509 or AT&T Multimode	62.5/125; various loose tube	3.0 dB/km
Corning 1517 or AT&T Multimode	50/125; various loose tube	2.5 dB/km
Corning 1517 or AT&T Multimode	50/125; tight buffer tube	4.0 dB/km

* This cable is guaranteed for 6 dB/km under specified applications for all temperature, humidity, and rated tension conditions.

**Loose tube cable is usually very difficult to terminate with a connector. Splicing a connector cable end onto the cable is recommended.

NOTE

To achieve the maximum link distances, premium grade fiber is recommended.

B.3.2.4 Single-Window and Dual-Window Fiber — When the goal of the installation is to provide a transmission medium for an 835-nanometer (nm) system, single-window fiber can be specified but dual-window fiber is recommended. If the cable contains extra fibers, or if future changes or upgrades are anticipated, the person measuring the cable should also specify or list the loss and bandwidth of the fiber at 1300 nm (the second window). Dual-window applications imply a high bandwidth and distance (200-500 MHz*km for LED applications). Similarly, low losses are implied (<2 dB/km) at this operating wavelength. Current DIGITAL general-purpose cable, BNE25B-xx, uses a dual-window fiber.

B.3.2.5 Single-Mode Fiber — Currently, single-mode fibers are for laser/telephony applications. Telephony and computer data communications are two entirely different communications systems. There is no practical method currently established that uses single-mode fiber in data communications. LEDs are the transmission device of choice for optimal cost and system performance.

B.3.3 Fiber-to-Pin Diode (Receiver)

The lowest-loss coupling type is fiber-to-receiver. Generally the receiver area is larger than the spot size of the light exiting the fiber (see Figure B-5). Losses are due to large-scale mechanical misalignment, surface quality of the polished fiber end, fresnel losses due to different materials with different refractive indexes, and efficiency of semiconductor devices. Fiber-to-receiver loss is accounted for in the engineering design phase and should not be a concern in the field other than for determining that the connector surface is not scratched, marred, smudged, and so on; conditions that would create greater losses than the design allows for.




Figure B-5: Fiber-to-Receiver Coupling


B.4 Numerical Aperture and Coupling Diameters

Coupling losses are not easily quantified in absolute terms. The following discussion describes the apparent variability in loss measurements and the range (optimistic to pessimistic) of the measurements.

When the connection is close to the source end, the fiber core is usually completely filled with light, resulting in more loss from misalignment. When the connection is at the far end, the center of the core is only partially filled with light, resulting in less loss from mechanical misalignment (see Figure B-6).

NOTE

Losses near the source are usually higher and give pessimistic results. Measurement of losses at distances greater than 100 meters (328 feet) from the source usually gives optimistic results.



LKG-1225-87


Figure B-6: Near-End and Far-End Coupling Losses

B.4.1 Bending Losses

Typical cables will incur an additional -0.5 dB loss because of changes in fiber bend radius. This loss occurs because the light in the outer edge of the core is more easily lost to the cladding, especially if the fiber is bent. In most systems, it is wise to assume the cable will be installed with several bends; therefore, bend losses are included as part of the the fixed cable losses and do not have to be separately included as part of the loss calculations (see Figure B-7).

NOTE

The manufacturer's recommendation for minimum bend radius must be adhered to for this to be valid.



MKV88-1839

Figure B-7: Minimum Bend Radius

B.5 System Losses and Variability

At the engineering level, a variety of factors must be quantified and specified. Many of these factors are not measurable in the field. Listed below are a number of factors which may incur losses in a system budget. These are calculated and taken into consideration for determining the available budget, and do not have to be recalculated when installing a fiber link. They are included here for reference purposes only.

- LED Specification
- LED Aging — Temperature Variations
- LED Drive — Circuit Variations
- Mechanical Coupling
- Connector Variations
- Fiber Variations/Polishing
- Temperature
- Cable Aging
- Connector/Polishing
- Coupling to Detector
- Receiver Specification
- Threshold Variations

B.6 Attenuation Measurement Procedures

This section provides three test methods that may be used to measure attenuation of a fiber-optic link. The tests performed in the field can help determine whether the attenuation (loss between the light launched into the fiber through the first connector and the light emitted at the other end of the fiber) is within specified loss limits.

Though many methods of testing fiber-optic links are available, Digital Equipment Corporation recommends using the first of the three methods described, which involves using an optical test set (FOTEC model T302D). This will allow for the accuracy necessary to achieve the greater distance capabilities of the higher power -RH/-RJ variations. The single-cable test (second method outlined), is similar to the first method, but is not the preferred method for testing all fiber types. This method is left here for informational purposes only. The third method is provided mainly for measuring splice losses through the fiber-optic cable.

As described in Section B.4, the value of connector-to-connector loss may change depending on the position of the connector relative to the light source and the direction in which the light is traveling. Though it is desirable to determine that a given connection is within acceptable limits, determining the value of “system loss” is the essential factor.

B.6.1 Field Measurement Procedure for Relative Power Loss

The following describes the procedure for determining the acceptance of an installed fiber-optic cable.

The test performed measures the relative loss of a system cable under test. It consists of launching a pre-set quantity of optical power into an installed cable under test, and then measuring the received power at the opposite end using an optical power meter. The value of the received power represents the loss associated with the installed cable link.

B.6.1.1 Equipment Required —

- FOTEC optical power meter equipped with ST and SMA adapters
- FOTEC optical LED source equipped with ST and SMA adapters
- 200/240 μm launch cable terminated with an ST and an SMA connector
- Connector cleaning paraphernalia
- ST and SMA barrel connector adapters

B.6.1.2 Preparation — Fiber-optic cables and connectors should always be handled with care. Cables should not be bent or twisted tightly, especially near the connectors, as damage or breakage may occur. Care should be taken to ensure that connectors are not dropped or physically impacted in any way. If this happens, the connector may become damaged, impairing optical transmission.

All connector faces should be cleaned before testing or inserting into any optical port. This is done by using Tex-wipes or similar cleaning pads dipped in alcohol to remove dust, debris, or fingerprints from the connector/fiber end. When removing and reinserting connectors, cleanliness should be checked. If it appears that there is dirt or debris on the connector endface, it should be re-cleaned.

Before any testing is performed, the equipment must be removed from the carrying case and brought to room temperature. If the temperature in the testing environment is approximately 20 degrees above or below the previous environment the test equipment came from, allow up to thirty minutes for this stabilization to occur.

B.6.1.3 Procedure — Perform the following set-up procedures

1. Install the adapter that matches the cable connector into the FOTEC optical power meter.
2. Install the opposite type of adapter on the FOTEC optical LED source (this allows for the use of the 200 μm launch cable).
3. Clean and connect the launch cable to the optical source.
4. The optical source contains a sticker that indicates the wavelength of the transmitter in that particular optical source. This wavelength is in the range of 790 to 850 nanometers (nm). Make a note of the wavelength value. This value is needed during the **Measurement Correction Calculation Procedure**.

5. To ensure that the launch cable-to-optical source connection is not disturbed, immobilize both the launch cable and optical source as follows:
 - a. Tape the optical source down to a stable surface using masking tape.
 - b. Tape the launch cable to the same surface as the optical source about 7.5 to 12.7 cm (3 to 5 in) away from the optical port of the source.

CAUTION

Take care not to bend the launch cable beyond its minimum bend radius of 12.7 cm (5 in).

B.6.1.4 Test Source Adjustment Procedure — Perform the test source adjustment procedure as follows:

1. Clean and connect the launch cable to the power meter.
2. Refer to Table B–6 and set the reading on the optical power meter to the correction setting that corresponds to fiber size of the cable being tested. Set the reading by adjusting the trim pot located on the source with a screwdriver.

NOTE

It is important that the trim pot not be readjusted or disturbed for the duration of the measurement.

Table B–6: Fiber Size Correction Settings for the Power Meter Reading

Fiber Size	Correction Setting
50 μm	15.5 dB μ
62.5 μm	11.0 dB μ
85 μm	9.3 dB μ
100 μm	7.0 dB μ

3. Disconnect the launch cable from the power meter and connect it to the cable being tested using the corresponding barrel adapter (for example, if ST connectors are being used, the ST barrel adapter must be used). The barrel adapter must be free of dirt and debris to ensure optimum coupling of power.

When using SMA connectors, attach the plastic full sleeve that goes inside the barrel adapter. The plastic sleeve must be free of dirt and debris and placed on the end of the first connector that is to be inserted into the adapter before the connector is inserted.

4. Take the power meter to the opposite (receive end) of the cable being tested.

- At the receive end of the cable, clean the connector on the cable being tested and attach it to the power meter.

NOTE

If a measurement is being made from one patch panel to another patch panel without the drop cables to the installed equipment, it will be necessary to use a jumper from the patch panel to the power meter. In this case, use a jumper with a fiber core that is larger than the fiber in the cable under test.

B.6.1.5 Measurement Correction Calculation Procedure — During this procedure, complete the End-to-End Cable Loss Calculation Worksheet. This worksheet is used to arrive at the corrected cable loss by recording the measured end-to-end cable loss and adding in correction values.

Figure B–8 is a blank End-to-End Cable Loss Calculation Worksheet. Before proceeding with the measurement correction calculation procedure, make at least one photocopy of this worksheet for each fiber optic cable to be certified.

END-TO-END CABLE LOSS WORKSHEET

MEASURED RELATIVE LOSS	(1) _____
LENGTH OF TESTED CABLE (TO NEXT .1 km)	(2) _____
WAVELENGTH CORRECTION VALUE	(3) _____
TOTAL WAVELENGTH CORRECTION (MULTIPLY LINE 2 BY LINE 3)	(4) _____
CORRECTED CABLE LOSS: SUBTRACT (SUBTRACT LINE 4 FROM LINE 1)	(5) _____

MKV88-1838

Figure B–8: End-to-End Cable Loss Calculation Worksheet

Take one copy of the blank worksheet (Figure B–8) and perform the following:

- On the **Measured Relative Loss** line of the worksheet, write the value that is displayed on the power meter.
- On the **Length of Tested Cable** line of the worksheet, write the rounded-off length of the cable under test by taking the cable length and rounding it upwards to the next tenth of a kilometer (for example; if the cable length is 1375 meters, round it off to 1.4 km)

3. Refer to Table B-7 and choose the wavelength correction value that applies to the optical source wavelength and to the Ethernet option being installed (the optical source wavelength is the wavelength value that was noted during Step 4 of the **Set-Up Procedure** in Section B.6.1.3). Write the value from Table B-7 on the **Wavelength Correction Value** line of the worksheet.

Table B-7: Wavelength Correction Values

Wavelength	DEBET-RC/RD/RH/RJ
790	-.4
795	-.3
800	-.2
805	-.1
810	.0
815	.1
820	.2
825	.3
830	.4
835	.45
840	.5
845	.6
850	.7

4. Multiply the **Wavelength Correction Value** line by the **Length of Tested Cable** line and write the result on the **Total Wavelength Correction** line of the worksheet.
5. Calculate the corrected output power by subtracting the **Total Wavelength Correction** line value from the **Measured Relative Loss** line value and write the result in the **Corrected Cable Loss** line of the worksheet.

NOTE

To subtract two negative values, drop the minus sign (–) on both numbers, subtract the smaller number from the larger value, and add a minus (–) to the result. For example –.5 subtracted from –7.4 becomes 7.4 minus .5 which results in –6.9 with the minus sign added.

To subtract a positive number from a negative number, drop the minus sign (–) from the negative number, add the two numbers together, and add a minus sign (–) to the result. For example, .5 subtracted from –7.4 becomes 7.4 plus .5 which results in –7.9 with the minus sign added.

B.6.1.6 System Certification Procedure — Table B-8 defines the minimum and recommended maximum loss values for fiber-optic cable by fiber size and Ethernet option.

Table B-8: Cable Certification Values

Fiber Size	Minimum Loss*	Recommended Maximum Loss**
<i>DEREP/DEBET-RC/RD</i>		
50 μm	None	-1.1 dB
62.5 μm	None	-6.3 dB
85 μm	None	-8.1 dB
100 μm	None	-10.5 dB
<i>DEREP/DEBET-RH/RJ</i>		
50 μm	None	-6.0 dB
62.5 μm	None	-10.0 dB
85 μm	None	-11.5 dB
100 μm	3 dB	-14.0 dB

* If the relative loss is less than this value, an optical attenuator must be installed on the transmit end of the system cable (see **Section B.9**).

**This recommended maximum loss value allows for 1 dB of additional loss for future repairs.

Compare the value on the **Corrected Cable Loss** line of the worksheet to the values defined in Table B-8. If the Worksheet value is equal or less than the values from Table B-8, the installation is acceptable.

If the planning and design of the cable plant was executed correctly, allowances should have been made for splice loss, connector coupling, and future maintenance of the installation. The recommended maximum loss in Table B-8 represents the loss that allows for future maintenance of the cable plant. The future maintenance can include splices or swapping of system boards.

B.6.2 Optical Test Set (Single-Cable) Method

This method uses an optical test set to measure optical attenuation through a fiber-optic link. A reference cable (1 meter in length) is used to calibrate the light meter relative to the light source. Once calibration is complete, the reference cable should not be disconnected from the light source.

Use the following procedure to measure attenuation of the fiber-optic link with an optical test set.

NOTE

Any reference or adapter cables used in the test procedures must have the same type of optical fiber (such as 100/140, 62.5/125, and so on) as the cable to be tested. In addition, the connectors used at the ends of the reference cable must match both the test equipment connector and the connector of the cable under test. Use of a reference cable that has a different fiber type than the cable under test yields inaccurate measurements.

1. Calibrate the test configuration as follows (see Figure B-9):
 - a. Attach the necessary connector adapters to the light source and the light meter.
 - b. Connect the reference cable to the light source and turn on the light source. Allow the light source several minutes to stabilize.
 - c. Connect the light meter to the other end of the reference cable.
 - d. Use the calibration knob on the light meter to set the reading to 0 dB (do not alter the calibration knob once it has been set).

NOTE

When measuring small fiber sizes such as 50/125, it may not be possible to calibrate the light meter to 0 dB. In this case, adjust the meter as close as possible to 0 dB and note the reading.

NOTE


When connecting cables to test equipment, use a half-length alignment sleeve. When connecting cables together using a barrel connector, use a full-length alignment sleeve. A procedure for using alignment sleeves is provided in the *Fiber-Optic Pocket Guide* EK-FIBOP-IS.

Be sure that the light meter is designed to measure the desired wavelength of light. That is, if the fiber system operates at 820 nm, the light meter must be designed to measure light loss at 820 nm.

- e. Disconnect the light meter from the reference cable.

NOTE

Once calibration is completed, do not disconnect the light source from the reference cable. Also do not move (tighten or loosen) the connector that is attached to the light source.



LKG-0790


Figure B-9: Calibrating the Optical Test Set

2. Measure each channel of the fiber-optic link in the direction that light will be transmitted under normal operating conditions (see Figure B-10).
 - a. Connect the light source and reference cable to the end of the cable under test that is labeled "Transmit".
 - b. Connect the light meter to the other end of the cable under test that is labeled "Receive".
 - c. Record the reading obtained from the light meter. If a calibration other than 0 dB was obtained in Step 1, subtract that number from the present reading on the light meter.

For example, if the reading obtained in Step 1 was -1.5 dB and the reading obtained in Step 2 is -11.9 dB, the corrected reading is -10.4 dB. The equation for this example is:

$$-11.9 - (-1.5) = -10.4$$

3. Repeat the measurement in Step 2 for the remaining channel(s).



LKG-0791

Figure B-10: Measuring the Cable Under Test

B.6.3 Optical Time Domain Reflectometry

This method uses an optical time domain reflectometer (OTDR) to measure the length and attenuation of a fiber-optic cable. This method may also help to identify the location and relative value of losses.

WARNING

Do not look into a fiber-optic connector or cable end while the cable is connected to an OTDR. Follow safety procedures recommended by the OTDR manufacturer.

NOTE

It is important to perform an OTDR measurement from each end (transmit and receive) of both channels. This ensures that a broken fiber at the far end of a channel is not missed (mistaken for the end of the fiber).

This procedure requires that all test equipment has been calibrated to the manufacturer's specifications within the recommended calibration period.

NOTE

Any reference or adapter cables used in the test procedures must have the same type of optical fiber (such as 100/140, 62.5/125, and so on) as the cable to be tested. In addition, the connectors used at the ends of the reference cable must match both the test equipment connector and the connector of the cable under test. Use of a reference cable that has a different fiber type than the cable under test yields inaccurate measurements.

B.6.3.1 Equipment Setup — Adjust any settings on the OTDR according to the manufacturer's instructions.

NOTE

The OTDR must measure light at the same wavelength used by the optical fiber system. Both the LAN Bridge 100 unit and the Remote Ethernet Repeater use a wavelength of 820 nm.


B.6.3.2 Link Measurement — Link measurement involves recording the readings obtained from the OTDR. A permanent record of fiber-optic link measurement may be possible using a strip chart or camera. Such records should be obtained whenever possible and stored for future reference.

Measure the fiber-optic link using the following procedure:

1. Connect the cable under test to the OTDR in accordance with the procedures outlined in the OTDR manual.

NOTE


The pigtail, if needed (see Figure B-11) must be of the same fiber type and have the same mating connectors as the cable under test and the OTDR. Some OTDRs may have a different connector than the cable under test.



LKG-0784

Figure B-11: OTDR Test Configuration

2. Measure the length and attenuation of each fiber (transmit and receive). Note the splice and connector losses (see Figure B-12). Look for abnormalities in the slope of the display.
3. Move the OTDR to the other end of the fiber system and measure the length and attenuation of each fiber (transmit and receive). Note that the splice and connector losses are similar to those obtained in Step 2.



LKG-0792

Figure B-12: OTDR Trace

B.7 Bandwidth Analysis and Derating Factors

B.7.1 Fiber Bandwidth Rating System

Bandwidth in a fiber cable is often discussed in terms of its dispersion characteristics because narrow light pulses launched into a fiber will be broadened (dispersed in time) at the fiber's end. This pulse broadening forces an increase in the amount of spacing between bits or pulses so as to avoid Inter-Symbol Interference (ISI) (overlapping of the dispersed bits). Consequently, the greater the amount of fiber dispersion, the lower the bandwidth of the fiber.

There are several types of fiber dispersion. For multimode fiber with LED light sources there are two dominant types: modal and chromatic.

1. **Modal dispersion** occurs because light modes within a fiber have different group velocities. Thus in multimode fiber, where hundreds of modes are excited, a narrow pulse at the input will have each mode or group of modes arriving at the output at slightly different times. This broadens or disperses the output pulse. Modal dispersion is one of the dominant forms of dispersion in multimode fibers with 1300-nm LED sources, thus limiting the bandwidth of such systems to 100-450 MHz·km.

A light-ray analysis supports the concept of modal dispersion. Light rays emitted from a LED each enter the fiber at a different angle and continue to travel through the fiber along separate paths. Thus, the pulse broadening is the difference between the time of the longest ray path and the time of the shortest ray path.

2. **Chromatic dispersion** occurs because each wavelength of light within a fiber travels at a slightly different speed. As a result, light from wide-spectrum sources (LEDs) is very susceptible to chromatic dispersion. The dominant form of dispersion in the 800- to 900-nm region is chromatic dispersion, thus limiting the fiber bandwidth to 25-70 MHz·km for LEDs. Chromatic dispersion lessens around the 1300-nm region, making modal dispersion the other major band-limiting factor.

The various forms of dispersion cause the fiber bandwidth to depend greatly upon the spectrum characteristics of the light source. Thus the bandwidth measurements reported have limited meaning until they are correlated with the LED spectrum data in the derived relating equations.

Fiber bandwidth is specified in units of megahertz times kilometers because both dispersion types continue to broaden (disperse) the light pulse as fiber length increases. The fiber bandwidth is traditionally specified at 1 km since it decreases with fiber length.

B.7.2 Distance Derating Factors for Ethernet Products

It is assumed that the cable is rated at 300 MHz·km (modal bandwidth) or greater. For LEDs, this translates to a fiber bandwidth of about 50 MHz at a distance of 1 km. Using cable with a higher rating does not decrease the bandwidth penalty (derating factor) enough to change the derating factor. **Bandwidth derating will be .5 dB/km.**

It is apparent that the fiber bandwidth limits cable length by attenuating the signal with bandwidth rolloff. Bandwidth rolloff is similar to the effect of a low-pass RC filter. The derating factor compensates for bandwidth attenuation by ensuring a larger signal at the input.

B.8 Suggested System Design Methodology

The worksheet shown in Figure B-13 can be used to calculate the LAN Bridge 100 budget loss for a system installation. All values needed for the worksheet are available in this document or are specified to be provided by the vendor or installer.

The fiber-optic system has a specification called allowable system loss. This is the maximum loss that can occur between the two cable connectors that attach to the receptacles of the fiber-optic interconnect devices (given in Tables B-2 and B-3). Note that the system is specified for the amount of loss **after** the light is in the cable. That is, the connectors that attach to the LED and to the fiber-optic receiver are already accounted for in the calculations of the maximum allowable system losses. **DO NOT ADD IN THEIR LOSSES AGAIN.**

Compare the sum of the losses to the allowable system losses listed in Table B-2 or B-3. The sum of the losses must be less than or equal to the allowable system losses. **MAKE SURE THE PROPER VALUE IS USED FOR THE LAN Bridge 100 OR DEREK VARIATION BEING USED.**

PRODUCT NAME =				
Fiber Type =				
Allowable System Loss = (see Tables B-2 or B-3 based on product variation)				
FACTOR	VALUE	X	QUANTITY	= LOSS
BARREL CONNECTIONS (B.3.2.1)	--- dB/conn.	X	-- # of connectors	=
SPLICES (B.3.2.1)	--- dB/splice	X	-- # of splices	=
CABLE ATTENUATION (B.3.2.3)	--- dB/km	X	--- km of fiber	=
BANDWIDTH DERATING (B.7.2)	0.5 dB/km	X	--- km of fiber	=
MARGIN	1.5 dB			= 1.5 dB
Sum of losses				=

MKV88-1799

Figure B-13: Sample Worksheet for Budget Loss Calculation

B.8.1 Example Worksheets

This section contains two sample worksheets. The first example (Figure B-14) shows sample values and loss calculations for a DEBET-RH using 62.5/125 micron fiber in a 2 km link. The second example (Figure B-15) shows system loss calculations for a DEBET-RC using 100/140 micron fiber in a 1 km link.

These loss calculations are for products that use the 54-16053-01 fiber-optic interface module. Digital Equipment Corporation products that use this interface include the following:

- LAN Bridge 100 (DEBET-RC)
- LAN Bridge 100 (DEBET-RD)
- LAN Bridge 100 (DEBET-RH)
- LAN Bridge 100 (DEBET-RJ)
- Remote Ethernet Repeater (DEREP-RA)
- Remote Ethernet Repeater (DEREP-RB)
- Remote Ethernet Repeater (DEREP-RC)
- Remote Ethernet Repeater (DEREP-RD)
- Remote Ethernet Repeater (DEREP-RH)
- Remote Ethernet Repeater (DEREP-RJ)

B.8.1.1 Sample Loss Calculation — Figure B-14 is a sample worksheet of a loss calculation for a 2 km link using 62.5/125 micron optical fiber cable and a DEBET-RH.

Using the allowable system losses listed in Table B-3, and the cable loss values listed in the manufacturer's specifications, Figure B-14 shows the allowable total loss.

PRODUCT NAME = DEBET - RH			
Fiber Type		= 62.5/125 micron	
Allowable System Loss		= 12 dB (see Tables B-2 or B-3 based on product variation)	
FACTOR	VALUE	QUANTITY	TOTAL LOSS
BARREL CONNECTIONS	1.5 dB/conn.	X 0 connectors	= 0.0 dB
SPLICES	0.2 dB/splice	X 4 splices	= 0.8 dB
CABLE ATTENUATION	3.0 dB/km	X 2.0 km of fiber	= 6.0 dB
BANDWIDTH DERATING	0.5 dB/km	X 2.0 km of fiber	= 1.0 dB
MARGIN	1.5 dB		= 1.5 dB
Sum of losses			= 9.3 dB

MKV88-1797

Figure B-14: Sample Worksheet for DEBET-RH Loss Calculation

The total system losses are less than the allowable system losses, therefore, the system design is valid.

Figure B-15 is a sample worksheet of a loss calculation for a 1 km link using 100/140 micron optical fiber cable and a DEBET-RC.

Using the allowable system losses listed in Table B-2, and the cable loss values listed in the manufacturer's specifications, Figure B-15 shows the allowable total loss.

PRODUCT NAME = DEBET - RC			
Fiber Type		= 100/140 micron	
Allowable System Loss		= 12.5 dB (see Tables B-2 or B-3 based on product variation)	
FACTOR	VALUE	QUANTITY	TOTAL LOSS
BARREL CONNECTIONS	0.8 dB/conn.	X 1 connectors	= 0.8 dB
SPLICES	0.2 dB/splice	X 2 splices	= 0.4 dB
CABLE ATTENUATION	4.5 dB/km	X 1.0 km of fiber	= 4.5 dB
BANDWIDTH DERATING	0.5 dB/km	X 1.0 km of fiber	= 0.5 dB
MARGIN	1.5 dB		= 1.5 dB
Sum of losses			= 7.7 dB

MKV88-1798

Figure B-15: Sample Worksheet for DEBET-RC Loss Calculation

The total system losses are less than the allowable system losses, therefore, the system design is valid.

It should be noted that a number of choices are available to link planners to minimize the link losses. Various types of splices produce different losses. Cables have a range of attenuation and lower values of attenuation may be used. Connectors may be eliminated in the link.

B.8.2 Worst-Case and Statistical Results

The data referenced in this section is based on a statistical analysis of fiber-optic systems. In a system with many variables, actual measured system loss will not be likely to match the worst-case calculations. In most cases, the calculated loss will exceed actual measured loss since it would be statistically improbable that every system component would have the worst-case characteristics.

However, if worst-case estimates are ignored so that the installed system depends entirely on a statistical spread, then the network installer or consultant is responsible for ensuring adequate performance of the system. The installer or consultant is responsible for diagnosing and modifying the system in the event of failure.

B.9 Attenuator Installation

The fiber-optic attenuator (P/N: 12-30068-01) is a device that induces a loss of 3 dB in a fiber-optic system. The attenuator (Figure B-16) is designed to be used with short cables and high power sources.




Figure B-16: Attenuator

B.9.1 When to Install the Attenuator

The attenuator can be used on the LAN Bridge 100 (DEBET-RH/RJ) and Remote Ethernet Repeater (DEREP-RH/RJ).

NOTE


The attenuator can be used only on LAN Bridge 100 or DEREPE devices that have the latest revision fiber-optic module (54-16053-01 Revision B1)

Use the attenuator when a:

- DEBET-RH/RJ or DEREPE-RH/RJ is connected to a DEBET-RH/RJ or DEREPE-RH/RJ (see Figure B-17).
 - Install the attenuator for 100/140 fiber-optic link that is 1000 m (3281 ft) or less.
 - Do NOT install the attenuator for 100/140 fiber-optic link beyond 1000 m (3281 ft).
 - Do NOT install the attenuator for 50/125, 62.5/125, and 85/125 optical fiber links.
- DEBET-RH/RJ is connected to a DEBET-RC/RD or DEREPE-RC/RD.
 - Install the attenuator on one side (DEBET-RH/RJ) by attaching the attenuator to the output (transmit) connector only (see Figure B-18).
 - If an older DEBET-RC/RD is being used with a newer DEBET-RH/RJ, the guidelines and budgets of the DEBET-RC/RD must be followed. Refer to *LAN Bridge 100 Installation/User's Guide*, Section 3.4 (Fiber-Optic Cables).


NOTE

If a special condition exists and more budget is required, the attenuator can be removed. This will increase the budget by 3 dB.



MKV88-1413

Figure B-17: Installing the Attenuator on Both Sides of the Link



MKV88-2046

Figure B-18: Installing the Attenuator on -RH/-RJ Side Only

B.9.2 Where to Install the Attenuator

The attenuator is:

- Installed ONLY on the SMA 906 connector that is on the fiber-optic cable end attached to the transmitter (output) of the bridge or repeater.
- Only one attenuator per fiber-optic cable.
- Do not install the attenuator onto the RX end of the fiber-optic cable.
- Do not install attenuators on both ends of the same optical fiber.
- Do not install the attenuator on a DEBET-RC/RD or DEREPC-RC/RD unit.
- If an attenuator is required for the second fiber, it will be installed on the SMA 906 connector attached to the transmitter of the bridge or repeater at the other end of the link.
- When the attenuator is installed, the cable should be marked with a label (36-18460-01) indicating “3 dB attenuator installed” and the loss of the cable (if known).

B.9.3 How to Install the Attenuator

1. Pull the protective caps from both the fiber-optic cable (SMA 906 connector) and the LAN Bridge 100 or DEREK fiber-optic connectors (see Figure B-19).




Figure B-19: Removing the Protective Caps from the LAN Bridge 100

2. Remove the attenuator (P/N: 12-30065-01) from the plastic bag.

CAUTION

The attenuator must be kept free of dust and dirt to ensure proper installation.

3. Install the attenuator between the source (TX) fiber-optic output connector and the fiber-optic cable connector (SMA 906). See Figure B-20.

NOTE

It is possible for the attenuator to fall out of the connector during installation. Before screwing the cable into the connector, ensure that the attenuator is still in place.

4. Connect the fiber-optic cable. Finger tighten the fiber-optic connectors.





Figure B-20: Connecting the Fiber-Optic Cables to the LAN Bridge 100

5. Place the attenuator label on the transmit cable of the fiber-optic cable (see Figure B-21).



MKV88-1417

Figure B-21: Fiber-Optic Cable Installed to the LAN Bridge 100

6. Refer to the *Fiber-Optic Attenuator Installation and Configuration Card* or the *DEBET/DEREP Fiber-Optic Upgrade* document for complete installation.

B.10 Reference Material

The following documents may be of general help in consulting on fiber-optic systems.

- *DEREP-RA Remote Ethernet Repeater Installation/Owner's Manual* (EK-DEREP-IN)
- *Ethernet Repeater Technical Manual* (EK-DEREP-TM)
- *Ethernet Installation Guide* (EK-ETHER-IN)
- *Digital Fiber-Optics Service Plan*, May 1984
- *Fiber-Optic Pocket Guide* (EK-FIBOP-IS)
- *FOTEC Practical Testing of Fiber-Optic Systems*
- *Siecor Fiber-Optic Handbook for Use With DEC Ethernet Repeaters*
- *Amphenol Fiber-Optic Designer's Handbook*
- *AT&T Lightguide Cable Specification*, Issue 5
- *ANSI Standard Z136.2: Fiber-Optic Safety*
- *DEBET/DEREP Fiber-Optic Upgrade* (EK-DEFOU-IN)
- *Fiber-Optic Attenuator Installation and Configuration Card* (EK-DEFOE-RC)

Index

A

Address

filtering, 3-7

Address Table

see also Circuit Descriptions

active station list, 3-21, 3-25

aged addresses, 3-25

entries, 3-21

inserting new addresses, 3-25

maintenance, 3-25

moved addresses, 3-24

new entries, 3-25

search time, 3-22

writing to, 3-24

B

Background Operations, 3-25, 3-38, 3-40

Binary Search, 3-21 to 3-23, 4-14, 4-31

Bridge access, 3-41

C

Circuit Descriptions

address comparator, 4-27 to 4-30

Ethernet address table, 4-30

fiber-optic module, 4-19

LANCE, 4-17

LANCE CSR registers, 4-17 to 4-19

LANCE register address port (RAP), 4-19

LANCE register data port (RDP), 4-19

NI ports, 4-16

NI subsystem, 4-15 to 4-23

packet memory, 4-20 to 4-23

packet memory address multiplexer, 4-20

packet memory refresh counter, 4-23

power supply, 4-32

processor subsystem, 4-5 to 4-14

search address multiplexer, 4-26

search control and status PALs, 4-31

Circuit Descriptions (Cont.)

search results register, 4-31

TLU, 4-26 to 4-31

Configuring Extended LANs, 1-16

Connectors, 1-9, 1-10, 1-14

Controls, 1-9, 1-10, 1-11

controlling bridge access, 3-41

Critical Path Operations, 3-39

D

Data Path, 4-3 to 4-4

Descriptor Rings, 3-5, 3-7

description, 3-9 to 3-17

in memory, 4-13

pointer to receive ring, 3-8

pointer to transmit ring, 3-9

receive descriptor ring entry, 3-11 to 3-13

transmit descriptor ring entry, 3-14 to 3-17

Designated Bridge, 3-28

examples, 3-29

Down-line loading, 3-6

E

Extended LANs, 1-1, 1-2 to 1-3

considerations, 1-16 to 1-31

F

Fiber-Optic Cable, 1-21 to 1-22

characteristics, 1-20

length, 1-20, 1-21

signal attenuation, 1-21

specifications, 1-31

with repeaters, 1-22

Fiber-Optic Link

see Fiber-Optic Cable

Foreground Operations, 3-38, 3-39

Forwarding, 3-23, 4-3

H

Hello Messages, 3-28, 3-30
description, 3-28

I

Initialization Block, 3-5, 3-7
field functions, 3-8
pointers, 3-5

Interrupts

hardware interrupts, 4-8
interrupt controller, 4-7

L

LAN Bridge 100

benefits to network, 1-1
connections allowed, 1-4
connectors, 1-9, 1-10
controls, 1-9, 1-10
description, 1-1, 1-4 to 1-14
limits on use of, 1-18
local, 1-18
packaging, 1-7
product designations, 1-7
remote, 1-20
status LEDs, 1-9, 1-10
transparent operation, 1-3
versions, 1-5
with repeaters
 see Repeaters
with routers
 see Routers

LAN Traffic Monitor

see LTM
remote access, 3-41

Local LAN Bridge 100

see LAN Bridge 100

Loop Detection, 1-18, 3-6, 3-27 to 3-38

and critical path operations, 3-39
benefits to network, 3-27
examples, 3-31 to 3-38

LTM, 1-3

description, 1-5, 1-14
software, 1-15
software installation, 2-4

M

Maintenance

corrective, 5-4 to 5-14
LAN Bridge 100 disassembly, 5-15 to 5-30
optional equipment, 5-2
philosophy, 5-1
preventive, 5-3

Maintenance (Cont.)

removal procedures

chassis cover removal, 5-18
fiber-optic interface, 5-21
logic module, 5-29
plastic enclosure removal, 5-16
power supply, 5-23

replacing the LAN Bridge 100, 5-14

required equipment, 5-1

Maintenance Operation Protocol (MOP), 3-7

description, 3-42

Memory

see also Circuit Descriptions

Ethernet address RAM test, 3-44

Ethernet address ROM checksum test, 3-44

Ethernet address table RAM, 4-30

map, 4-10 to 4-14

NVRAM, 4-9

NVRAM checksum test, 3-44

NVRAM write test, 3-44

packet memory, 4-13, 4-20 to 4-23

Packet memory refresh test, 3-44

processor memory, 4-9

RAM, 4-9

status, 4-9, 4-13

RAM hidden during power-up, 4-10

RAM test, 3-44

ROM, 4-9

ROM test, 3-44

status RAM test, 3-47

Microprocessor, 4-7

O

Operating Voltage Range, 4-32

see also Specifications

Operational States, 3-1 to 3-7

BACKUP, 3-1, 3-7

BROKEN, 3-1, 3-5

FORWARDING, 3-1, 3-7

INITIALIZATION, 3-1, 3-5

PREFORWARDING, 3-1, 3-6, 3-20

SELF-TEST, 3-1, 3-5

transition between, 3-1

OWN Bit, 3-12, 3-15, 4-20

OWN Interrupt, 4-20

see also Interrupts

P

Packet

stale, 3-24

Packet memory

see Circuit Descriptions

Path Cost, 3-28

Product Designations
see LAN Bridge 100

R

RBMS, 3-7, 3-21, 3-24
remote access, 3-41
setable parameter, 3-6

Regionalization, 3-24

Register Address Port (RAP)
see Circuit Descriptions

Register Data Port (RDP)
see Circuit Descriptions

Registers

see also Circuit Descriptions
compare-and-move, 4-13
LANCE CSR, 4-13
miscellaneous control, 4-10, 4-12

Remote LAN Bridge 100
see Lan Bridge 100

Removal Procedures
see Maintenance

Repair

see Maintenance

Repeaters, 1-16, 1-18, 1-22

Reset Circuitry, 4-14

Root Bridge, 3-28

Routers, 1-27

S

Search Results Register, 3-22

Self-Test

basic tests, 3-43
description, 3-42 to 3-47
external loop tests, 3-46
internal loop tests, 3-45
LANCE tests, 3-45
loop-on-self-test, 3-43
memory tests, 3-43
normal mode, 3-43
TLU tests, 3-47

Single-Entry Compare, 3-26

Spanning Tree Algorithm, 3-27, 3-28
see Loop Detection
examples, 3-31 to 3-38
parameters, 3-30

Specifications

electrical, 1-33
environmental, 1-32
physical, 1-31

States

see Operational States

Status LEDs, 1-9, 1-10, 1-13

Status Word, 3-21, 3-25

Subsystems

concurrent operation of, 4-2, 4-19, 4-23
hardware overview, 4-1 to 4-2
NI, 4-1, 4-15 to 4-23
power supply, 4-1, 4-32
processor, 4-1, 4-5 to 4-14
TLU, 4-1, 4-23 to 4-31

T

Topology

changes, 3-25

Troll feature, 3-18, 3-24

W

Watchdog Timer, 4-14