

SunFDDI/P™ 1.0 Adapter User's Guide



THE NETWORK IS THE COMPUTER®

Sun Microsystems Computer Company

A Sun Microsystems, Inc. Business
2550 Garcia Avenue
Mountain View, CA 94043 USA
415 960-1300 fax 415 969-9131

Part No.: 805-0809-10
Revision A, May 1997

Copyright 1997 Sun Microsystems, Inc. 2550 Garcia Avenue, Mountain View, California 94043-1100 U.S.A.

All rights reserved. This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Portions of this product may be derived from the UNIX[®] system and from the Berkeley 4.3 BSD system, licensed from the University of California. UNIX is a registered trademark in the United States and in other countries and is exclusively licensed by X/Open Company Ltd. Third-party software, including font technology in this product, is protected by copyright and licensed from Sun's suppliers.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 and FAR 52.227-19.

Sun, Sun Microsystems, the Sun logo, Solaris, SunOS, SunFDDI/P, and SunNet Manager are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and in other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the United States and in other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc. Netware[®] is a registered trademark of Novell, Inc. in the United States and other countries.

The OPEN LOOK[®] and Sun[™] Graphical User Interfaces were developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox Corporation in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a nonexclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

X Window System is a trademark of X Consortium, Inc.

THIS PUBLICATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

Copyright 1997 Sun Microsystems, Inc., 2550 Garcia Avenue, Mountain View, Californie 94043-1100 U.S.A.

Tous droits réservés. Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie et la décompilation. Aucune partie de ce produit ou de sa documentation associée ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a.

Des parties de ce produit pourront être dérivées du système UNIX[®] et du système Berkeley 4.3 BSD licencié par l'Université de Californie. UNIX est une marque enregistrée aux Etats-Unis et dans d'autres pays, et licenciée exclusivement par X/Open Company Ltd. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Sun, Sun Microsystems, le logo Sun, Solaris, SunOS, SunFDDI/P, et SunNet Manager sont des marques déposées ou enregistrées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC, utilisées sous licence, sont des marques déposées ou enregistrées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc. Netware[®] est une marque enregistrée de Novell, Inc. aux Etats-Unis et dans d'autres pays.

Les utilisateurs d'interfaces graphiques OPEN LOOK[®] et Sun[™] ont été développés de Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox Corporation pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique, cette licence couvrant aussi les licenciés de Sun qui mettent en place les utilisateurs d'interfaces graphiques OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

Le système X Window est un produit du X Consortium, Inc.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" SANS GARANTIE D'AUCUNE SORTE, NI EXPRESSE NI IMPLICITE, Y COMPRIS, ET SANS QUE CETTE LISTE NE SOIT LIMITATIVE, DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DES PRODUITS A REpondre A UNE UTILISATION PARTICULIERE OU LE FAIT QU'ILS NE SOIENT PAS CONTREFAISANTS DE PRODUITS DE TIERS.



Please
Recycle



Adobe PostScript

Regulatory Compliance Statements

Your Sun product is marked to indicate its compliance class:

- Federal Communications Commission (FCC) — USA
- Department of Communications (DOC) — Canada
- Voluntary Control Council for Interference (VCCI) — Japan

Please read the appropriate section that corresponds to the marking on your Sun product before attempting to install the product.

FCC Class A Notice

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions:

1. This device may not cause harmful interference.
2. This device must accept any interference received, including interference that may cause undesired operation.

Note: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

Shielded Cables: Connections between the workstation and peripherals must be made using shielded cables in order to maintain compliance with FCC radio frequency emission limits. Networking connections can be made using unshielded twisted-pair (UTP) cables.

Modifications: Any modifications made to this device that are not approved by Sun Microsystems, Inc. may void the authority granted to the user by the FCC to operate this equipment.

FCC Class B Notice

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions:

1. This device may not cause harmful interference.
2. This device must accept any interference received, including interference that may cause undesired operation.

Note: This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio/television technician for help.

Shielded Cables: Connections between the workstation and peripherals must be made using shielded cables in order to maintain compliance with FCC radio frequency emission limits. Networking connections can be made using unshielded twisted pair (UTP) cables.

Modifications: Any modifications made to this device that are not approved by Sun Microsystems, Inc. may void the authority granted to the user by the FCC to operate this equipment.

DOC Class A Notice - Avis DOC, Classe A

This Class A digital apparatus meets all of requirements the Canadian Interference-Causing Equipment Regulations.
Cet appareil numérique de la classe A respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

DOC Class B Notice - Avis DOC, Classe B

This Class B digital apparatus meets all of requirements the Canadian Interference-Causing Equipment Regulations.
Cet appareil numérique de la classe B respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

VCCI 基準について


第一種VCCI基準について

第一種VCCIの表示があるワークステーションおよびオプション製品は、第一種情報装置です。これらの製品には、下記の項目が該当します。

この装置は、第一種情報装置(商工業地域において使用されるべき情報装置)で商工業地域での電波障害防止を目的とした情報処理装置等電波障害自主規制協議会(VCCI)基準に適合しております。したがって、本製品を、住宅地域または住宅地域に隣接した地域でご使用になりますと、ラジオ、テレビジョン受信機等に受信障害を与えることがあります。

取り扱い説明書に従って正しくお取り扱いください。

第二種VCCI基準について

第二種VCCIの表示  があるワークステーションおよびオプション製品は、第二種情報装置です。これらの製品には、下記の項目が該当します。

この装置は、第二種情報装置(住宅地域または住宅地域に隣接した地域において使用されるべき情報装置)で住宅地域での電波障害防止を目的とした情報処理装置等電波障害自主規制協議会(VCCI)基準に適合しております。しかし、本製品を、ラジオ、テレビジョン受信機に近接してご使用になりますと、受信障害の原因となることがあります。

取り扱い説明書に従って正しくお取り扱いください。

Declaration of Conformity

Compliance ID: PCI-S10-100

Product Name: SunFDDI/P SAS Adapter

This equipment complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: 1) This equipment may not cause harmful interference and 2) This equipment must accept any interference that may cause undesired operation.

In addition this equipment complies with the following requirements of the EMC Directive 89/336/EEC and Low Voltage Directive 73/23/EEC;

EMC:	EN55022 / CISPR22 (1991)	Class B	
	EN50082-1	IEC801-2 (1991)	4 kV (Direct), 8 kV (Air)
		IEC801-3 (1984)	3 V/m
		IEC801-4 (1988)	1.0 kV Power Lines, 0.5 kV Signal Lines
	EN61000-3-2/IEC1000-3-2(1994)		Pass (Class D)

Supplementary Information:

This product was tested and complies with all the requirements for the CE Mark.

_____/ S /_____
Dennis P. Symanski DATE
Manager, Product Compliance

Sun Microsystems Computer Company
2550 Garcia Avenue, M/S UMPK15-102
Mt. View, CA 94043, USA
Tel: 415-786-3255
Fax: 415-786-3723

_____/ S /_____
Stephen McGoldrick DATE
Quality Assurance Manager

Sun Microsystems Limited
Springfield, Linlithgow
West Lothian, EH49 7LR
Scotland, United Kingdom
Tel: 0506 670000
Fax: 0506 760011

Declaration of Conformity

Compliance ID: PCI-D10-100

Product Name: SunFDDI/P DAS Adapter

This equipment complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: 1) This equipment may not cause harmful interference and 2) This equipment must accept any interference that may cause undesired operation.

In addition this equipment complies with the following requirements of the EMC Directive 89/336/EEC and Low Voltage Directive 73/23/EEC;

EMC:	EN55022 / CISPR22 (1991)	Class B
	EN50082-1 IEC801-2 (1991)	4 kV (Direct), 8 kV (Air)
	IEC801-3 (1984)	3 V/m
	IEC801-4 (1988)	1.0 kV Power Lines, 0.5 kV Signal Lines
	EN61000-3-2/IEC1000-3-2(1994)	Pass (Class D)

Supplementary Information:

This product was tested and complies with all the requirements for the CE Mark.

_____/ S /_____
Dennis P. Symanski DATE
Manager, Product Compliance

Sun Microsystems Computer Company
2550 Garcia Avenue, M/S UMPK15-102
Mt. View, CA 94043, USA
Tel: 415-786-3255
Fax: 415-786-3723

_____/ S /_____
Stephen McGoldrick DATE
Quality Assurance Manager

Sun Microsystems Limited
Springfield, Linlithgow
West Lothian, EH49 7LR
Scotland, United Kingdom
Tel: 0506 670000
Fax: 0506 760011

Contents

Preface.....	xix
<i>Part 1 —Installing and Configuring SunFDDI/P</i>	
1. SunFDDI/P Overview	1
Feature Summary.....	1
FDDI Conformance	2
2. Hardware Installation	3
Device Instances and Device Names.....	4
Installation Items	4
Installing SunFDDI/P Adapter Cards.....	5
Connecting Fiber Optic Cables.....	5
▼ To Connect the Fiber Optic Cable	6
MIC-to-MIC Coupler.....	7
Cleaning the Fiber Optic Cable	7
▼ To Clean the Fiber Optic Cable	8
Link Status Indicator (Diagnostic LED)	9

3. Installing SunFDDI/P Software	11
Installation Requirements	12
Declaring IP Addresses for SunFDDI/P	12
Installing and Configuring SunFDDI/P	13
▼ To Load and Mount the CD-ROM	13
▼ To Install SunFDDI/P Using pkgadd	13
▼ To Finish the Installation	17
▼ To Check the Installation	18
▼ Unloading Drivers	19
Installing SunNet Manager Agents for SunFDDI/P	20
▼ To Use pf_install_agents	20
Removing SunFDDI/P	22
▼ To Remove SunFDDI/P	22
4. Improving Network Performance	23
High Performance FDDI Networks	24
Ring Latency Improvement	24
Data Throughput Improvement	25
Send and Receive Buffers	25
Performance Across Bridges	25
Target Token Rotation Time (TTRT)	26
Improving Performance	27
Changing the High and Low Water Marks	27
▼ To Tune the High Water Mark	27
▼ To Tune the Low Water Mark	28

Tuning the High and Low Water Marks at Boot Time.	29
Tuning the Socket Options	29
MTU Path Discovery	30
▼ To Tune the Maximum Transfer Unit (MTU)	30
▼ To Disable MTU Path Discovery	31
▼ To Tune the Target Token Rotation Time (TTRT)	31
5. Troubleshooting and Diagnostics	33
Troubleshooting Checklist	34
▼ To Check the Connection to the Station	34
▼ To Check the Connections Between Stations	34
▼ To Check the Link Status Indicator	34
Checking the MAC Address.	35
▼ To Check the SunFDDI/P Driver.	36
▼ To Check the IP Addresses	36
▼ To Check the IP Routing	37
▼ To Check the Protocol Statistics	37
▼ To Check the SMT Traffic	38
Solving Common Problems	39
Running the Hardware Self-Test Program	44
▼ To Run the Hardware Self-Test	44
Loading the Device Driver Manually	45
▼ To Configure and Load the Device Driver	45
 <i>Part 2 —Planning and Implementing SunFDDI Networks</i>	
6. FDDI Network Architecture	49

FDDI Network Model	50
Physical Medium Dependent (PMD) Layer	51
Physical (PHY) Layer	51
Media Access Control (MAC) Layer	52
Station Management (SMT) Layer	52
Communication Between FDDI Layers	53
FDDI Network Architecture	55
FDDI Stations	56
FDDI Concentrators	58
FDDI Failure Recovery	60
Station Wrapping	60
Optical Bypass Switches	62
FDDI Ring Operation	65
Unrestricted Asynchronous Transmission	65
FDDI Performance	66
7. FDDI Network Topologies	69
Basic Dual-Ring Network	70
Advantages	70
Disadvantages	71
Standalone Concentrator	72
Advantages	72
Disadvantages	73
Concentrators with Dual-Homing	74
Advantages	75

Disadvantages	75
Tree of Concentrators.....	75
Advantages.....	76
Disadvantages	77
Ring of Trees.....	77
Advantages.....	78
Disadvantages	78
Mixed FDDI/Ethernet Networks.....	79
FDDI/Ethernet IP Router	79
A Tree of Concentrators in an Ethernet Network	80
A Ring of Trees in an Ethernet Network.....	81
Routes in Mixed FDDI/Ethernet Networks.....	81
8. Using the SunFDDI	
Network Utilities	85
Changing the Default MAC Address (pf_macid)	86
▼ To Change the Default MAC Address with pf_macid	86
Displaying SunFDDI/P Statistics (pf_stat).....	87
Displaying Local Interface Statistics	88
Displaying Statistics from Neighboring Stations.....	94
Monitoring SMT Frames (pf_smtmon).....	98
SMT Frame Classes and Types.....	100
Filtering Output from pf_smtmon	101
9. Managing FDDI Stations	
Using SunNet Manager	103
Installing the SunNet Manager Agents.....	104

▼ To Install the Standard SNM Agents on a Client	104
▼ To Install the FDDI Schemas on a Client	105
Installing the FDDI Schemas on the Console	106
▼ To Share <code>pf_install_agents</code> From a Client	107
Managing FDDI Networks with SunNet Manager	109
▼ To Set Up the SunNet Manager Console.	109
Using the SunFDDI/P Local Agent (<code>pf_fddi</code>)	110
▼ To Use the SunFDDI/P Proxy Agent (<code>pf_fddismt7_2</code>)	112
Interpreting Traps and Errors.	114
MAC Address Formats	117
10. Developing Applications that Run over SunFDDI/P	119
Developing Applications with DLPI	120
DLPI Connectionless Service Primitives.	121
DLPI Addressing	122
11. Setting Up Servers and Diskless Clients.	125
Setting Up a Server and Diskless Clients	126
▼ To Set Up a Server	127
▼ To Introduce Diskless Clients to a Server	127
▼ To Update the NIS Server	132
Installing the SunFDDI/P Driver for a Diskless Client	133
Booting a Diskless Client Over SunFDDI/P.	133
▼ To Boot a Diskless Client	133
A. Frequently Asked Questions About SunFDDI	135

Glossary	141
Index.....	147

Figures

Figure 2-1	Aligning SC-Type Fiber Optic Cable Connectors	6
Figure 2-2	Cleaning the End of a Fiber Optic Cable	8
Figure 6-1	FDDI Architectural Model	50
Figure 6-2	Communication Between FDDI Layers	54
Figure 6-3	Basic FDDI Network Architecture	55
Figure 6-4	Single-Attached Station (SAS)	56
Figure 6-5	Dual-Attached Station (DAS)	57
Figure 6-6	Single-Attached Concentrator (SAC)	58
Figure 6-7	Dual-Attached Concentrator (DAC)	59
Figure 6-8	Isolating a Single Station Failure	60
Figure 6-9	Isolating a Single Cable Failure	61
Figure 6-10	Dual-Attached Station in Wrap Mode	62
Figure 6-11	Dual-Ring Network Divided by Two Faults	63
Figure 6-12	Optical Bypass Switch	63
Figure 6-13	Optical Bypass Switches used in a Network	64
Figure 7-1	Basic Dual-Ring Network	70

Figure 7-2	Basic Dual-Ring Network with Two Faults.	71
Figure 7-3	Standalone Concentrator	72
Figure 7-4	Standalone Concentrator With Dual-Homing	74
Figure 7-5	Hierarchical Tree of Concentrators.	76
Figure 7-6	Ring of Trees.	77
Figure 7-7	Basic Mixed FDDI/Ethernet Topology	80
Figure 7-8	Tree of Concentrators Attached to an Ethernet Network	80
Figure 7-9	Ring of Trees Attached to an Ethernet Network.	81
Figure 7-10	Routing in an FDDI/Ethernet Internetwork	82
Figure 8-1	Normal Sequence of PCM States.	92
Figure 9-1	Quick Dump (Local Statistics)	110
Figure 9-2	Quick Dump (Neighbor Statistics)	111
Figure 9-3	SunFDDI/P Proxy Agent	112
Figure 9-4	Specifying the Target Station.	112
Figure 9-5	Quick Dump (SMT Frames).	113
Figure 9-6	Single-Attached Station Disconnected	114
Figure 9-7	Single-Attached Station Reconnected.	115
Figure 9-8	Dual-Attached Station Halted	116
Figure 9-9	Dual-Attached Station Restarted.	116
Figure 10-1	DLPI Access to FDDI.	120
Figure 10-2	DLPI Request/Response Negotiation	121
Figure 10-3	SunFDDI/P DLSAP Format.	122

Tables

Table 5-1	Problems Installing SunFDDI/P	39
Table 5-2	Problems Running SunFDDI/P	40
Table 5-3	Problems Using the SNM Agents for SunFDDI/P	43
Table 8-1	pf_stat States Under the ECM Heading	90
Table 8-2	pf_stat States Under the RMT Heading	90
Table 8-3	pf_stat States Under the PCM Heading	91
Table 10-1	DL_INFO_REQ Frame	122

Preface

The *SunFDDI/P™ 1.0 Adapter User's Guide* describes how to turn your system into an FDDI station attached to an FDDI network. It is intended for experienced network administrators who are familiar with the Solaris 2.5.1 Hardware: 4/97 operating environment. The SunFDDI/P 1.0 Adapter is supported on Sun-4u PCI-based system architectures.

How this Book is Organized

The *SunFDDI/P™ 1.0 Adapter User's Guide* is divided into two parts:

If you are unfamiliar with FDDI network terminology, or if you are implementing an FDDI network for the first time, read Part 2 first.

- *Part 1 — “Installing and Configuring SunFDDI/P,”* describes how to turn your machine into a SunFDDI/P station. It tells you how to install and configure the SunFDDI/P device driver, and how to tune your FDDI station to get the best performance. It also includes troubleshooting advice to help you detect and resolve problems with your FDDI network.
- *Part 2 — “Planning and Implementing SunFDDI Networks,”* includes a brief overview of the FDDI protocols, and describes common FDDI network topologies. It tells you how to use the SunFDDI/P network utilities, how to use a SunNet Manager™ console to manage your SunFDDI/P stations, how to develop applications over SunFDDI/P, and how to boot a diskless station across an FDDI connection.

The connector type has been changed for SunFDDI/P to an SC-type connector. If you are connecting the SunFDDI/P card to a network that has a MIC connector, an SC-MIC converter cable is required.

Part 1 — “Installing and Configuring SunFDDI/P”

Chapter 1, “SunFDDI/P Overview,” describes the SunFDDI/P 1.0 implementation of the FDDI protocols and includes a list of the specifications to which it conforms.

Chapter 2, “Hardware Installation,” tells you where to find detailed instructions on how to install your SunFDDI/P PCI card in your machine and describes how to connect the FDDI cable.

Chapter 3, “Installing SunFDDI/P Software,” describes how to install the SunFDDI/P device driver and utilities on machines running a Solaris™ 2.5.1 Hardware: 97 environment. It includes a description of the post-installation scripts that configure your machine as an FDDI station.

Chapter 4, “Improving Network Performance,” describes the configurable network parameters and provides general advice on how to obtain the best performance from your network.

Chapter 5, “Troubleshooting and Diagnostics,” tells you how to detect and resolve problems with your FDDI network.

Part 2 — “Planning and Implementing SunFDDI Networks”

Chapter 6, “FDDI Network Architecture,” provides an overview of the FDDI network model and describes how FDDI networks operate.

Chapter 7, “FDDI Network Topologies,” describes various pure FDDI network topologies and discusses their relative strengths and weaknesses. It also describes how to implement routing between mixed FDDI and Ethernet networks.

Chapter 8, “Using the SunFDDI Network Utilities,” describes the network utilities delivered with SunFDDI/P and tells you how to use them to recover network statistics.

Chapter 9, “Managing FDDI Stations Using SunNet Manager,” describes how to install the SunNet Manager agents for SunFDDI/P and how to set up the SunNet Manager console to manage them.

Chapter 10, “Developing Applications that Run over SunFDDI/P,” describes how to create applications that run over SunFDDI/P, using the DLPI interface for a Solaris 2.x environment.

Chapter 11, “Setting Up Servers and Diskless Clients,” describes how to set up a server so that you can boot and run diskless clients over an FDDI network.

Appendix A, “Frequently Asked Questions About SunFDDI,” contains a list of common questions and their responses.

UNIX Commands

This document may not include specific software commands or procedures. Instead, it may name software tasks and refer you to operating system documentation or the handbook that was shipped with your new hardware.

The type of information that you might need to use references for includes:

- Shutting down the system
- Booting the system
- Configuring devices
- Other basic software procedures

See one or more of the following:

- *Solaris 2.x Handbook for SMCC Peripherals* contains Solaris™ 2.x software commands.
- Online AnswerBook™ for the complete set of documentation supporting the Solaris 2.x software environment.
- Other software documentation that you received with your system.

Typographic Conventions

The following table describes the typographic changes used in this book.

Typeface or Symbol	Meaning	Example
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. machine_name% You have mail.
AaBbCc123	What you type, contrasted with on-screen computer output	machine_name% su Password:
<i>AaBbCc123</i>	Command-line placeholder: replace with a real name or value	To delete a file, type <code>rm filename</code> .
<i>AaBbCc123</i>	Book titles, new words or terms, or words to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be root to do this.

Shell Prompts

The following table shows the default system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

Shell	Prompt
C shell	machine_name%
C shell superuser	machine_name#
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Related Documents

For a more detailed description of FDDI technology and the relevant FDDI protocols, see these documents:

Title	Author/Publisher	Part Number
<i>Handbook of Computer-Communications Standards, Volume 2,</i> <i>FDDI Technology and Applications</i>	William Stallings, Macmillan Publishing Company: 1987	NA
<i>FDDI A High Speed Network</i>	Edited by Sonu Mirchandani and Raman Khanna, John Wiley & Sons: 1993	NA
<i>STREAMS DLPI Specification</i>	Amit Shah and G. Ramakrisnan, Prentice Hall, Inc.: 1994	NA
	Sun Microsystems	800-6915

Ordering Sun Documents

SunDocsSM is a distribution program for Sun Microsystems technical documentation. Easy, convenient ordering and quick delivery is available from SunExpressTM. You can find a full listing of available documentation on the World Wide Web: <http://www.sun.com/sunexpress/>

Country	Telephone	Fax
United States	1-800-873-7869	1-800-944-0661
United Kingdom	0-800-89-88-88	0-800-89-88-87
France	0800-90-61-57	0800-90-61-58
Belgium	02-720-09-09	02-725-88-50
Luxembourg	32-2-720-09-09	32-2-725-88-50
Germany	01-30-81-61-91	01-30-81-61-92
The Netherlands	06-022-34-45	06-022-34-46
Sweden	020-79-57-26	020-79-57-27
Switzerland	0800-55-19-26	0800-55-19-27
Japan	0120-33-9096	0120-33-9097

Getting Help

For technical assistance in the United States, call 1-800-872-4786.

To get the latest patches and patch revisions, contact your local Sun Service provider. For additional information, access Sun on the World Wide Web: <http://www.sun.com> and select Sales & Service ► [On-line support](#) ► [SunSolve Online™](#) ► Patches.

Sun Welcomes Your Comments

Please use the *Reader Comment Card* that accompanies this document. We are interested in improving our documentation and welcome your comments and suggestions.

If a card is not available, you can email or fax your comments to us. Please include the part number of your document in the subject line of your email or fax message.

- Email: `smcc-docs@sun.com`
- Fax: SMCC Document Feedback
1-415-786-6443

*Part 1 — Installing and Configuring
SunFDDI/P*

<i>Feature Summary</i>	<i>page 1</i>
<i>FDDI Conformance</i>	<i>page 2</i>

This chapter describes the Sun FDDI (Fiber Distributed Data Interface) implementation of the FDDI protocols, including a list of the ANSI/FDDI standards to which it conforms. See Chapter 6, “FDDI Network Architecture” for more information on FDDI architecture.

Feature Summary

The SunFDDI/P 1.0 product is a combination of hardware and software that turns your system into an FDDI station. There are two hardware options:

- **SunFDDI/P SAS** — single attached station on fiber
- **SunFDDI/P DAS** — dual attached station on fiber

SunFDDI/P is a high-speed networking product that provides significantly greater bandwidth (up to 10 times) compared to 10-megabit Ethernet networks. It is designed for Sun-4u PCI-based system architectures running Solaris 2.5.1 Hardware: 97.

SunFDDI/P provides the following capabilities:

- Implements FDDI single-attached station (SunFDDI/P) and FDDI dual-attached station (SunFDDI/P Dual)

- Provides connection to multimode fiber (SunFDDI/P and SunFDDI/P Dual) networks
- Supports data transfer rates of up to 100 Mbps
- Provides power-up self test and system diagnostic tests
- Provides high reliability and availability through distributed, automatic recovery (DAS adapter only)
- Supports up to four SunFDDI/P PCI cards installed in one PCI bus and up to eight SunFDDI/P PCI cards installed in one machine
- Conforms to the 32-bit Peripheral Component Interconnect (PCI) specification for short form adapter cards
- Supports 33 MHz operating frequency and 5.0 volt I/O signaling
- Supports 62.5/125 micrometer, multimode fiber (MMF)
- Provides SunNet Manager™ agents for the effective management of FDDI stations and networks from a centralized SunNet Manager console; (requires SunNet Manager 2.2 or later in addition to SunFDDI/P)
- Complies with the ANSI X3T9.5 standard for single-attached LAN fiber rings and with revision 7.3 of the ANSI/FDDI SMT specification

FDDI Conformance

SunFDDI/P conforms to the following standards and specifications:

- ANSI/FDDI Media Access Control (MAC) X3.139-1987
- ANSI/FDDI physical sublayer (PHY) X3.148-1988
- ANSI/FDDI physical medium dependent (PMD) X3.166-1990
- ANSI/FDDI Station Management (SMT) X3.299 R7.3 (formally R7.2.99)
- PCI Spec. Rev 2.1 (Part Number 802-2387-02)

<i>Device Instances and Device Names</i>	<i>page 4</i>
<i>Installing SunFDDI/P Adapter Cards</i>	<i>page 5</i>
<i>Connecting Fiber Optic Cables</i>	<i>page 5</i>
<i>Link Status Indicator (Diagnostic LED)</i>	<i>page 9</i>

This chapter tells you where to find step-by-step procedures for installing PCI cards. It describes the convention used to assign device names to SunFDDI/P adapter cards, and tells you how to connect fiber optic cables.



Caution – For your personal safety and to protect your machine, please observe the following precautions when you install an PCI card in a SPARC™ workstation or server:

Before opening your machine, always ensure that the power switch is off (O position). When the power is off, the green light on the front of the machine is unlit and the fan in the power supply is not running. Always take care to protect your equipment from the effects of static electricity.

A lithium battery is molded into the NVRAM–TOD chip on the main logic board. This battery is not a customer-replaceable part. Do not attempt to disassemble it or recharge it.

Device Instances and Device Names

You can install a SunFDDI/P PCI card in any available master PCI slot. SunFDDI/P supports up to four SunFDDI/P PCI cards installed in one PCI, and up to eight SunFDDI/P PCI cards installed in one machine.

SunFDDI/P 1.0 PCI cards are assigned device names of the form `pf<inst>`, where the instance number `<inst>` is determined by the number and relative positions of the cards installed. When you install a SunFDDI/P card in your machine for the first time, it has instance number 0 and a device name `pf0`. If multiple cards are installed, the device name `pf0` is associated with the card in the lowest numbered slot, the device name `pf1` is associated with the card in the next highest numbered slot, and so on.

Refer to the hardware installation manual that came with your system for instructions on how to identify the PCI slots in your machine.

Note – For machines running Solaris 2.5.1: If a SunFDDI/P PCI card is moved from its initial slot to a different one, the operating system will find the device in the new slot and assume that it is a new instance of the device. As a result, the operating system will assign the next instance number to it and two devices will appear to be installed in the machine.

Installation Items

Before proceeding with the installation, make sure you have the following items. Contact your sales representative if any of these items are missing.

- SAS or DAS SunFDDI PCI adapter
- FDDI converter cable
- MIC (Media Interface Connector) coupler
- Disposable grounding wrist strap
- CD-ROM containing SunFDDI/P 1.0 Adapter software
- *SunFDDI/P 1.0 Adapter User's Guide*

Installing SunFDDI/P Adapter Cards



Caution – The SunFDDI/P PCI card is sensitive to static electricity. Always use the wrist strap supplied with SunFDDI/P when handling the card and ensure that the wrist strap is properly grounded. Handle the SunFDDI/P PCI card by the edges, and avoid touching any of the components.

Refer to the hardware installation manual that came with your system for detailed instructions on how to install an adapter card in your machine.

There are no user-configurable jumpers or switches on the SunFDDI/P PCI card. Altering any of the components on the card may affect its operation adversely and will render any product warranty void.

Connecting Fiber Optic Cables

Note – If you are connecting the SunFDDI/P card to a network that has a MIC connector, a converter cable is required. Earlier versions of SunFDDI used MIC connectors.

Sun provides converter cables that enable existing MIC-based networks to connect to the SunFDDI/P adapter without making cabling changes. Networks that have already been designed for SC connections do not need the converter cable. Less than one decibel of signal integrity is lost when using the adapter. Borderline applications may need to reconfigure the network directly for SC connection if the signal loss materially affects the error rate.

Fiber optic cable is fragile and should not be subjected to mechanical stress or impact. Always observe the following precautions when handling fiber optic cable and connectors:

- Fiber optic cable connectors consist of a ceramic or steel ferrule through which the optical fiber passes. Do not touch the tip of the ferrule, as marring or depositing oil or dirt on its delicate surface will destroy its ability to transmit light.
- Use the protective caps provided on the fiber optic cable connectors and optic data links whenever the cable is not connected.

- Fiber optic cable has a limited turn radius. Do not bend it in a way that could damage the fiber or cause signal loss.

▼ To Connect the Fiber Optic Cable

1. **Unpack the cable, remove the plastic protective caps or plugs from each end, and remove the plug from the transceiver unit on the card.**
2. **Orient the fiber optic cable to the connector on the SunFDDI/P card.**
Most fiber optic cables have raised keys on one side of the connectors (see Figure 2-2). There are also keyway notches on the inside of the connector on the card; orient the fiber optic cable so that the keys on the cable connectors line up with the keyway notches on the card. Also make sure to connect the cable to the proper port, either Port A or Port B (see Figure 2.2).

Warning – Not all cables have raised keys to help you orient the connection.




Figure 2-1 Aligning SC-Type Fiber Optic Cable Connectors

3. **Plug the cable into the card connector.**
A “click,” tells you that the cable is properly mated with the card connector.

MIC-to-MIC Coupler

A separate MIC (Media Interface Connector) coupler is shipped with the SunFDDI/P 1.0 Adapter product to support MIC-to-MIC network connections.

Two different versions of the coupler are available:

- One with raised keys to help you properly orient the connection
- One without raised keys, which has specific connecting instructions

Note – Make sure you properly connect the coupler. An improper connection will “twist” or “cross” the ring.

Cleaning the Fiber Optic Cable

Fiber optic cable connectors can be cleaned if they become contaminated with dirt, dust or oil, which can cause signal noise on the network.

▼ **To Clean the Fiber Optic Cable**

- 1. Gently wipe the end of the ferrule with an alcohol-moistened cotton swab.**




Figure 2-2 Cleaning the End of a Fiber Optic Cable

- 2. Blow clean, dry compressed air around the end of the connector and around the ferrule to dislodge loose dust and grit.**
- 3. Clean the adapter by running a pipe cleaner moistened in alcohol running through it.**
Repeat with a dry pipe cleaner.
- 4. Dry the adapter with clean, dry compressed air.**


Link Status Indicator (Diagnostic LED)

A link status indicator is mounted on the SunFDDI/P PCI card. The color of the LED indicates the current status of the connection to the FDDI network as follows:

- **No LED lit:** Either the driver is loaded and the interface is *not* configured or the driver has not been loaded.
- **Amber:** Interface configured, no valid line state to an active FDDI network exists.
- **Green:** Connection established to an active FDDI network.
- **Blinking green:** In a dual-ring environment, this is an invalid connection. In a dual-homing environment, this is normal. The connection is in Standby mode.
- **Red:** Adapter fails FDDI path test.

Refer to Chapter 5, “Troubleshooting and Diagnostics,” for more information on diagnosing problems with your FDDI network.

Installing SunFDDI/P Software

3 

<i>Installation Requirements</i>	<i>page 12</i>
<i>Declaring IP Addresses for SunFDDI/P</i>	<i>page 12</i>
<i>Installing and Configuring SunFDDI/P</i>	<i>page 13</i>
<i>Installing SunNet Manager Agents for SunFDDI/P</i>	<i>page 20</i>
<i>Removing SunFDDI/P</i>	<i>page 22</i>

This chapter assumes that you have already installed the SunFDDI/P adapter card in your machine and that you are loading the software from a local CD-ROM drive.

This chapter tells you how to install the SunFDDI/P device drivers, utilities, and SunNet Manager agents in a Solaris 2.5.1 Hardware: 4/97 environment. It includes instructions on how to check your installation and how to remove the SunFDDI/P software, if necessary.

Installation Requirements

- **Hardware and Software Platforms**

SunFDDI/P is designed for Sun-4u PCI-based system architectures (32 Mbytes RAM minimum) running a Solaris 2.5.1 Hardware: 4/97 environment.

- **Disk Space**

A machine with 1600 KBytes of disk space is required to install SunFDDI/P.

- **Host names and IP Addresses**

A unique host name and IP address must be assigned to each SunFDDI/P device installed in your machine. If you are installing more than one interface in a single machine, each interface must be connected to a different subnet of a network.

Declaring IP Addresses for SunFDDI/P

You must assign a unique IP address to each SunFDDI/P interface. If you have more than one networking interface installed on your machine—either multiple SunFDDI/P interfaces or some other LAN interface (for example, an Ethernet interface)—each interface must be connected to a different subnet.

Before installing the SunFDDI/P software, update the files on your NIS/NIS+ server to assign IP addresses and host names for the SunFDDI/P interfaces that you are going to configure. If you are not using a name service, update the `/etc/hosts` file on each machine in the network to add the IP addresses and host names of the stations attached to the FDDI ring.

The installation script updates the `/etc/hosts` file on the local machine and enters the IP addresses and host names assigned to the SunFDDI/P interfaces in which it is installed.

Installing and Configuring SunFDDI/P

Use `pkgadd(1M)` to install unbundled software. See the *Solaris 2.5 Software and AnswerBook Packages Administration Guide* for detailed information on installing software products using `pkgadd`.

For detailed instructions on how to load software from a CD-ROM drive mounted on a remote directory, see the *Solaris 2.5 Software and AnswerBook Packages Administration Guide*.

▼ To Load and Mount the CD-ROM

1. **Log in as root or become superuser.**
2. **Place the CD-ROM in its caddy and load the caddy into the CD-ROM drive.**
3. **Mount the CD-ROM on a local directory.**
 - If the Volume Manager (`vold`) is running on your machine, then the CD-ROM is mounted automatically under `/cdrom/sunfddip_1_0`.
 - If the Volume Manager (`vold`) is *not* running on your machine, create a directory called `/cdrom/sunfddip_1_0` and mount the CD-ROM manually:

```
# mkdir /cdrom/sunfddip_1_0
# /usr/sbin/mount -o ro -F hsfs /dev/dsk/c0t6d0s0 /cdrom/sunfddip_1_0
```

▼ To Install SunFDDI/P Using `pkgadd`

1. **Log in as root or become superuser.**
2. **Start `pkgadd` by typing:**

```
# /usr/sbin/pkgadd -d /cdrom/sunfddip_1_0/Product
```

3. Enter the number that corresponds to the package you want to install and press Return. Respond to any prompts that are displayed.

Two packages are associated with SunFDDI/P. They contain the device drivers and utilities used to manage your SunFDDI/P station.

```
The following packages are available:
```

```
1 SUNWpfr SunFDDI/P (Driver)
   (sparc) 1.0
2 SUNWpft SunFDDI/P (Man Pages/Utilities)
   (sparc) 1.0
```

```
Select package(s) you wish to process (or "all" to process all
packages). (default: all) [?,??,q]:
```

▼ **To Install the Device Driver (SUNWpfr)**

The SUNWpfr package contains the SunFDDI/P device driver and the SunNet Manager agent daemon and start-up scripts. By default, the base directory *<basedir>* for this package is *.* You can change the default base directory for this package when you start `pkgadd(1M)` with the `-R` option. This package must be installed on every machine that runs SunFDDI/P.

An installation script is executed automatically when you install the SunFDDI/P device driver (SUNWpfr). This script checks that the hardware is installed and displays the number of SunFDDI/P interfaces that it finds. If there is no hardware installed, the script halts the installation process without installing the software.

1. Indicate whether you are installing the driver for a diskless client.

- Type `y` if you want to install the driver in the `/export/root/<client>` file system for a diskless client
- Type `n` if you are installing the driver for the machine on which you are running the script. You must use `pkgadd` with the `-R` option to change the base directory when you install the driver for a diskless client.


```
Is this install for a diskless client [n] [y,n,?]
```

(See Chapter 11, “Setting Up Servers and Diskless Clients” for more information on booting diskless clients across an FDDI network.)

The number of SunFDDI/P interfaces found is displayed and you are asked how many interfaces you want to configure.

2. Enter a number from 1 to 16 or press Return for the default.

```
How many FDDI (pf) interfaces do you want to configure?
```

3. Type the host name for which you want the interface associated.

```
What host name do you want to use for interface 1: hysop-pf
```

4. Type the IP address that corresponds to the host name or press Return for the default IP address.

If an IP address is already associated with the host name (either in the NIS/NIS+ database or in the `/etc/hosts` file) the script offers it as a default IP address. You must assign a unique IP address to each SunFDDI/P interface. If you have multiple IP interfaces installed on one machine, each interface must be attached to a different FDDI subnetwork.

```
What ip address do you wish to use for hysop-pf [<ip_address>]:
```

5. Repeat Steps 3 and 4 until you have assigned a host name and IP address for each SunFDDI/P PCI card installed in your machine.

6. Indicate whether you want the SunNet Manager agents for SunFDDI/P to start when you reboot the machine.

- Type `y` to start the SunNet Manager agents for SunFDDI/P.

- Type `n`, or press Return, if you are not using SunNet Manager to manage your network; if you are running SunNet Manager agents on another machine; or if you want to start the SunNet Manager agents manually.

```
Do you want to start the SunNet Manager daemons for SunFDDI
at boot time? [n] [y,n,?,q] y
```

7. Type `y` to use the installation scripts with superuser permission.

```
This package contains scripts which will be executed with superuser
permission during the process of installing the package.
```

```
Do you want to continue with the installation [y,n,?] y
```

The script displays a list of the files that it installs, loads the driver, and configures the IP interfaces automatically.

If the installation is successful, your SunFDDI/P interfaces are now active. The link status indicator (diagnostic LED) will be:

- **Amber**—if the interface is configured and no valid line state to an active FDDI network exists
- **Green**—if the interface is configured and connected to an active network

▼ **To Install the Utilities, man Pages, and SNM Agents** (`SUNWpft`)

The `SUNWpft` package contains the SunFDDI/P utilities and man pages. By default, the base directory `<basedir>` for this package is `/opt`. You can change the default base directory when you start `pkgadd(1M)` with the `-R` option. This package can be installed on a server and shared between multiple machines, if required.

An installation script is executed automatically when you install the SunFDDI/P utilities and man pages (`SUNWpft`):

1. **Type `y` to accept the installation of files with `setuid/setgid` permission.**
If you are installing the `SUNWpft` package for the first time, this step will not occur.

```
Do you want to install these setuid/setgid files [y,n,?,q] y
```

2. **Type `y` to accept the installation scripts with superuser permission.**

```
This package contains scripts which will be executed with superuser  
permission during the process of installing the package.
```

```
Do you want to continue with the installation [y,n,?] y
```

▼ To Finish the Installation

1. **Type `q` to stop `pkgadd` when the list of available packages is displayed again.**

```
The following packages are available:
```

```
 1 SUNWpfr  SunFDDI/P (Driver)  
            (sparc) 1.0  
 2 SUNWpft  SunFDDI/P (Man Pages/Utilities)  
            (sparc) 1.0
```

```
Select package(s) you wish to process (or "all" to process all  
packages). (default: all) [?,??,q]:
```

2. **Eject the CD-ROM from the CD-ROM drive.**

- If Volume Manager (`vold`) is running on your machine, the CD-ROM is unmounted automatically when it is ejected:

```
# eject cdrom
```

- If Volume Manager (`vold`) is not running on your machine, unmount the CD-ROM before you eject and remove the directory that you created:

```
# /usr/sbin/umount /cdrom/sunfddip_1_0
# eject cdrom
# rmdir /cdrom/sunfddip_1_0
```

3. Reboot your machine.

▼ To Check the Installation

1. Use `netstat(1M)` to check for SunFDDI/P interfaces.

SunFDDI Interface
pf<inst> ---->

```
% /usr/bin/netstat -i
Name Mtu  Net/Dest Address  Ipkts Ierrs Opkts Oerrs Collis
Queue
lo0  8232  loopback localhost  21    0    21    0    0    0
le0  1500  our-lan  hysop    2146  0    950   1    13   0
pf0  4352  our-fddi hysop-pf 1086  0    907   0    0    0
```

2. Use `ifconfig(1M)` to check the current state of the SunFDDI/P interfaces.

```
% /usr/sbin/ifconfig pf0
pf0: flags=863<UP,BROADCAST,NOTRAILERS,RUNNING,MULTICAST> mtu 4352
    inet <host_address> netmask <netmask> broadcast <ip_address>
    ether <mac_address>
```

3. Use `ping(1M)` to verify that you can send and receive data across the FDDI connection.

If your SunFDDI/P station is attached to an active FDDI network, try to access a remote station.

```
% /usr/sbin/ping hemlock-pf
hemlock-pf is alive
```

Your SunFDDI/P station should now be active, and you should be able to send and receive data across an FDDI connection.

For optimum performance from your FDDI network, you may need to tune the configurable network parameters. See Chapter 4, “Improving Network Performance,” for detailed instructions.

If you have problems configuring or using your SunFDDI/P station, see Chapter 5, “Troubleshooting and Diagnostics,” for help.

▼ Unloading Drivers

The driver for the SunFDDI/P 1.0 adapter can be unloaded while the system is running by using the `modunload(1M)` command. If the SunNet Manager daemons are running on your system, you must kill that process *before* you unload the SunFDDI/P driver.

◆ If SunNet Manager is configured and running on your system:

```
% ps -ef | grep pf_snmd
% kill -9 <process IDs from above command>
% modinfo | grep smt
% modunload -i <module ID for smt driver>
```

◆ To unload the SunFDDI/P driver:

```
% modinfo | grep pf
% modunload -i <module ID for pf driver>
```

Installing SunNet Manager Agents for SunFDDI/P

Use `pf_install_agents` to install the SunNet Manager agents for SunFDDI/P. This script copies the FDDI schema files to the directory in which the standard agents are installed and updates the configuration files for SunNet Manager. Run `pf_install_agents` on each SunNet Manager Console, and on each SunFDDI/P station you want to manage using SunNet Manager.

▼ To Use `pf_install_agents`

1. **Log in as root or become superuser.**
2. **Start `pf_install_agents` by typing:**

```
# /opt/SUNWconn/bin/pf_install_agents
```

3. **Type the host names of the console machine(s) that you want to receive SMT Request Frame (SRF) trap reports from the local station.**
You can send SRF trap reports to more than one console. Each console must have the schemas for SunFDDI/P installed, for example:

Host name of Console -----> Host name of Console -----> Press Return to terminate --->	<pre>Enter hostnames (if any) to send SRF Trap reports to: (enter blank to terminate) enter hostname:opus enter hostname:rigel enter hostname:</pre>
--	--

4. Specify the destination directory for the SunFDDI/P schemas.

SunNet Manager looks for the schemas in the directory `<basedir>/SUNWconn/snm/agents`. The variable `<basedir>` is set to `/opt` by default. If you installed the SunNet Manager agents and libraries on the FDDI station, or if you used `getagents` to recover the agents from the console, then this directory exists.

Enter the base directory under which you installed the SunNet Manager agents as the destination directory for the SunFDDI/P schemas.

```
What is the name of the root directory under which the SunNet
Manager is installed? [/opt]: [?]
```

5. Specify the current location of the SunFDDI/P schemas.

This is the base directory under which you installed the SunFDDI/P software package (SUNWpf) using `pkgadd(1M)`.

```
What is the name of the root directory under which the pf FDDI
software is installed? [/opt]: [?]
```

The script copies the SunFDDI/P schemas to the destination directory and modifies the local configuration files for SunNet Manager.

6. Start the SunNet Manager agents for SunFDDI/P by typing:

```
# /usr/sbin/pf_snm
```

Note – If you typed `n` in response to the question “Do you want to start the SunNet Manager daemons for SunFDDI at boot time?” when you installed the SunFDDI/P software, you must start the SunNet manager agents for SunFDDI/P *manually* each time you reboot your machine. Alternatively, you can install the start-up script into the `/etc/rc2.d` directory.

```
# cd /opt/SUNWconn/SUNWpf/utilities
# cp pf_fddidaemon /etc/rc2.d/S98pf_fddidaemon
```

Removing SunFDDI/P

If you remove the SunFDDI/P hardware interfaces from your machine, you also need to remove the SunFDDI/P software to prevent the device driver from starting each time you reboot the machine.

Use `pkgrm(1M)` to remove unbundled software from machines running a Solaris 2.x environment. Refer to *Installing Solaris Software* for detailed information on removing software packages using `pkgrm`.

▼ To Remove SunFDDI/P

1. Log in as root or become superuser.

2. Use `pkgrm` to remove SunFDDI/P.

The post-remove script removes the modifications made to the system files and updates the package information on your machine.

```
# /usr/sbin/pkgrm SUNWpfr
# /usr/sbin/pkgrm SUNWpft
```

3. Modify the `/etc/hosts` file to remove the IP addresses and host names for the SunFDDI/P interfaces that were inserted by the post-installation script.

4. Reboot the machine to ensure correct system operation.

Improving Network Performance

<i>High Performance FDDI Networks</i>	<i>page 24</i>
<i>Ring Latency Improvement</i>	<i>page 24</i>
<i>Data Throughput Improvement</i>	<i>page 25</i>
<i>Performance Across Bridges</i>	<i>page 25</i>
<i>Target Token Rotation Time (TTRT)</i>	<i>page 26</i>
<i>Improving Performance</i>	<i>page 27</i>

This chapter tells you how to optimize your FDDI network and assumes that you are familiar with FDDI network architecture and related terminology. See Chapter 6, “FDDI Network Architecture” for a detailed description of the FDDI protocols, and Chapter 7, “FDDI Network Topologies” for a detailed description of the FDDI network architecture.



Caution – The procedures for improving network performance are discussed in this chapter; however, these should be undertaken only by experienced system administrators. Tuning your network incorrectly can affect network performance adversely.

High Performance FDDI Networks

FDDI networks increase the potential to carry *more* information, rather than create faster connections. If the applications running over the network do not use the available bandwidth efficiently, you will not see much improvement in the performance of your network by using FDDI.

You can obtain optimum performance by balancing the complementary objectives of maximum throughput and minimum access delay:

- **Throughput** is a measure of the ring use. Network efficiency is defined as the ratio of actual throughput to theoretical maximum bandwidth (100 Mbps).
- **Access Delay** is the time a station waits before it can transmit on the network. This is largely dependent on how fast the permission token rotates around the ring.

For some applications (for example, bulk data transfers), throughput is critical; other applications (for example, voice or video applications) are more sensitive to access delay.

Ring Latency Improvement

The ring latency is the physical delay inherent in the FDDI network. It is dependent on the number of stations in the ring and the overall length of the ring. The speed of the microprocessor in the FDDI station also has a significant effect on ring latency.

The FDDI specifications define a network that supports up to one thousand physical connections with a total ring length of 200 km (100 km per ring in a dual-ring topology). There are two physical connections for each dual-attached station; therefore, in theory, up to five hundred dual-attached stations could be connected in a single ring. In reality, optimum performance is obtained by limiting the number of connections to less than one hundred stations in an unsegmented ring.

Performance improvements are possible by changing the topology of the network, trading off the total number of stations in comparison to the overall length of the ring, and selecting faster machines to act as FDDI stations. However, these modifications are rarely practical in an established network.

Data Throughput Improvement

The most significant improvements in throughput are achieved by maximizing the rate at which data is transferred to and from the FDDI network. If you are running TCP/IP or UDP/IP applications, you can improve throughput by modifying the parameters that control the rate at which these protocols transfer data to and from the SunFDDI/P device driver—that is, between user space and kernel space.

Send and Receive Buffers

The send and receive buffers control the maximum rate at which data is transferred between user space and kernel space by applications that use sockets. The size of these buffers defines the amount of data that can be transferred at one time.

To optimize performance, consistent values must be assigned for the send and receive buffers. The default value (4 Kbytes) assigned to these parameters is optimized for transmission across Ethernet connections. For applications running over FDDI, the send and receive buffers should be set to 32 Kbytes.

Performance Across Bridges

The *maximum transmission unit* (MTU) for each physical layer in the network defines the maximum size (in octets) of the transmitted packets. A larger MTU means that more data can be transmitted in a single packet; a larger MTU usually means higher throughput.

For SunFDDI/P, the MTU is set to 4352 by default¹, and for an Ethernet interface, the MTU is 1500. This difference can cause performance problems in mixed FDDI/Ethernet networks.

Packets are transmitted between the two networks across a bridge, which handles the translation between the two protocols. When data is transmitted across the bridge from the FDDI network, the larger FDDI packets need to be

1. This value makes allowance for the frame header and has been demonstrated to be the optimum value of MTU for pure FDDI networks (see RFC 1390)

fragmented so that they do not exceed the MTU of the Ethernet network. Some low-cost bridges that do not support fragmentation will reject the 4352-byte packets.

Fragmenting and reassembling the packets introduces a considerable overhead that affects performance. It also increases the risk of out-of-sequence or dropped packets. Therefore, if there is heavy network traffic across the bridge, it may be more efficient to force the MTU in the FDDI network to be 1500 to match the MTU on the Ethernet side of the bridge. This argument can also be applied to a mixed FDDI/NetWare® network. In this case, the MTU in the NetWare LAN is 4178 octets.

Target Token Rotation Time (TTRT)



Caution – Modifying the target token rotation time (TTRT) can affect network performance adversely. In the worst case, it can reduce network throughput to zero. In general, the TTRT should not be set to less than 8000 microseconds (ms).

The target token rotation time (TTRT) is the key FDDI parameter used for network performance tuning. In general, increasing the TTRT increases throughput and increases access delay.

For SunFDDI/P, the TTRT must be between 4000 and 165000 microseconds, and is set to 8000 ms by default. The optimum value for the TTRT is dependent on the application and the type of traffic on the network:

- If the network load is very irregular (*bursty* traffic), the TTRT should be set as high as possible to avoid lengthy queueing at any one station.
- If the network is used for the bulk transfer of large data files, the TTRT should be set relatively high to obtain maximum throughput, without allowing any one station to monopolize the network resources.
- If the network is used for voice, video, or real-time control applications, the TTRT should be set low to decrease access delay.

The TTRT is established during the *claim process*. Each station on the ring bids a value (T_{req}) for the operating value of the TTRT (T_{opr}) and the station with the lowest bid wins the claim. Setting the value of T_{req} on a single station does not guarantee that this bid will win the claim process.

Improving Performance

The following section describes how to modify the user-configurable network parameters. To obtain optimum network performance, you may need to tune these parameters, depending on your network configuration and the type of network traffic.

Changing the High and Low Water Marks

The maximum rate at which data is transferred between user space and kernel space by applications that use STREAMS is controlled by the *high water marks*. These parameters define the maximum amount of data that can be queued for transmission to the TCP or UDP STREAMS modules. There are independent send and receive high water marks for TCP and UDP applications.

When the number of bytes queued exceeds the high water mark, transmission is halted temporarily the backlog can be cleared. The *low water mark* specifies the level to which the queue must drop before transmission is restarted.

▼ To Tune the High Water Mark

Assign consistent values for the transmit and receive high water marks. The default value (8192) assigned to these parameters is optimized for transmitting across Ethernet connections. For applications running over FDDI, the high water marks should be set to 32 Kbytes. Use `ndd(1M)` to change the TCP and UDP high water marks:

1. **Log in as root or become superuser.**
2. **Use `ndd -get` (the default) to check the current value of the TCP high water marks (`tcp_xmit_hiwat` and `tcp_recv_hiwat`).**

```
# ndd /dev/tcp tcp_xmit_hiwat
8192
# ndd /dev/tcp tcp_recv_hiwat
8192
```

3. Use `ndd -set` to modify the current value of the TCP high water marks (`tcp_xmit_hiwat` and `tcp_rcv_hiwat`).

For optimum performance over FDDI connections, the TCP high water marks should both be set to 32 Kbytes.

```
# ndd -set /dev/tcp tcp_xmit_hiwat 32768
# ndd -set /dev/tcp tcp_rcv_hiwat 32768
```

4. Use `ndd -get` (the default) to check the current value of the UDP high water marks (`udp_xmit_hiwat` and `udp_rcv_hiwat`).

```
# ndd -get /dev/udp udp_xmit_hiwat
8192
# ndd -get /dev/udp udp_rcv_hiwat
8192
```

5. Use `ndd -set` to modify the current value of the UDP high water marks (`udp_xmit_hiwat` and `udp_rcv_hiwat`).

For optimum performance over FDDI connections, the UDP high water marks should both be set to 32 Kbytes.

```
# ndd -set /dev/udp udp_xmit_hiwat 32768
# ndd -set /dev/udp udp_rcv_hiwat 32768
```

▼ **To Tune the Low Water Mark**

By default, the low water mark is set to 2 Kbytes. Provided the high water mark is also tuned, increasing the low water mark can improve the performance slightly. The low water mark must never be set higher than the high water mark. Use `ndd(1M)` to change the TCP and UDP low water marks:

1. Log in as root or become superuser.

2. Use `ndd -set` to modify the current value of the TCP and UDP low water marks (`tcp_xmit_lowat` and `udp_xmit_lowat`).

For optimum performance over FDDI connections, set the TCP and UDP low water marks to 24 Kbytes.

```
# ndd -set /dev/tcp tcp_xmit_lowat 24576
# ndd -set /dev/udp udp_xmit_lowat 24576
```

Tuning the High and Low Water Marks at Boot Time

Each time you reboot your machine, the high and low water marks are reset to their default values. If you want to configure your system so that the high water marks are modified automatically each time you reboot the machine, you must put these commands in a start-up script.

An example script is located on the SunFDDI/P CD-ROM. Once installed on your machine, this script sets the TCP and UDP high and low water marks each time you reboot. The TCP and UDP high water marks are set to 32 Kbytes; the TCP and UDP low water marks are set to 24 Kbytes. You can modify the example script to enter your own values, if required.

Tuning the Socket Options

For applications that use Sockets, the application developer can make more efficient use of the increased bandwidth provided by an FDDI connection by increasing the size of the send and receive buffers used by a specific application. This manipulation is restricted to applications and does not affect the other socket connections.

Solaris 2.5.1 environments support program calls to `getsockopt(3N)` and `setsockopt(3N)`. The options `SO_SNDBUF` and `SO_RCVBUF` are used to set the size of the TCP send and receive buffers for a specified Socket. See the man page for more detailed information.

To optimize performance, assign consistent values for the transmit and receive buffers. The default value (4 Kbytes) assigned to these parameters is optimized for transmitting across Ethernet connections. For applications running over FDDI, the send and receive buffers should be set to 32 Kbytes.

MTU Path Discovery

Machines running a Solaris 2.x environment support MTU path discovery, which allows the optimum MTU to be negotiated. Under most circumstances, this ensures efficient use of the network resources. However, to enable MTU path discovery to work, the “don't fragment” bit in the packet is set, which causes a problem with some bridges and routers that do not support this feature. In this case, packets are not transmitted across the bridge, and the generated error messages as a result can be misleading.

Refer to “To Disable MTU Path Discovery” for instructions on how to accommodate bridges and routers that do not support this feature.

▼ **To Tune the Maximum Transfer Unit (MTU)**

The maximum transfer unit (MTU) has been optimized for pure FDDI networks for FDDI/P. Follow these steps to change the value of the MTU on a station running SunFDDI/P:

- 1. Log in as root or become superuser.**
- 2. Edit the `/etc/system` file to add the following line:**

```
set pf:pf_mtu=<packet_size>
```

The variable `<packet_size>` is set to the desired value for MTU entered in octets. The default value for MTU is 4352 octets. It must be in the range 512 to 4500 octets.

- 3. Reboot the machine to apply the changes to the system.**

Note - Additionally, the `ifconfig(1M)` command can be used to change the MTU.

▼ To Disable MTU Path Discovery

Many bridges do not support MTU path discovery.

◆ **To disable this feature in order to transmit packets across a bridge, type:**

```
# ndd -set /dev/ip ip_path_mtu_discovery 0
```

MTU path discovery is re-enabled each time the machine is rebooted.

▼ To Tune the Target Token Rotation Time (TTRT)

The target token rotation time (TTRT) for the network is the lowest value of `T_req bid` during the claim process. Follow these steps to change the value of `T_req bid`:

1. **Log in as root or become superuser.**
2. **Edit the `/etc/system` file to add the following line:**

```
set pf:pf_treq=8000
```

3. **Reboot the machine to apply the changes to the system.**

<i>Troubleshooting Checklist</i>	<i>page 34</i>
<i>Solving Common Problems</i>	<i>page 39</i>
<i>Running the Hardware Self-Test Program</i>	<i>page 44</i>
<i>Loading the Device Driver Manually</i>	<i>page 45</i>

This chapter describes how to detect and resolve common problems with your FDDI network. It includes instructions on how to load and configure the SunFDDI/P device driver manually, if required. Note that SunFDDI/P supports almost all of the standard SunOS™ network diagnostic utilities, including `snoop(1M)`, `netstat(1M)`, `ping(1M)`, and `etherfind(8C)`.

For a Solaris 2.5.1 Hardware: 4/97 environment, the default base directory `<basedir>` is:

```
/opt/SUNWconn/bin
```

Troubleshooting Checklist

Use the following checklist to verify the major components of your SunFDDI/P station, and to check that it is installed, configured, and attached to the network correctly.

▼ **To Check the Connection to the Station**

◆ **Check that the cable connector is seated firmly into the plug on the SunFDDI/P adapter card.**

You should feel the connector “click” into place.

▼ **To Check the Connections Between Stations**

◆ **Check that the cables are connected correctly between ports of different types.**

The keys provided with most FDDI cables will help you to avoid “crossing” or “twisting” the ring. See “Connecting Fiber Optic Cables” on page 5 for a detailed description.

- For single attached stations (SAS), cables should be connected by either S↔M (to a concentrator) or S↔S (back-to-back).
- For dual attached stations (DAS), cables should be connected by either A↔B B↔A (dual-ring) or A↔M M↔B (dual-homing).

▼ **To Check the Link Status Indicator**

◆ **Check the status of the SunFDDI/P interface by looking at the Link Status Indicator (diagnostic LED) mounted on the backplate of the SunFDDI/P PCI card.**

- If the LED is **green**, the SunFDDI/P driver is loaded, the pf interface is configured, and the station is attached to an active FDDI network. This is the operating state.

- If the LED is **amber**, the SunFDDI/P driver is loaded, the `pf` interface is configured, but no valid line state to an active FDDI network exists.
 - a. **Check the physical connections between the station and the ring, or the station and the concentrator.**
 - b. **Check the status of the neighboring station.**

If this station is down, it will bring down the link between the stations.
- If the LED is **red**, the adapter failed the FDDI path test.
- If no LED is lit, either the SunFDDI/P driver is loaded and the interface is *not* configured, or the driver has not been loaded.
- - a. **Configure the `pf` interface with the `ifconfig(1M)` command of the form:**

```
# ifconfig plumb pf<inst> <hostname> netmask <mask> up
```

- If the LED is **red**, the SunFDDI/P driver is not loaded, or is loaded incorrectly.
 - a. **Check that the software is installed.**
 - b. **Remove the existing version, if necessary, and reload the software.**

Checking the MAC Address

By default, the first SunFDDI/P interface installed in the machine adopts the host-resident Media Access Control (MAC) address stored on the motherboard. Each subsequent SunFDDI/P interface takes the card-resident MAC address stored in its IDPROM.

▼ To Check the SunFDDI/P Driver

- ◆ Use `netstat(1M)` to check that the SunFDDI/P (pf) driver is installed correctly, and to check for an excessive number of errors and collisions:

```

SunFDDI interface
pf<inst>---->
% netstat -i
Name Mtu Net/Dest Address Ipkts Ierrs Opkts Oerrs Collis
Queue
lo0 8232 loopback localhost 21 0 21 0 0 0
le0 1500 our-lan hysop 2146 0 950 1 13 0
pf0 4352 our-fddi hysop-pf 1086 0 907 0 0 0

```

A sudden increase in the number of errors could indicate a noisy connection caused by a dirty cable.

▼ To Check the IP Addresses

You must assign a unique IP address to each SunFDDI/P interface. If you have more than one IP interface installed in your machine—either multiple SunFDDI/P interfaces, or some other LAN interface (for example, an Ethernet interface)—each interface must be attached to a different subnetwork.

- ◆ Use `ifconfig(1M)` to check the IP addresses assigned to the two SunFDDI/P interfaces:

```

# ifconfig -a
pf0: flags=843(UP,BROADCAST,RUNNING,MULTICAST) mtu 4352
    inet 123:123:201:71 netmask ffffffff broadcast
    123:123:201:255
pf1: flags=843(UP,BROADCAST,RUNNING,MULTICAST) mtu 4352
    inet 123:123:201:75 netmask ffffffff broadcast
    123:123:201:255

```

If the interfaces have the same subnetwork number, as shown in this example, you must change the IP address assigned to one of the interfaces.

Ensure that the IP addresses and host names for each SunFDDI/P interface are entered in the NIS map, or in `/etc/hosts` on the local machine if you are not using NIS/NIS+.

▼ To Check the IP Routing

1. Check the IP routing table using the `netstat(1M)` command:

```
% netstat -nr
```

2. Check that the FDDI subnetwork is featured in the routing table.

If the routing table is empty, check that the routing daemon (`in.routed`) is running on your machine.

▼ To Check the Protocol Statistics

♦ Use `netstat(1M)` to check the per-protocol (IP, TCP, UDP, etc.) statistics for errors:

```
# netstat -s
UDP
  udpInDatagrams      = 1423udpInErrors      = 0
  udpOutDatagrams     = 540

TCP tcpRtoAlgorithm    = 4tcpRtoMin          = 200
  tcpRtoMax           = 60000tcpMaxConn     = -1
  tcpActiveOpens      = 38tcpPassiveOpens   = 1
  tcpAttemptFails     = 0tcpEstabResets     = 33
  tcpCurrEstab        = 2tcpOutSegs         = 427
  tcpOutDataSegs      = 264tcpOutDataBytes  = 15917
  tcpRetransSegs      = 38tcpRetransBytes   = 0
  tcpOutAck           = 143tcpOutAckDelayed = 56
  tcpOutUrg           = 1tcpOutWinUpdate    = 0
  tcpOutWinProbe      = 0tcpOutControl       = 111
  tcpOutRsts          = 36tcpOutFastRetrans = 0
  tcpInSegs           = 606
  tcpInAckSegs        = 292tcpInAckBytes     = 15949
  tcpInDupAck         = 62tcpInAckUnsent     = 0
  tcpInInorderSegs    = 311tcpInInorderBytes = 99169
  tcpInUnorderSegs    = 0tcpInUnorderBytes    = 0
.
.
.
```

◆ Use `netstat(1M)` to check the driver statistics:

```
# netstat -k
pfo
    ipacket = x, ierror = y, opackets = z
    .
    .
    .
```

▼ To Check the SMT Traffic

If the target station is not running SunFDDI/P, it does not necessarily support the same set of SMT frames. If the target station receives an SMT request for an unsupported service, it issues a Request Denied Frame (RDF).

◆ Use `pf_smtmon(1M)` to examine the SMT frames:

```
# <basedir>/pf_smtmon [-i pf<inst>] [-x] [-h] [<frameclass>]
```


Solving Common Problems

Use the information in Table 5-1 and Table 5-2 to diagnose and resolve some of the common problems that can occur when installing, configuring, or running SunFDDI/P.

Table 5-1 Problems Installing SunFDDI/P

Problem	Action
The SunFDDI/P software package (SUNWpfr) cannot be found.	Check that you have inserted the CD-ROM in the CD-ROM drive and that the CD-ROM is mounted on a local directory. If the Volume Manager (vold) is running on your machine, the SunFDDI/P software is located in the following directory: /cdrom/sun_fddip_1_0/Product. If the Volume Manager (vold) is not running on your machine, you must create a directory and mount the CD-ROM as described on page 13.
The SunFDDI/P man pages/utilities package (SUNWpft) cannot be found.	
The configuration script cannot configure the IP address. “What ip address do you want to use for xyz [] [?] 0.0.0.0 Illegal ip address. . . Try again”	Check that you have entered the valid IP address. The script will keep prompting you until you enter a valid address.
The configuration script cannot find the hardware devices installed in the machine.	Check that the SunFDDI/P PCI card is installed correctly and is seated firmly in the PCI slot. Use the prtconf command to see if the pf card exists.
The configuration script cannot load the SunFDDI/P driver. The procedure fails with the message, “no available major numbers.”	The number of major numbers—that is, the maximum number of device drivers that can be installed—is limited to 127. This number is quickly exhausted by the large number of drivers installed in a typical system. Remove an existing driver before loading SunFDDI/P.
The configuration script cannot load the device driver.	Check whether the driver is already installed on the system. If you removed a previous version of the package using pkgrm(1M), you must reboot the system before attempting another add_drv.
The configuration script cannot configure the device driver. The procedure fails with the message, “Duplicate MAC address.”	Two or more interfaces have been assigned the same MAC address. The most likely conflict lies between the first SunFDDI/P 1.0 interface (pf0) and one or more SunFDDI/P interfaces (pf1, etc) installed in the same machine.

Table 5-2 Problems Running SunFDDI/P

Problem	Action
The link status indicator (diagnostic LED) remains red.	Check that the SunFDDI/P software is installed correctly and that the driver is loaded and configured.
The link status indicator (diagnostic LED) remains amber.	Check that the SunFDDI/P PCI card is connected to the network.
The link status indicator (diagnostic LED) remains amber, even after the local station is connected to the network and the driver is configured and loaded.	Check that the other end of the cable is connected and that the neighboring station or concentrator is configured correctly. Some concentrators have diagnostic LEDs that indicate if the ring is operating correctly. Check that the ring is not crossed.
The SunFDDI/P interface (pf) is not displayed by netstat.	Check that the SunFDDI/P device driver is loaded, using modinfo(1M) to display information about loaded kernel modules. If the device driver is not loaded, see “Loading the Device Driver Manually” on page 45 for detailed instructions.
The SunFDDI/P interface (pf) is not displayed by ifconfig.	Check that the SunFDDI/P device driver is loaded, using modinfo(1M) to display information about loaded kernel modules. If the device driver is not loaded, see “Loading the Device Driver Manually” on page 45 for detailed instructions.
The local station can reach FDDI stations located on the same subnetwork, but these stations cannot reach the local station.	Check that the IP address and host name of the local station is entered in the NIS map or NIS+ tables (or in /etc/hosts on each remote station if you are not running NIS or NIS+).
The local station cannot reach other FDDI stations located on the same subnetwork.	<p>Check that the link status indicator is green, indicating that the SunFDDI/P interface is attached to an active network. Use ifconfig(1M) to check that the SunFDDI/P interface (pf) is up.</p> <p>Check that the IP address and host name of the remote station is entered in the NIS map or NIS+ tables (or in /etc/hosts on the local station if you are not running NIS or NIS+).</p>
Running pf_stat without the -m option (to display information about the local station) shows the ring alternating frequently between UP and DOWN.	<p>Check that the link status indicator is green, indicating that the SunFDDI/P interface is attached to an active network.</p> <p>Check the connections to the concentrator. Some concentrators have diagnostic LEDs that indicate if the ring is operating correctly.</p>

Table 5-2 Problems Running SunFDDI/P (Continued)

Problem	Action
<p>Running <code>pf_stat</code> without the <code>-m</code> option (to display information about the local station) shows that the ring is <code>ISOLATED</code>.</p>	<p>Check that the link status indicator is green, indicating that the SunFDDI/P interface is attached to an active network. Check the connections to neighboring stations or the concentrator. Some concentrators have diagnostic LEDs that indicate if the ring is operating correctly.</p> <p>Check that the neighboring stations are configured correctly. If everything else seems to be working, check the cable itself. Fiber optic cable is fragile and should not be subjected to stress or impact. If the turning curve around objects is too tight, the light path is interrupted. Clean the connectors.</p>
<p>The local station cannot reach FDDI stations located on a remote FDDI network.</p>	<p>Check that the IP address and host name of the remote station is entered in the NIS map or NIS+ tables (or in <code>/etc/hosts</code> on each remote station if you are not running NIS or NIS+).</p> <p>Check that the routing tables are using <code>netstat -r</code>. Check for an <code>pf</code> interface that shows a route to a gateway with the <code>UG</code> flags set. If you can reach the local gateway, then the problem probably lies with the IP routing.</p> <p>Try to confirm that the remote station is configured and running, and if possible try to reach the local station from the remote station. If you can, it indicates that the problem lies with the local IP routing.</p> <p>Try to reach the router nearest the remote station. If you can, it indicates that the fault lies between the remote station and the router. If you cannot, try to reach a router that is closer to the local station. Continue in this way until you have isolated the router that is dropping or misrouting packets.</p> <p>If the dynamic routing protocol is not adding routes, try adding a static route to the remote station. This method is not recommended for large networks with a large number of nodes.</p>

Table 5-2 Problems Running SunFDDI/P (Continued)

Problem	Action
<p>The local station cannot reach FDDI stations located on a remote Ethernet network.</p>	<p>Check that the IP address and host name of the remote station is entered in the NIS map or NIS+ tables (or in <code>/etc/hosts</code> on each remote station if you are not running NIS or NIS+).</p> <p>If you are operating in a Solaris 2.x environment, which uses MTU path discovery, check that packets are being transmitted across the bridge between the networks. Many bridges do not yet support MTU path discovery and do not fragment the large packets sent by the FDDI stations. In this case you may need to disable this feature on your station by typing:</p> <pre># ndd -set /dev/ip ip_path_mtu_discovery 0</pre> <p>Check that the routing tables are using <code>netstat -r</code>. Check for an <code>pf</code> interface that shows a route to a gateway with the <code>UG</code> flags set. If you can reach the local gateway, then the problem probably lies with the IP routing.</p> <p>Try to confirm that the remote station is configured and running, and if possible try to reach the local station from the remote station. If you can, it indicates that the problem lies with the local IP routing.</p> <p>Try to reach the router nearest the remote station. If you can, it indicates that the fault lies between the remote station and the router. If you cannot, try to reach a router that is closer to the local station. Continue in this way until you have isolated the router that is dropping or misrouting packets.</p> <p>If the dynamic routing protocol is not adding routes, try adding a static route to the remote station. This method is not recommended for large networks with a large number of nodes.</p>
<p>Running <code>pf_stat</code> with the <code>-m</code> option (to display information about the neighboring stations) shows frequent <code>error frames</code> and <code>lost frames</code>.</p>	<p>This could indicate a noisy network. For fiber connections, noise is probably caused by dirt or grease on the optical surface or by a damaged cable. Unshielded twisted-pair is sensitive to electrical and electromagnetic interference. Note that the receipt of <code>error frames</code> and <code>lost frames</code> does not give any indication of where the problem is occurring on the ring.</p>

Table 5-2 Problems Running SunFDDI/P (Continued)

Problem	Action
Running <code>pf_smtmon</code> shows the frequent occurrence of request denied frames (RDF).	The SMT entity on one of the remote stations does not support the same set of SMT services. This should not occur on a station running SunFDDI/P. You can use this facility to detect problems communicating with a remote station using the SunNet Manager proxy agent.

Table 5-3 Problems Using the SNM Agents for SunFDDI/P

Problem	Action
Cannot select <code>pf_fddi</code> or <code>pf_fddismt7_2</code> agent for quick dump.	Check that the SNM agents for SunFDDI/P are enabled in the properties window for the proxy station. See “Using the SunFDDI/P Local Agent (<code>pf_fddi</code>)” on page 110 for detailed instructions.
Cannot recover SMT statistics from the proxy agent. The procedure fails with the error message, “Cannot create RPC client: program is not registered.”	<p>Check that you specified the host name or MAC address for the target machine in the properties window for the target station. See “Using the SunFDDI/P Local Agent (<code>pf_fddi</code>)” on page 110 for detailed instructions.</p> <p>Check that the SNM daemons are started on the proxy and the target machines:</p> <pre data-bbox="829 991 1365 1104"># ps -ef grep snm root <pid> <date&time> <path>/pf_snmd root <pid> <date&time> <path>/pf_snmd root <pid> <date&time> <path>/pf_snmd</pre> <p>On a Solaris 2.x client, you can start the SunNet Manager daemons for SunFDDI/P by typing:</p> <pre data-bbox="829 1230 1105 1256"># /usr/sbin/pf_snmd</pre>

Running the Hardware Self-Test Program

If you suspect that there may be a problem with the SunFDDI/P adapter card, you can use the built-in hardware self-test to check the state of its primary components and the connection to the network.

▼ To Run the Hardware Self-Test

1. Log in as `root` or become superuser.
2. Halt the machine.

```
# sync;sync;halt
```

3. At the boot prompt, type:

```
ok test /pci/pf

Local MAC Address 0:80:d8:10:3:ed (Canonical)
Interrupt register read/write test .... PASS
CMT Processor read/write test ..... PASS
S Port register test ..... PASS
MAC register test ..... PASS
DMA register test ..... PASS
DMA bus master test ..... PASS
Performing FDDI path test ..... PASS
Connection Management (CMT) test ..... PASS
Local loopback packet test ..... PASS
Network loopback packet test ..... connection not active

/pci/pf selftest failed. Return code = -1
```

This example shows that the primary components on the board are working, but that the board is not currently connected to the network.

Loading the Device Driver Manually

Normally, the SunFDDI/P device driver is loaded and configured by the post-installation script, which is launched automatically when the software is installed. If you encounter problems when running this script, or if you want to customize the installation, you may need to load the device driver manually.

▼ To Configure and Load the Device Driver

1. Log in as `root` or become `superuser`.
2. Check that there are no partially installed drivers for SunFDDI/P.
 - a. Check for existing SunFDDI/P devices in the `/dev` directory.

```
# ls /dev | grep pf
pf
```

- b. Use `rem_drv(1M)` to remove any installed SunFDDI/P devices and remove the device directories.

```
# /usr/sbin/rem_drv pf
# /bin/rm -f /dev/pf
```

- c. Check the `/etc/name_to_major` file for entries of the form `pf<num>` and `smt<num>` and remove these if they exist.

3. Use `add_drv(1M)` to inform the system about the new drivers.

```
# /usr/sbin/add_drv pf
```

4. Check that the device drivers have been added to the system correctly.
 - a. Look in the `/devices/pseudo` directory for entries of the form:

```
ls -l /devices/pseudo | grep pf
crw----- 1 root sys 11,119 <date> clone@0:pf
```

b. Look in the `/dev` directory for links to these entries:

```
ls -l /dev | grep pf
lrwxrwxrwx 1 root other <date> pf -> ../devices/pseudo/clone@0:pf
```

5. Create a file called `/etc/hostname.pf<inst>` for each SunFDDI/P IP interface that you configure.

Each file must contain the host name assigned to the IP interface. These files are used to configure the interfaces when the system is rebooted. If you assign the primary host name to one of the FDDI interfaces, this host name must be entered in the file `/etc/nodename`.

6. Use `ifconfig(1M)` to configure the SunFDDI/P IP interfaces (`pf<inst>`). The host name assigned to the IP interface must be entered in the NIS map, NIS+ tables, or in `/etc/hosts` on the local machine. Configure the interface using an `ifconfig` command of the form:

```
# /usr/sbin/ifconfig pf<inst> plumb <hostname> netmask + up
```

Note that the modifier `netmask +` takes the netmask defined in the `/etc/netmasks` file. If this entry is missing, you must specify the netmask explicitly using dot notation.

Provided the installation and configuration was successful, your SunFDDI/P station is now active and you can send and receive data transparently across an FDDI connection.

*Part 2 — Planning and Implementing
SunFDDI Networks*

<i>FDDI Network Model</i>	<i>page 50</i>
<i>FDDI Network Architecture</i>	<i>page 55</i>
<i>FDDI Failure Recovery</i>	<i>page 60</i>
<i>FDDI Ring Operation</i>	<i>page 65</i>
<i>FDDI Performance</i>	<i>page 66</i>

This chapter provides a brief introduction to the Fiber Distributed Data Interface (FDDI), and the network architecture described by the following specifications:

- ANSI/FDDI Physical Media Dependent (PMD) X3.166-1990
- ANSI/FDDI Physical Layer (PHY) X3.148-1988
- ANSI/FDDI Medium Access Control (MAC) X.3.139-1987
- ANSI/FDDI Station Management (SMT) X3.299 R7.3 (formally R7.2.99)

FDDI Network Model

FDDI provides high-performance, multistation networking at data transfer rates of up to 100 Mbps. It is based on a Token-Ring network architecture, and provides communication over optical fiber or copper twisted-pair connections.

The ANSI/FDDI specifications define a network model that consists of the following components:

- Physical Medium Dependent (PMD) Layer
- Physical (PHY) Layer
- Medium Access Control (MAC) Layer
- Station Management (SMT) Protocol

Figure 6-1 shows the organization of these components, and their relationship with other network protocols that use them.

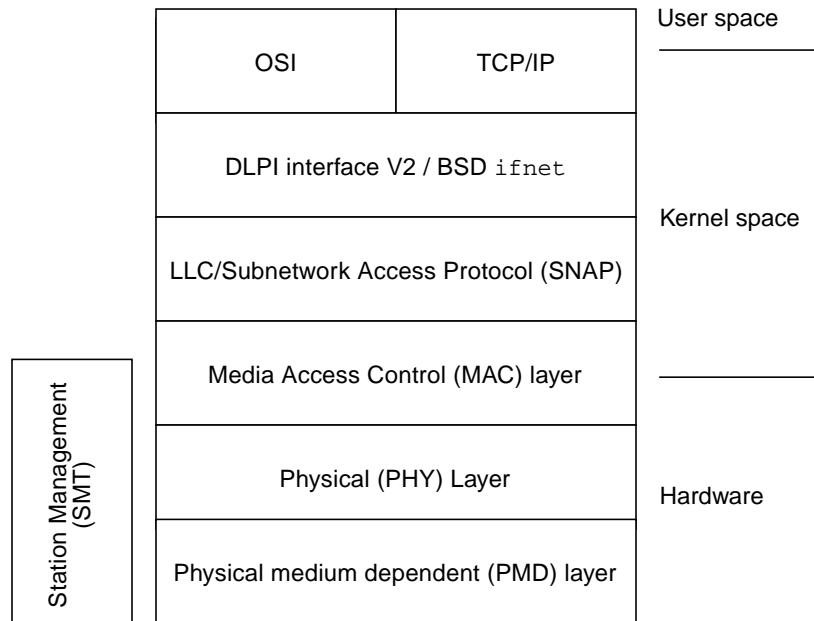


Figure 6-1 FDDI Architectural Model

Physical Medium Dependent (PMD) Layer

This portion of the FDDI network model defines the physical medium used to carry the encoded digital signal. It is also referred to as the *media* layer.

The PMD layer determines the characteristics of the transmitters, receivers, connectors, and cables used to attach the FDDI station to the network. SunFDDI/P supports only the Multimode fiber (MMF)—single and dual-attached PMD option.

Multi-Mode Fiber (MMF)

A fiber optic FDDI connection supports data transmission rates of up to 100 Mbps and, because attenuation in the fiber is very low, larger distances between FDDI stations are possible without the use of repeaters. In addition, fiber does not emit electromagnetic radiation and is electrically nonconductive. Therefore, it neither causes, nor suffers from, problems due to signal interference between adjacent equipment or cabling.

SunFDDI/P supports 62.5/125 micrometer, multimode fiber (MMF). Multimode fiber cannot support such large distances between FDDI stations as single-mode fiber (SMF), which is not supported by SunFDDI/P. However, multimode fiber is cheaper, and uses lower cost terminations, connectors, and light-sources.

Physical (PHY) Layer

The physical (PHY) layer handles the efficient encoding and decoding of digital data. It is also referred to as the *signal* layer.

At the transmitting end, the PHY layer encodes the digital data into FDDI *symbols* and passes them to the physical medium dependent (PMD) layer for transmission. At the receiving end, the PHY layer decodes the FDDI symbols and passes the digital data to the medium access control (MAC) layer. An FDDI symbol is the basic transmission unit in an FDDI network.

The PHY layer also provides the network synchronization mechanism. It uses a distributed clocking scheme under which each FDDI station has its own local clock to synchronize the transmission of outgoing data. The PHY layer extracts clocking information from incoming data as it is received.

Media Access Control (MAC) Layer

The Media Access Control layer specifies the access mechanism used to transmit and receive data on the FDDI network. It packages digital data in *frames*.

The MAC layer specifies three classes of digital data traffic:

- Synchronous (guaranteed) traffic
- Asynchronous (priority-based) traffic
- Restricted (dialogue-based) traffic

The MAC layer uses a *timed token rotation protocol* that regulates how much digital data can be sent at one time. This protocol ensures that the network bandwidth is used predictably and efficiently. It also gives the FDDI network its distinctive “ring” topology.

The FDDI stations connected on the network use a token to control the right to transmit data for a predefined time, determined by the local *timed target rotation timer* (T_{TRT}) on each station. When an FDDI station completes transmission, it releases the token for use by downstream stations. In this way, the token rotates continuously around the ring.

Each attachment to an FDDI network is identified by a unique MAC address. The first SunFDDI/P card installed in a machine takes its identity from the host-resident MAC address that is stored in nonvolatile memory on the motherboard of the machine in which it is installed. Each subsequent SunFDDI/P SBus card takes its identity from the card-resident MAC address stored in its own IDPROM.

Use the `pf_macid(1M)` utility to display the card-resident MAC address. See Chapter 8, “Using the SunFDDI Network Utilities” for detailed instructions.

Station Management (SMT) Layer

The Station Management layer provides services to manage, control, and configure the FDDI network. It is also referred to as the *network management* layer.

At its lowest level, the SMT protocol handles *connection management* (CMT). This includes station initialization, the insertion and removal of stations on the network, and connection compatibility between stations. At a higher level, the SMT protocol handles *ring management* (RMT). This includes the detection of duplicate addresses and the isolation of fault conditions.

The SMT protocol also defines the FDDI *management information base* (MIB). This is a set of managed objects and associated attributes that includes the MAC entity (MAC), data path (PATH), attachment type (ATTACHMENT), and port identifier (PORT).

Communication services are used to exchange objects and information between peer SMT entities through special management frames that are carried at the same time as normal network traffic. This frame-based management is used to:

- Examine and modify FDDI station configuration
- Schedule synchronous (guaranteed) and asynchronous (priority-based) traffic
- Gather network statistics
- Generate status reports

SunFDDI/P includes SunNet Manager agents that access the SMT entity to collect and return FDDI statistics to a SunNet Manager console. See Chapter 9, “Managing FDDI Stations Using SunNet Manager” for detailed instructions on how to manage FDDI networks using a SunNet Manager console.

Communication Between FDDI Layers

The MAC layer transfers data between peer entities in the form of frames that are encoded as FDDI symbols by the PHY layer. The PMD layer transmits FDDI symbols across the network as a bit stream of light pulses (for fiber connections) or electrical signals (for twisted-pair connections).

The communication between layers in the FDDI architectural model is summarized in Figure 6-2.

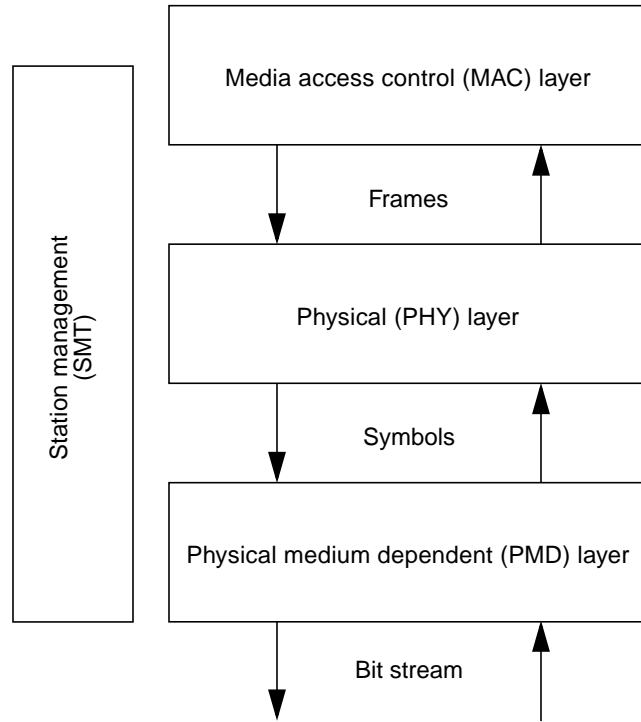


Figure 6-2 Communication Between FDDI Layers

FDDI Network Architecture

A typical FDDI network is based on a dual, counter-rotating ring, as illustrated in Figure 6-3. Each FDDI station is connected in sequence to two rings simultaneously—a primary ring and a secondary ring. Data flows in one direction on the primary ring, and in the other direction on the secondary ring.

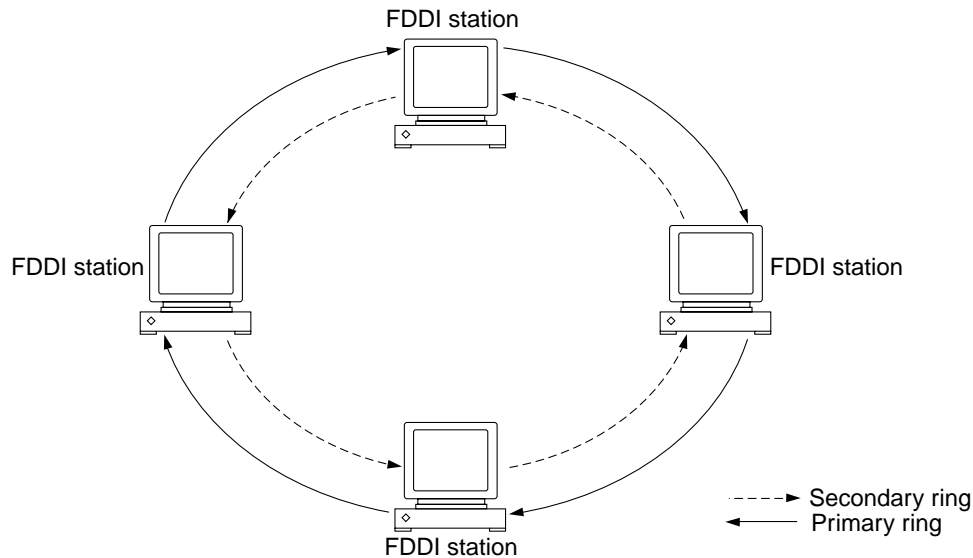


Figure 6-3 Basic FDDI Network Architecture

The secondary ring serves as a redundant path. It is used during station initialization, and may be used as a backup to the primary ring in the event of a station or cable failure. When a failure occurs, the dual-ring is “wrapped” around to isolate the fault and to create a single, one-way ring. The components of a typical FDDI network and the failure recovery mechanism are described in more detail in the following sections.

FDDI Stations

An FDDI station is any device that can be attached to a fiber or copper twisted-pair FDDI network through an FDDI interface. The FDDI protocols define two types of FDDI stations:

- Single-attached station (SAS)
- Dual-attached station (DAS)

Single-Attached Station (SAS)

A single-attached station (SAS) is attached to the FDDI network through a single connector called the S-port. The S-port has a primary input (Pin) and a primary output (Pout). Data from an upstream station enters through Pin and exits from Pout to a downstream station, as shown in Figure 6-4.

Single-attached stations are normally attached to single- and dual-attached concentrators, as described in “FDDI Concentrators” on page 58.

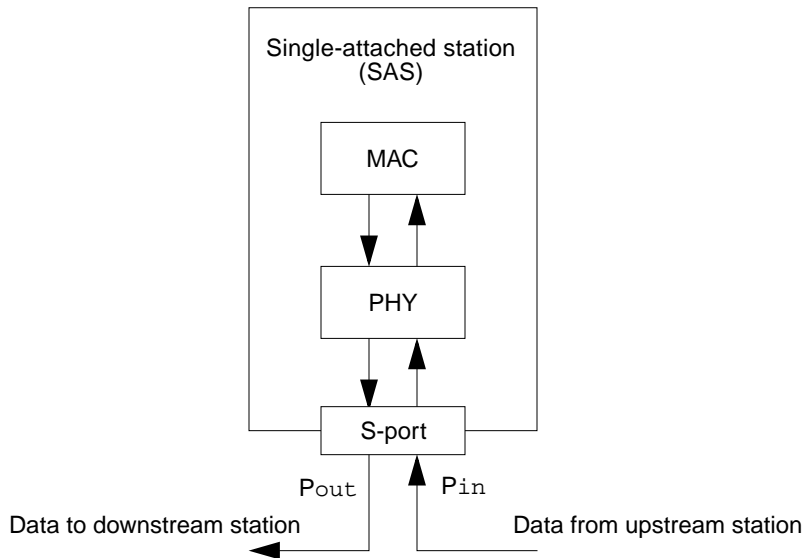


Figure 6-4 Single-Attached Station (SAS)

Dual-Attached Station (DAS)

A dual-attached station (DAS) is attached to the FDDI network through two connectors called the A-port and the B-port, respectively. The A-port has a primary input (Pin) and a secondary output (Sout); the B-port has a primary output (Pout) and a secondary input (Sin).

The primary input/output is attached to the primary ring and the secondary input/output is attached to the secondary ring. The flow of data during normal operation is shown in Figure 6-5.

To complete the ring, you must ensure that the B-port of an upstream station is always connected to the A-port of a downstream station. For this reason, most FDDI DAS connectors are keyed to prevent connections between two ports of the same type.

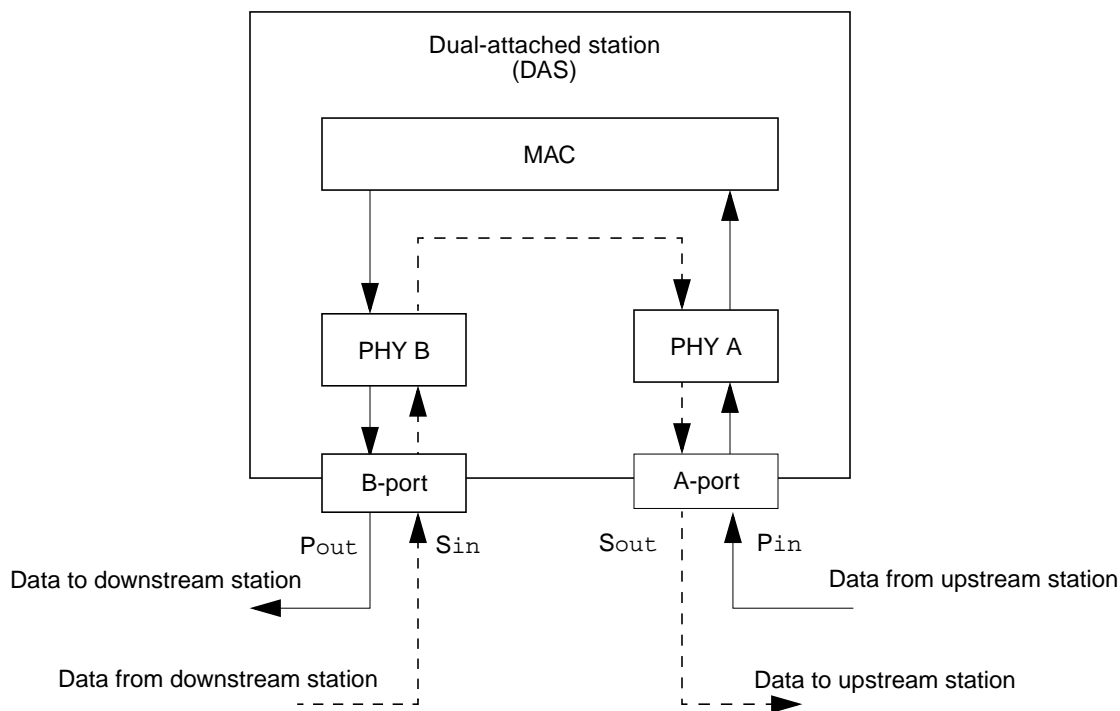


Figure 6-5 Dual-Attached Station (DAS)

FDDI Concentrators

FDDI concentrators are multiplexers that attach multiple single-attached stations to the FDDI ring. An FDDI concentrator is analogous to an Ethernet hub.

The FDDI protocols define two types of concentrators:

- Single-attached concentrator (SAC)
- Dual-attached concentrator (DAC)

Single-Attached Concentrator (SAC)

A single-attached concentrator (SAC) is attached to the FDDI network through a single connector, which is identical to the S-port on a single-attached station. It has multiple M-ports, to which single-attached stations are connected, as shown in Figure 6-6.

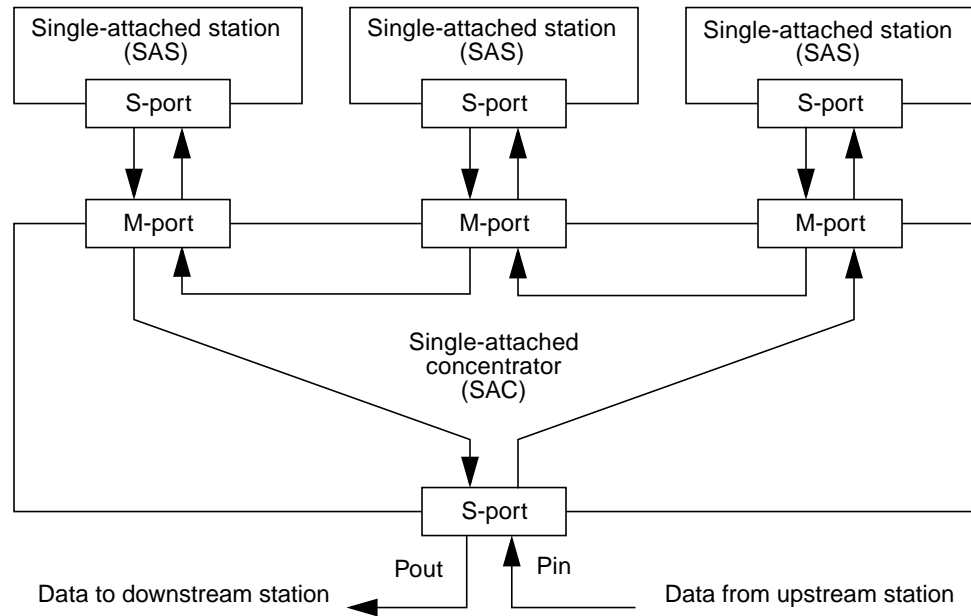


Figure 6-6 Single-Attached Concentrator (SAC)

Dual-Attached Concentrator (DAC)

A dual-attached concentrator (DAC) is attached to the FDDI network through two ports—the A-port and the B-port, which are identical to the ports on a dual-attached station. It has multiple M-ports, to which single-attached stations are connected, as shown in Figure 6-7.

Dual-attached concentrators and FDDI stations are often arranged in a very flexible network topology called the *ring of trees*. This configuration is discussed in more detail in Chapter 7, “FDDI Network Topologies.”

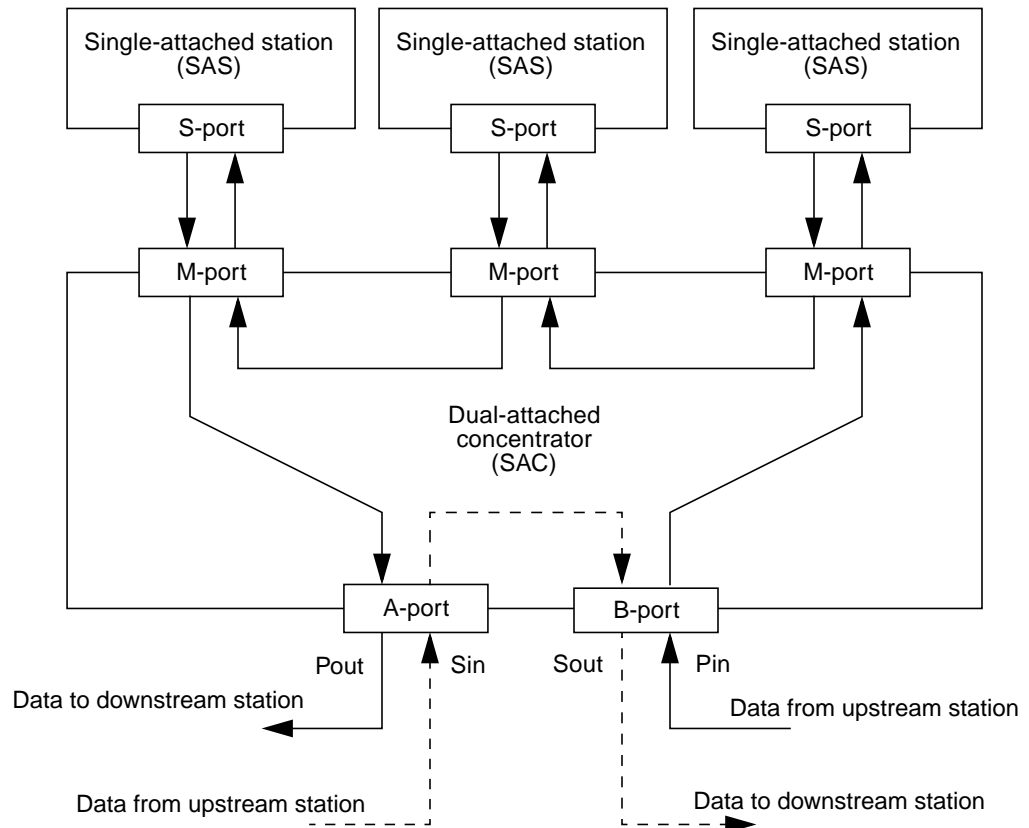


Figure 6-7 Dual-Attached Concentrator (DAC)

FDDI Failure Recovery

One of the primary advantages of FDDI is its ability to recover reliably from failures in stations and cables. The failure mechanism is implemented and controlled by the Station Management (SMT) entity described in the section “Station Management (SMT) Layer” on page 52.

A failure could be something as insignificant as someone switching off their workstation; the ability of FDDI to recover from such an event increases the reliability of the network significantly.

Station Wrapping

Figure 6-8 shows the effect of a single-station failure in an FDDI network that comprises four dual-attached stations connected in a basic ring configuration. When a failure occurs, the SMT entities on the stations on either side of the fault reconfigure the network dynamically to isolate it. In this condition, the primary and secondary rings are wrapped on to each other to form a single, one-way ring.




Figure 6-8 Isolating a Single Station Failure

This failure recovery mechanism is only supported by dual-attached stations and concentrators; a single-attached station connected directly to the ring cannot wrap around the fault because it is only connected to one ring at a time.

The effect of a cable or link failure on a basic FDDI network is very similar to a station failure, as shown in Figure 6-9. The SMT entities on the stations on either side of the failed connection reconfigure the network dynamically so that the primary ring is wrapped onto the secondary ring.

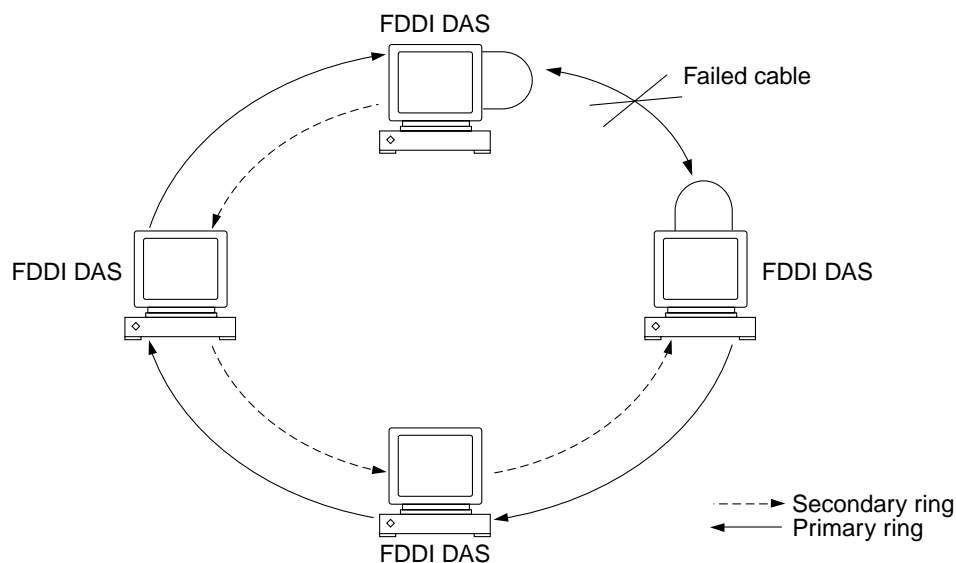


Figure 6-9 Isolating a Single Cable Failure

Wrapping occurs within the MAC layer. Figure 6-10 on page 62 illustrates how the dual-attached station is reconfigured to wrap the primary ring onto the secondary ring. Dual-attached concentrators wrap in a similar way. In this case, all of the single-attached stations connected to the concentrator are also wrapped onto the secondary ring.

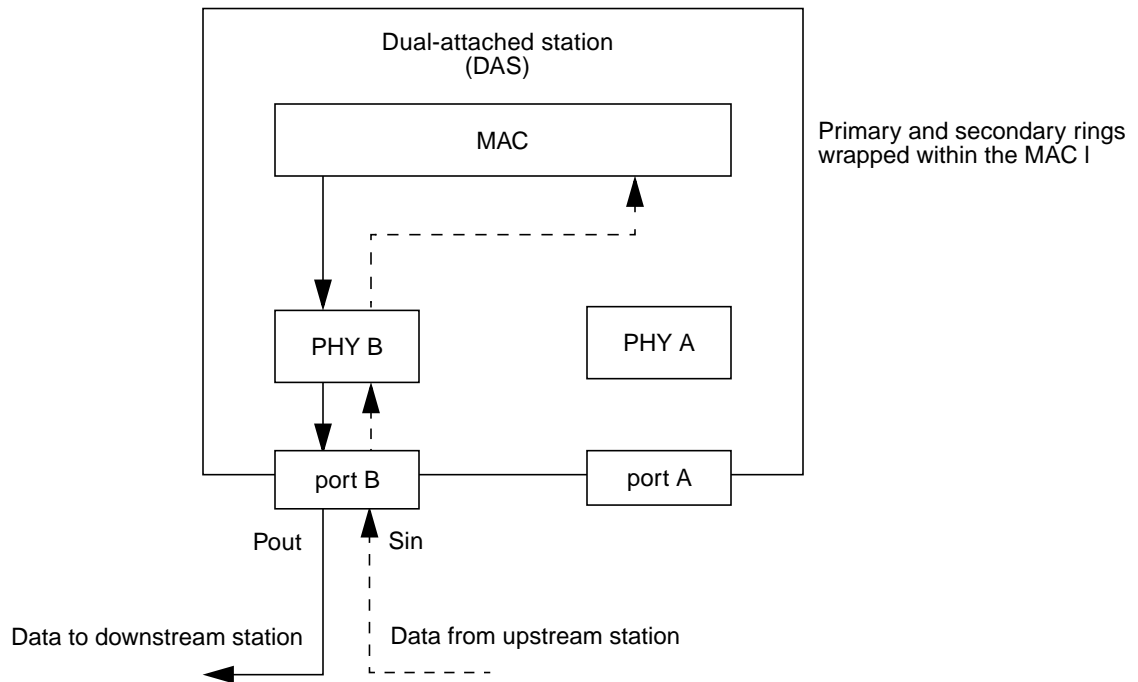


Figure 6-10 Dual-Attached Station in Wrap Mode

Optical Bypass Switches

Station wrapping provides effective network recovery in the event of a single station or cable failure. However, two or more failures in the ring will isolate portions of the network, as shown in Figure 6-11.

This limitation can be overcome to a certain extent by fitting *optical bypass switches* to the dual-attached stations. As their name suggests, these switches provide an optical connection that bypasses the station in the event of failure, or if the station needs to be removed from the ring. Optical bypass switches can only be used to overcome station faults; they have no effect in the event of a cable fault.

The attenuation in an optical bypass switch is far greater than in a normal FDDI connection; therefore, the number of optical bypass switches that can be connected in series in a single ring is limited.




Figure 6-11 Dual-Ring Network Divided by Two Faults

The maximum number of active optical bypass switches that can be connected in a series is four. This assumes that the maximum distance between stations in the ring is no more than 400m; otherwise, the aggregate attenuation in the ring exceeds the total optical power budget.

Figure 6-12 shows how connecting an optical bypass switch to a dual-attached station is used to conserve the connection when the station is switched off, or removed from the ring.

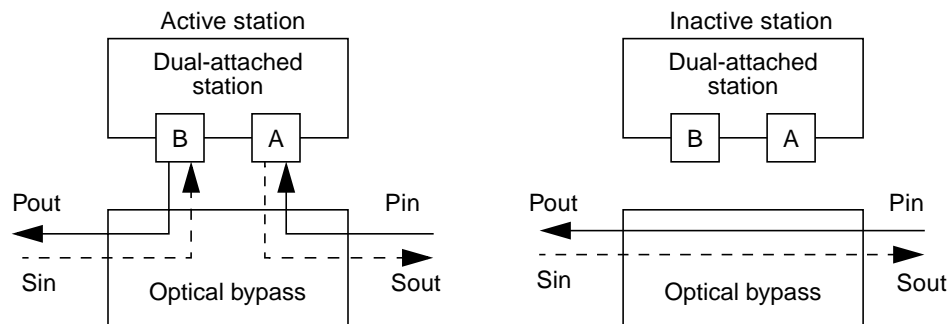


Figure 6-12 Optical Bypass Switch

Figure 6-13 shows the occurrence of two different fault conditions in an FDDI network that includes optical bypass switches. The station fault is bypassed effectively to conserve the majority of the network intact; however, the cable fault still causes station wrapping.

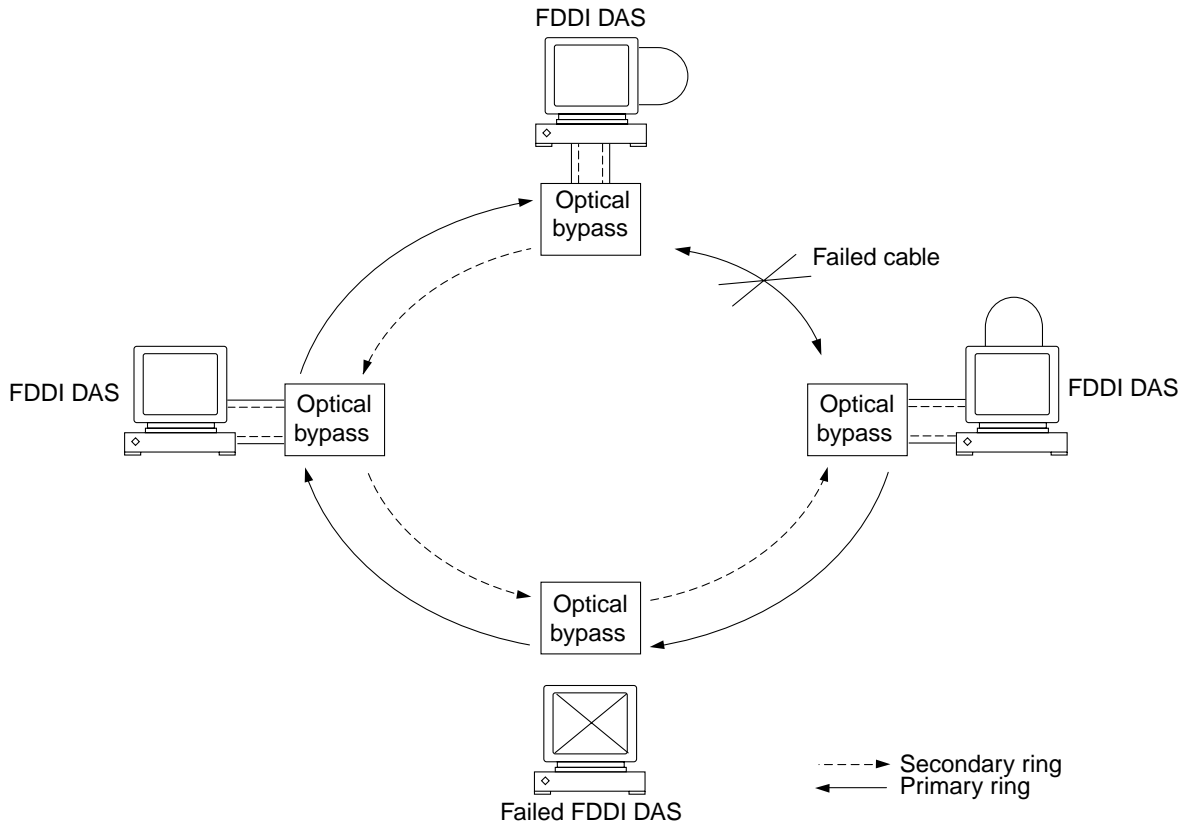


Figure 6-13 Optical Bypass Switches used in a Network

FDDI Ring Operation

Two types of traffic are allocated bandwidth on an FDDI ring:

- **Asynchronous traffic** (unrestricted and restricted)

The unrestricted asynchronous service allocates bandwidth dynamically based on the *timed token rotation protocol*, and according to the priority assigned to the traffic. It does not guarantee bandwidth; therefore, unrestricted asynchronous traffic may be delayed on a heavily loaded network.

The restricted asynchronous service allocates bandwidth for extended transmission based on *restricted token exchange* between communicating stations. During restricted transmission, the timed token rotation protocol is suspended until the restricted token is released. The use of the restricted asynchronous service is relatively rare.

Unrestricted Asynchronous Transmission

The FDDI network is most commonly used for unrestricted asynchronous traffic. This is the default mode for SunFDDI and is the only mode supported in a Solaris 1.x environment. Access to the FDDI ring is controlled by a *permission token*, or right to transmit. The permission token is a special frame that is passed between stations. Unlike the tokens used in a Token Ring network, the FDDI permission token contains no additional information.

The station holding the permission token completes transmission either when it has no more packets to transmit, or when its right to hold the token expires. It then releases the token, which is now available for use by the next station on the ring. In this way, the permission token rotates around the ring, at a rate determined by the overall *target token rotation time* (TTRT).

The TTRT is established when the ring is initialized, based on a bidding procedure called the *claim process*. During the claim process, each station puts in a request (T_{req}) for the TTRT. The lowest bid or, the fastest rotation time wins the claim process and is stored in 2's complement form as the operating value of TTRT (T_{opr}). The maximum value of TTRT (T_{max}) allowed by the MAC layer is implementation dependent, but is always in the ~165 ms range. SunFDDI/P has $T_{max}=167.874$ ms.

The maximum time for which a station can hold the permission token, and therefore the time for which a station can transmit on the ring, is determined by these two station timers that work together to maintain the target token rotation time (TTRT), irrespective of the load on the network:

- **Token Rotation Timer (TRT)**

The TRT measures the time between successive arrivals of the token or, the time taken for the token to rotate once around the ring. The TRT is reset to the operating value of TTRT (T_{opr}) each time the token is received. If the TRT expires before the token arrives, the token is considered late and the flag `Late_Ct` is set. Only synchronous (high priority) traffic can be transmitted when `Late_Ct` is set.

- **Token Hold Timer (THT)**

The THT determines how long a station can hold a token or, how long a station can transmit asynchronous traffic. It is loaded with the value remaining in the token rotation timer (TRT) each time the token arrives. Thus, the faster the token rotates, the more transmission time is allocated to each station. If the token is delayed because the network is heavily loaded, the amount of transmission time is reduced. If THT expires during transmission, the current transmission is completed before the token is released.

The interaction of these two timers causes token rotation to increase (reduced transmission time per station) in a heavily loaded network, and to decrease (increased transmission time per station) in a lightly loaded network.

FDDI Performance

Although the FDDI standards define a medium that supports data transfer rates of up to 100 Mbps, this does not translate directly to an increase in overall system performance. In fact, it frequently transfers the bottleneck elsewhere, so that the FDDI connection is not used to full capacity.

An FDDI ring has the *potential* to carry more information rather than to be a faster connection. If the applications running over the network do not use the available bandwidth efficiently, you will not see much improvement in the performance of your network above that of traditional Ethernet connections.

Among the factors that affect FDDI network performance are:

- Network topology

-
- Network load
 - Network efficiency
 - Processor speed (on the FDDI stations)
 - Bus architecture

There are a number of ways of improving the overall performance of the network, some of which are discussed in Chapter 4, “Improving Network Performance.” However, the majority of these suggestions should only be undertaken by experienced system administrators. Any improvement made by modifying the actions of the FDDI network is negligible compared to what is gained by making more efficient use of the available bandwidth.

<i>Basic Dual-Ring Network</i>	<i>page 70</i>
<i>Standalone Concentrator</i>	<i>page 72</i>
<i>Concentrators with Dual-Homing</i>	<i>page 74</i>
<i>Tree of Concentrators</i>	<i>page 75</i>
<i>Ring of Trees</i>	<i>page 77</i>
<i>Mixed FDDI/Ethernet Networks</i>	<i>page 79</i>

FDDI networks can be arranged in a variety of ways, depending on the placement of stations (SAS and DAS) and the use of concentrators (SAC and DAC). The optimum arrangement for a particular installation is dependent on several factors, including:

- Cost
- Network size
- Required Bandwidth
- Type of network traffic
- Fault resistance and network reliability

This chapter describes some of the common FDDI network topologies and discusses the primary advantages and disadvantages of each.

Basic Dual-Ring Network

The dual ring (or dual, counter-rotating ring) is one of the simplest FDDI network topologies. It clearly illustrates the distinctive ring architecture most commonly associated with the FDDI standards, as shown in Figure 7-1. Each station is critical to the operation of the network; therefore, the basic dual-ring topology is best adapted to small, stable networks that are not subject to frequent reconfiguration.

In a dual-ring network, dual-attached stations are connected directly to the primary and secondary rings. Data, and the token that controls the flow of data, are transmitted in one direction on the primary ring. Data flows in the other direction on the secondary ring, which is used during ring initialization and as a backup in case of a ring failure.

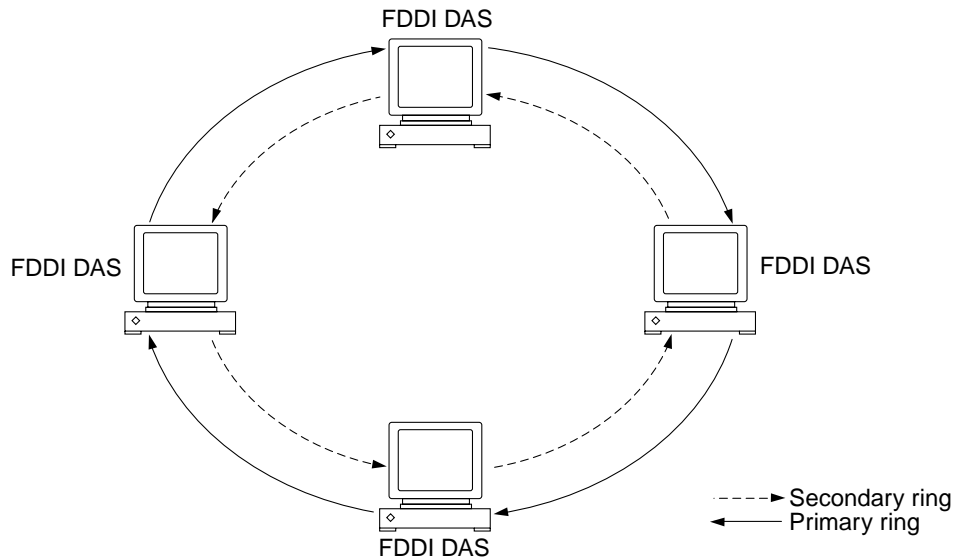


Figure 7-1 Basic Dual-Ring Network

Advantages

The primary advantages of the dual-ring network topology are its simplicity and its ability to recover from simple station and line faults. The secondary ring provides an effective backup in the event of a *single* failure in the ring.

When a ring fails, the primary ring is wrapped automatically on either side of the fault so that the primary and secondary rings are combined to form a single, one-way ring. This mechanism is described in more detail in “FDDI Failure Recovery” on page 60. A dual-ring network does not require a concentrator.

Disadvantages

Although the dual-ring topology is resistant to single failures in the ring, two or more failures break the network into parts. Small fragments of the network can still function, but they are isolated from the other stations. Figure 7-2 shows how two faults in a network, with five dual-attached stations, isolates two parts of the network. Cable and connection costs can be high in large installations because there are two cables between each station.

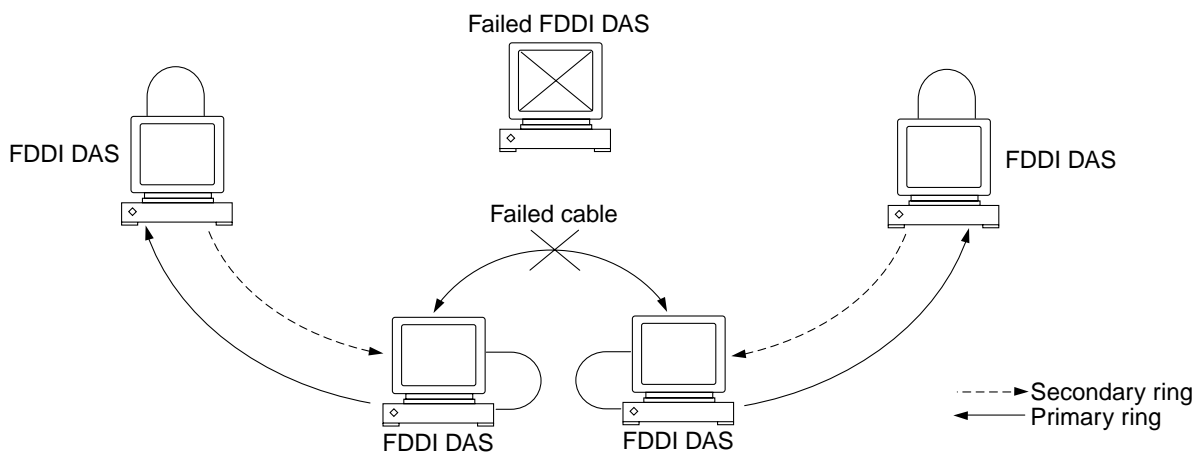


Figure 7-2 Basic Dual-Ring Network with Two Faults

Standalone Concentrator

Figure 7-3 shows multiple single-attached stations connected to a single, dual-attached concentrator through its M-ports. The concentrator can also be connected to an external dual ring through its A- and B-ports.

A standalone concentrator provides a stable, low-cost alternative for small work groups that do not require the fault recovery facility provided by the dual-ring configuration.

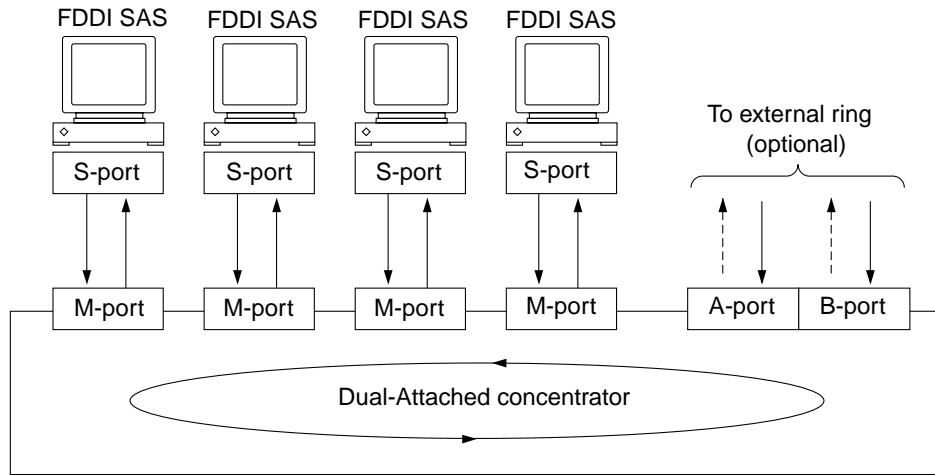


Figure 7-3 Standalone Concentrator

The typical ring architecture of the FDDI network is less obvious in this topology because it exists within the concentrator itself. For this reason, this arrangement of stations is usually described as a *tree*, with the concentrator as the root.

Advantages

In the standalone concentrator configuration, individual stations have less influence on the operation of the network, which is controlled by the concentrator. Concentrators are inherently more stable than FDDI stations. They do not have monitors, or disk drives, are subject to more predictable

usage, and are less likely to be switched off. As a result, a standalone concentrator provides a more reliable network than the basic dual-ring configuration described on page 70.

Concentrators are equipped with built-in electrical bypass facilities that isolate single-station faults. Unlike the station optical bypass facility described on page 62, there is no limit to the number of stations that can be bypassed using the electrical switches in concentrators.

The majority of stations attached to the concentrator are single-attached stations; therefore, only one cable is required for each station.

More flexibility is allowed in the physical location and wiring configuration. Since the stations do not have to be attached in any fixed order and all cables return to a central concentrator, this configuration is useful at sites where FDDI cable has already been installed.

The A- and B-ports on a concentrator can be used to connect it to an external dual-ring configuration. This is a common configuration called the *ring of trees*, which is discussed on page 77.

Disadvantages

The number of stations that can be attached to a single concentrator is limited by the number of M-ports. This is typically in the 2 to 32 range.

The cost of a concentrator is significantly higher than that of a single-attached or dual-attached station; however, there are some low cost concentrators available that do not support all SMT management functions.

Although concentrators are more stable than FDDI stations, when the concentrator goes down, the entire network goes down with it.

Concentrators with Dual-Homing

Figure 7-3 shows two dual-attached stations connected to two dual-attached concentrators in a dual-homing configuration. In this case, each dual-attached station is connected to both DACs. This topology is typically used for connecting critical systems such as file and name servers.

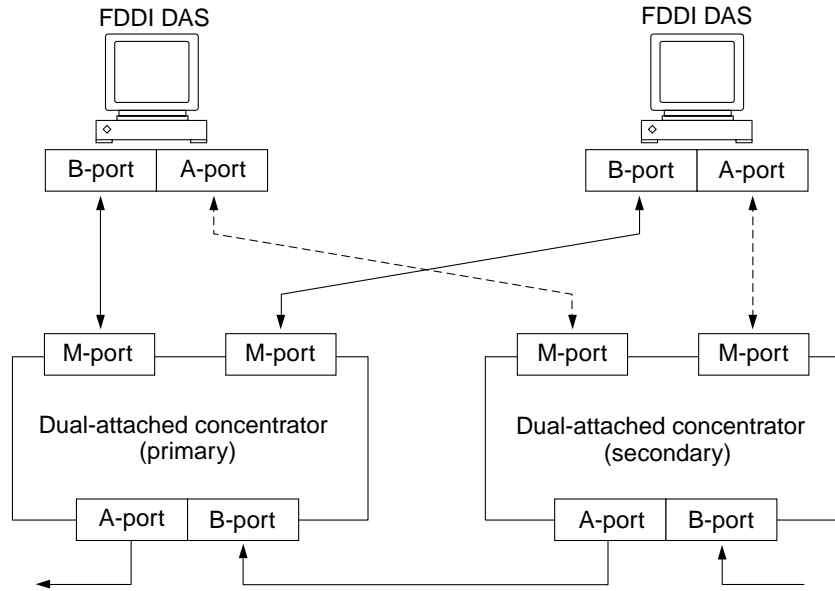


Figure 7-4 Standalone Concentrator With Dual-Homing

Dual-homing provides two independent data paths for each dual-attached station. Under normal conditions, the station communicates on its primary path through the B-port. In the event of a cable or concentrator failure, the station switches to the secondary path connected through the A-port.

Dual-homing is equivalent to the *redundant single-attached station (RSAS)* configuration, which was supported by SunFDDI 2.0. In the RSAS configuration, two single-attached interfaces are used to emulate a dual-attached interface connected in a dual-homing configuration. RSAS is not supported by SunFDDI/P 1.0.

Advantages

Dual-homing offers the same advantages as a standalone concentrator, described on page 72. It also offers improved resistance against cable faults and concentrator failure.

Disadvantages

The number of stations that can be attached to the concentrator is limited by the number of M-ports, which is typically between 2 and 32. Since each dual-homed station requires two M-ports, only a limited number of stations can be connected in this way.

The dual-homing configuration requires a minimum of two concentrators, which are more expensive than single-attached or dual-attached stations.

Tree of Concentrators

The *tree of concentrators* is an incremental expansion of the standalone concentrator topology, described on page 72. Two or more concentrators are connected in a hierarchical topology, with one concentrator designated as the root of the tree, as shown in Figure 7-5. This topology is typically used to connect a large number of stations within a single building or office.

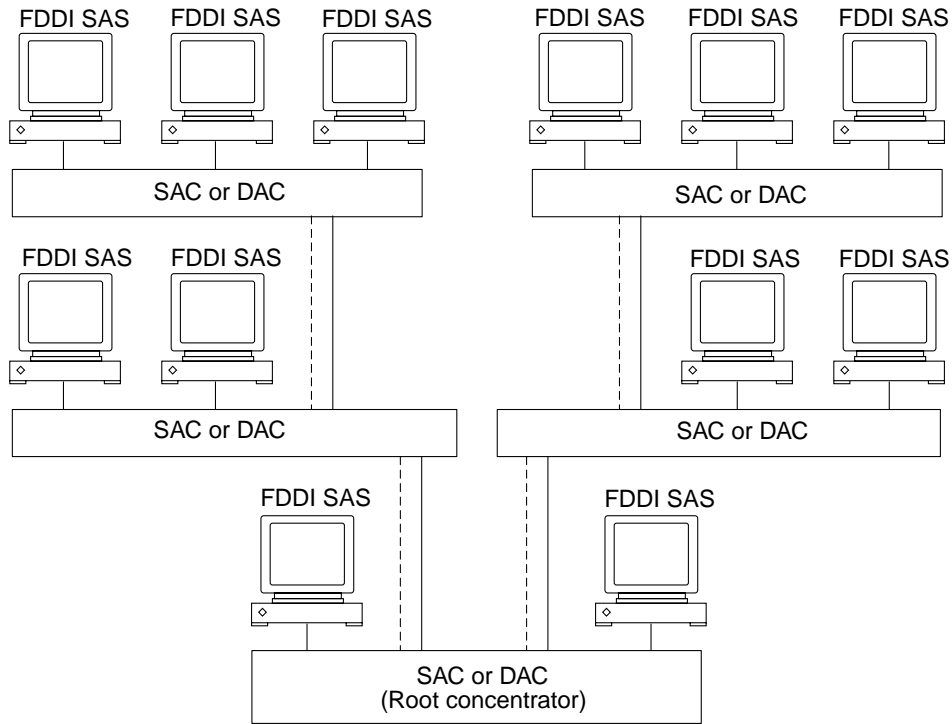


Figure 7-5 Hierarchical Tree of Concentrators

Advantages

The tree of concentrators configuration offers many of the advantages offered by the standalone concentrator, described on page 72; however, it allows a much larger number of stations to be connected.

The cumulative length of the network is increased significantly because the limitation on distance occurs between the end-station and the nearest concentrator, and not between the end-station and the root concentrator.

Disadvantages

This configuration suffers from the same disadvantages as the standalone concentrator, described on page 72. The number of stations connected to each concentrator is limited by the number of M-ports, and the loss of a concentrator takes down all of the stations to which it is attached.

Ring of Trees

The *ring of trees* is a derivative of the tree of concentrators topology described on page 75. The root concentrator is replaced by a dual-ring to which dual-attached concentrators are connected, as shown in Figure 7-6. This topology is ideal for connecting a large number of stations in a campus or multiple-building network.

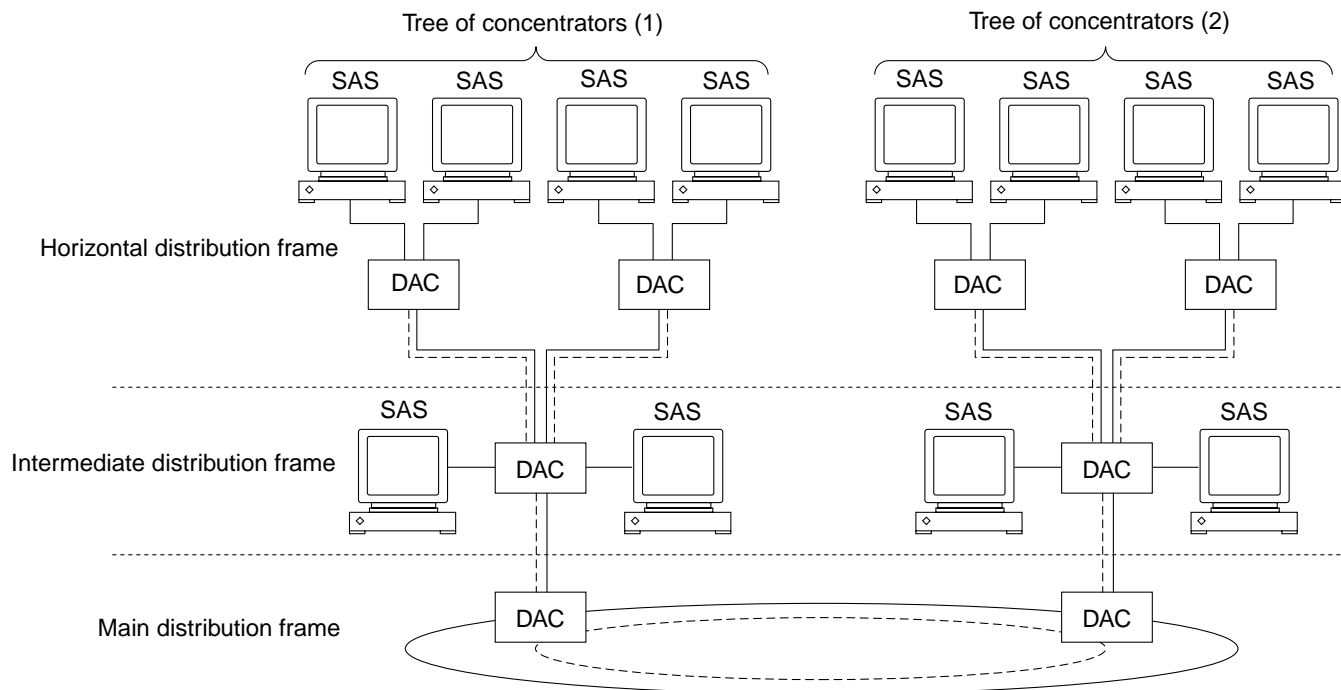


Figure 7-6 Ring of Trees

The dual-ring, and the concentrators attached to it directly, are called the *main distribution frame*. Only one main distribution frame per network is allowed.

The second level in the hierarchy is called the *intermediate distribution frame* and consists of the root concentrators in each building, and any stations connected to them directly.

The third level in the hierarchy is called the *horizontal distribution frame* and consists of concentrators connected to distributed stations, usually through wall boxes.

Advantages

The ring of trees configuration combines many of the advantages of the basic dual-ring and the concentrator-based topologies. It offers good fault tolerance and is a very flexible configuration.

Since only a few stations (the root concentrators) are attached to the dual ring, optical bypass switches can be used to increase the reliability of the main distribution frame, as described in “Optical Bypass Switches” on page 62. The small number of stations and the increased reliability of the concentrators means that the chances of having more than three adjacent stations bypassed at one time are minimal.

The cumulative length of the network can be increased because the limitation on distance occurs between the end-station and the nearest concentrator, and not between the end-station and the root concentrator.

Disadvantages

Only the concentrators attached directly to the dual-ring benefit from its inherent redundancy.

Mixed FDDI/Ethernet Networks

This section concentrates on the implementation of mixed FDDI/Ethernet networks. The same principles could be applied to mixed FDDI/802.3 and mixed FDDI/802.5 (Token Ring) networks. FDDI and Token Ring networks are based on a ring topology; Ethernet and 802.3 networks are based on a bus topology.

FDDI/Ethernet IP Router

The machines at the boundaries between the FDDI and Ethernet networks in the following examples are assumed to be Ultra systems that have at least one SunFDDI/P interface and one Ethernet interface installed. These machines operate as IP routers.

Special software is not required. The combination of the SunFDDI/P driver and the SunOS IP routing implementation handles the routing and relaying of packets between FDDI and Ethernet stations. Note that the packets on the FDDI ring can be as long as 4500 octets, whereas Ethernet packets are restricted to a maximum of 1500 octets.

Figure 7-7 illustrates the basic topology of a mixed FDDI/Ethernet network. In this example, there are two FDDI/Ethernet IP routers that represent the boundaries between the three networks.

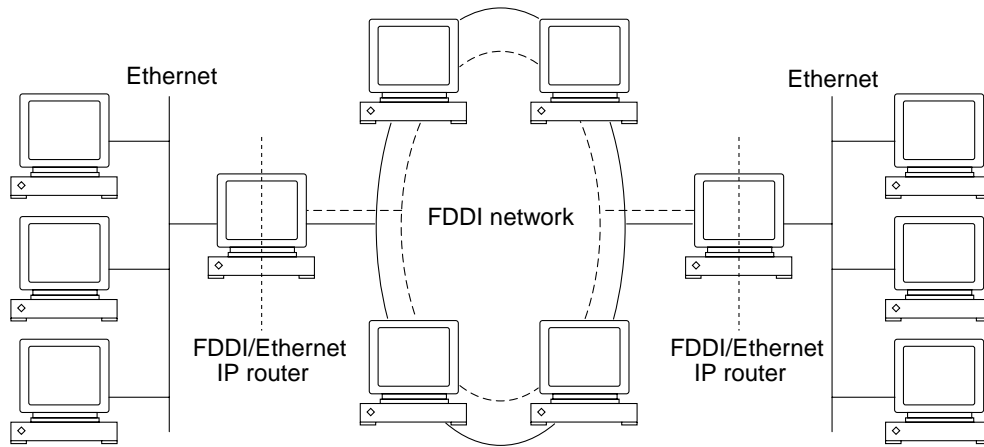


Figure 7-7 Basic Mixed FDDI/Ethernet Topology

A Tree of Concentrators in an Ethernet Network

A tree of concentrators of the type described on page 77 is attached to an Ethernet network by connecting your system to the root concentrator as shown in Figure 7-8

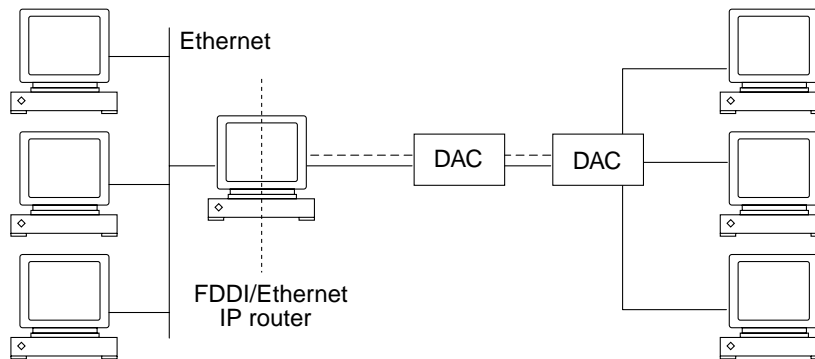


Figure 7-8 Tree of Concentrators Attached to an Ethernet Network

A Ring of Trees in an Ethernet Network

A ring of trees of the type described on page 77 is attached to an Ethernet network by connecting your system to one of the concentrators in the main distribution frame, as shown in Figure 7-9.




Figure 7-9 Ring of Trees Attached to an Ethernet Network

Routes in Mixed FDDI/Ethernet Networks

The FDDI/Ethernet IP router is recognized by the IP address and host name of its FDDI interface by the FDDI network; it is recognized by the IP address and host name of its Ethernet interface by the Ethernet network. Note that an FDDI/Ethernet router usually has the same MAC address on both networks.

In general, stations in an internetwork can be reached via multiple routes. The standard routing daemon (`routed`) included in the SunOS determines the optimum route between stations based on relative hop count, and not on the speed of the interface. Therefore, in cases where the hop count is equal, the use of the faster FDDI network is not guaranteed.

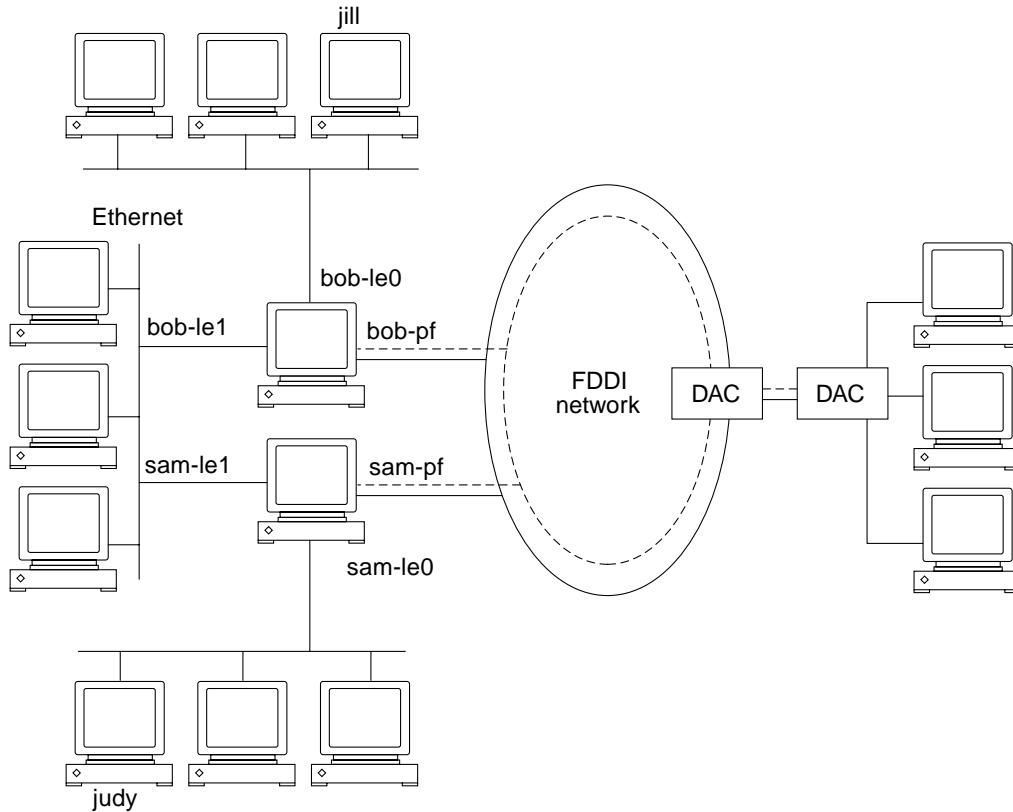


Figure 7-10 Routing in an FDDI/Ethernet Internetwork

For example, in the FDDI/Ethernet internetwork shown in Figure 7-10, there are parallel routes between jill and judy across relatively faster and slower networks:

- **Route 1** (across FDDI link):
jill → bob-le0 → bob-pf → sam-pf → sam-le0 → judy
- **Route 2** (across Ethernet link):
jill → bob-le0 → bob-le1 → sam-le1 → sam-le0 → judy

An identical number of hops between the two stations exists regardless of the route taken; therefore, routed may send all of the traffic across the slower Ethernet link in preference to the faster FDDI link.

To avoid this situation, set up a static route that forces the traffic between jill and judy to pass over the FDDI link; however, static routing is not recommended in a network beyond a few nodes in size. For efficient IP routing in large FDDI/Ethernet networks, use a routing package that provides routing metrics *in addition* to hop count (for example, `gated`, public domain software available by the Center for Theory and Simulation in Science at Cornell University).

To improve the end-to-end system performance across the boundary between the networks, you may need to alter the default MTU for the FDDI network (see “Performance Across Bridges” on page 25) or to disable the MTU path discovery feature (see “Improving Performance” on page 27).

Using the SunFDDI Network Utilities



<i>Changing the Default MAC Address (pf_macid)</i>	<i>page 86</i>
<i>Displaying SunFDDI/P Statistics (pf_stat)</i>	<i>page 87</i>
<i>Monitoring SMT Frames (pf_smtmon)</i>	<i>page 98</i>

This chapter describes the network utilities of SunFDDI/P.

Throughout this chapter, it is assumed that you have installed the SunFDDI/P software under the default base directory *<basedir>* for your operating system:

The default base directory *<basedir>* is:

```
/opt/SUNWconn/bin
```

Changing the Default MAC Address (`pf_macid`)

Each attachment to an FDDI network is identified by a unique 48-bit MAC address. By default, the first SunFDDI/P card takes the host-resident MAC address, which is stored in nonvolatile memory (NVRAM) on the motherboard of the machine. Each subsequent SunFDDI/P card adopts the card-resident MAC address stored in its own IDPROM.

In general, this convention is sufficient to ensure that each SunFDDI/P card installed in the machine has a unique MAC address. However, there may be a conflict with other LAN interfaces that also take the host-resident MAC address—for example, an Ethernet (1e) interface, or a SunFDDI/P 2.0 (bf) interface. In this event, change the default MAC address assigned to the first SunFDDI/P card installed in the system.

▼ To Change the Default MAC Address with `pf_macid`

Use the `pf_macid(1M)` utility to recover the card-resident MAC address, and then modify the system files to override the default MAC address:

1. **Log in as root or become superuser.**
2. **Use the `pf_macid(1M)` utility to recover the MAC address from the IDPROM on the SunFDDI/P interface identified by the instance number `<inst>`.**

```
# <basedir>/pf_macid pf<inst>
<mac_address>
```

3. **Modify the start-up file on your machine so that the MAC address is assigned correctly when the system is rebooted.**

- a. **Edit the `/etc/rcS.d/S30rootusr.sh` file to add the following `if` statement immediately after the `ifconfig` command that initializes the interface `pf<inst>`.**

If you are changing the MAC address of more than one interface, add one `if` statement for each interface.

`if` statement to initialize
interface `pf<inst>` ----->

```
ifconfig $1 plumb
  if [ $1 = "pf<inst>" ]; then
    ifconfig pf<inst> ether <mac_address>
  fi
```

On most systems, the `/etc/rcS.d/S30rootusr.sh` file is a hard link to the `/etc/rootusr` file.

4. Reboot your machine to assign the new MAC address to the SunFDDI interface.

When a SunFDDI/P card takes the host-resident MAC address, it can be swapped to another system without affecting the existing network. However, once a station starts sending packets on the network, the Address Resolution Protocol (ARP) updates the ARP tables on other stations to include the MAC address of its interface. The ES-IS protocol performs the same function for SunFDDI OSI running over FDDI. If you swap SunFDDI/P cards that use the card-resident MAC address, you must wait until the ARP entries time-out, or remove the ARP entries from every active station manually before packets can be routed correctly.

Displaying SunFDDI/P Statistics (`pf_stat`)

The `pf_stat(1M)` utility interrogates a specified SunFDDI/P interface and displays the accumulated statistics. This command must be executed as `root` (or `superuser`) and has the general form:

```
# <basedir>/pf_stat [-m] pf<inst> [<interval>] [<count>]
```

`pf<inst>` specifies the SunFDDI/P interface

`<interval>` is the elapsed time (in seconds) between interrogations

`<count>` the total number of interrogations

The `pf_stat` utility displays information using column headings that conform to SMT revision 7.3, which differ from SMT revision 5.1 and 4.2 headings in the following cases:

- The ECM heading corresponds to the 5.1 MIM heading.
- The RMT heading does not have an analog in SMT revision 4.2. If you run SunFDDI/P at revision level 4.2, ignore any data displayed under the RMT heading of `pf_stat`.

Displaying Local Interface Statistics

When you enter the `pf_stat` command *without* the `-m` option, it displays statistics recovered from the local interface `pf<inst>`.

For example, to display the accumulated statistics for the interface `pf0`, type:

```
# <basedir>/pf_stat pf0
Ring   ECM      RMT      PCMS   Ring_OP   XmitP     RecvP
UP     IN    RING_OP  ACTIVE      c     16fde     1862d
```

You can also monitor the interface dynamically (active monitor), by specifying the *interval* (the elapsed time between interrogations) and *count* (the total number of interrogations). This displays the incremental difference between the current state and the previous state. The minimum interval is one second and the accumulated statistics are displayed after every tenth interrogation.

For example, to monitor the interface `pf0` once every 60 seconds for 3 minutes (a total of 3 interrogations), type:

```
# <basedir>/pf_stat pf0 60 3
Ring   ECM      RMT      PCMS   Ring_OP   XmitP     RecvP
UP     IN    RING_OP  ACTIVE      c     131a0     131aa
UP     IN    RING_OP  ACTIVE      0      0          0
UP     IN    RING_OP  ACTIVE      0      1          1
```

Interpreting Local Statistics

Running the `pf_stat` utility *without* the `-m` option displays information about the various SMT state machines and the network to which the local station is attached:

Ring (Ring Status)

The `Ring` status shows the current state of the physical connection to the FDDI network. The following states may be returned by `pf_stat` under the `Ring` heading

- UP—SunFDDI/P interface physically connected to the active network
- DOWN—SunFDDI/P interface disconnected, or connected to the inactive network

ECM (Entity Coordination Management)

ECM shows the current state of the Entity Coordination Management state machine, which controls the following features and facilities:

- Media availability
- Trace
- Path Test
- Optical Bypass (optional)
- Hold Policy (optional)

Table 8-1 lists the states that may be returned by `pf_stat` under the ECM heading.

Table 8-1 `pf_stat` States Under the ECM Heading

State	Meaning
OUT	ECM is inactive and is waiting for a connect request (initial state).
IN	ECM is active; normal state after successful connection request.
TRACE	ECM is propagating a trace request to the appropriate entity.
LEAVE	ECM is closing all connections prior to the station leaving the ring.
PATH_TEST	ECM is entering a path test state following trace completion.
INSERT	ECM is sending a request to the optical bypass switch to indicate that the station is entering the ring. This disengages the switch.
CHECK	ECM is verifying that symbols are being received from the network.
DEINSERT	ECM is sending a request to the optical bypass switch, to indicate that the station is leaving the ring. This engages the switch.

RMT (Ring Management)

RMT shows the current state of the Ring Management state machine, which controls the following features and facilities:

- MAC availability
- Detection and resolution of duplicate addresses
- Identification of stuck beacon and initiation of trace

Table 8-2 lists the states that may be returned by `pf_stat` under the RMT heading.

Table 8-2 `pf_stat` States Under the RMT Heading

State	Meaning
ISOLATED	RMT is inactive (initial state).
NON_OP	RMT is waiting for an operational ring.
RING_OP	RMT is operating normally.
DETECT	RMT is checking for duplicate addresses (transient state during initialization).

Table 8-2 pf_stat States Under the RMT Heading (Continued)

NON_OP_DUP	RMT has detected that its address is duplicated and is initiating recovery. The ring is not operational in this state.
RING_OP_DUP	RMT has detected that the MAC address is duplicated and flagged the error. The ring is operational in this state.
DIRECTED	RMT has been beaconing for an extended period of time and is transmitting a stream of directed beacons prior to initiating recovery.
TRACE	RMT has initiated a trace to recover a stuck beacon.

PCM (Physical Connection Management)

PCM shows the current state of the Physical Connection Management state machine that controls the following features and facilities:

- Connection initialization
- Maintenance support

This heading is modified to indicate the type of port that is being managed:

- PCMS: single-attached station, S-port
- PCMA: dual-attached station, A-port
- PCMB: dual-attached station, B-port

Table 8-3 lists the states that may be returned by pf_stat under the PCM heading.

Table 8-3 pf_stat States Under the PCM Heading

State	Meaning
OFF	PCM is inactive (initial state).
BREAK	PCM is starting the connection synchronization phase.
CONNECT	PCM is synchronizing the connection end-points prior to the signaling sequence.
NEXT	PCM is transmitting PDUs prior to entering SIGNAL state.
SIGNAL	PCM is transmitting and receiving signal bits (information) following a NEXT state.
JOIN	First state in the sequence leading to a synchronized connection.

Table 8-3 pf_stat States Under the PCM Heading (Continued)

VERIFY	Second state in the sequence leading to a synchronized connection.
ACTIVE	Final state indicating that the port is successfully incorporated in the token path.
TRACE	PCM is localizing a stuck beacon condition.

The normal sequence of PCM states leading to a fully synchronized connection and incorporation of the port into the token path is shown in Figure 8-1. Note that the minimum interval between interrogations is one second and that this is not always fast enough to recover and display the complete sequence of PCM states.




Figure 8-1 Normal Sequence of PCM States

Ring_OP (Ring Operational)

Ring_OP shows the number of Ring_OP (Ring Operational) signals received. This signal is generated when the station is incorporated into an operational network.

XmitP (Transmit Packets)

Running pf_stat without an interval and count, displays the total number of packets transmitted since the interface was activated. Running pf_stat with an interval and count, displays the number of packets transmitted since the last interrogation.

RecvP (Receive Packets)

Running `pf_stat` *without* an interval and count displays the total number of packets received since the interface was activated. Running `pf_stat` *with* an interval and count displays the number of packets received since the last interrogation.

Example Local Statistics

The following output was recovered from a single-attached station using the command shown. A temporary fault condition was simulated by disconnecting the FDDI cable from the SunFDDI/P card and then reconnecting it.

```

Start pf_stat -----> # <basedir>/pf_stat pf0 1 20
Accumulated statistics -> Ring   ECM      RMT      PCMS   Ring_OP   XmitP   RecvP
                        UP     IN   RING_OP  ACTIVE   2         26      1d
                        UP     IN   RING_OP  ACTIVE   0         0       0
                        UP     IN   RING_OP  ACTIVE   0         0       0
                        UP     IN   RING_OP  ACTIVE   0         0       0
Path disconnected -----> DOWN   IN   ISOLATED  CONNECT  0         1       1
                        DOWN   IN   ISOLATED  CONNECT  0         0       0
                        DOWN   IN   ISOLATED  NEXT     0         0       0
Path re-established ----> UP     IN   RING_OP  ACTIVE   1         0       0
(Ring_OP received)      UP     IN   RING_OP  ACTIVE   0         1       1

Accumulated statistics -> Ring   ECM      RMT      PCMS   Ring_OP   XmitP   RecvP
                        UP     IN   RING_OP  ACTIVE   3         29      20
                        UP     IN   RING_OP  ACTIVE   0         0       0
                        UP     IN   RING_OP  ACTIVE   0         1       1
                        UP     IN   RING_OP  ACTIVE   0         0       0
                        UP     IN   RING_OP  ACTIVE   0         0       0
                        UP     IN   RING_OP  ACTIVE   0         0       0
                        UP     IN   RING_OP  ACTIVE   0         0       0
                        UP     IN   RING_OP  ACTIVE   0         0       0
                        UP     IN   RING_OP  ACTIVE   0         0       0
                        UP     IN   RING_OP  ACTIVE   0         0       0
Finish pf_stat -----> UP     IN   RING_OP  ACTIVE   0         1       1

```

Note the following observations regarding this example:

- Accumulated statistics are displayed automatically after every tenth interrogation.
- The combination of `Ring=DOWN` and `RMT=ISOLATED` indicate that the station is disconnected from the network.

- The minimum interval of one second is not fast enough to recover and display the complete sequence of PCM states during the path re-establishment phase.
- A Ring_OP signal is received when the path is re-established indicating that the ring is operational.

The link status indicator mounted on the SunFDDI/P card displays the following sequence of events:

Green (connected) —> **Amber** (disconnected) —> **Green** (connected)

Displaying Statistics from Neighboring Stations

When you use the `pf_stat` command *with* the `-m` option, it displays information about the neighboring stations attached to the local interface `pf<inst>` and the frames received from the network.

For example, to display information about the neighboring stations attached to the interface `pf0`, type:

```
# <basedir>/pf_stat -m pf0
PhyS   Frame  Error  Lost      SA          UNA          DNA
      M  b43eb2    0     3    <mac_addr1> <mac_addr2> <mac_addr3>
```

You can also monitor the neighboring stations dynamically (active monitor), by specifying the *interval* (the elapsed time in seconds between interrogations) and *count* (the total number of interrogations). The minimum interval is one second and the accumulated statistics are displayed after every tenth interrogation.

For example, to monitor the stations attached to `pf0` once every 10 seconds for 1 minute (a total of 6 interrogations), type:

```
# <basedir>/pf_stat -m pf0 10 6
PhyS   Frame  Error  Lost      SA          UNA          DNA
      M  c460a6d    0     3    <mac_addr1> <mac_addr2> <mac_addr3>
      M   27224    0     0    <mac_addr1> <mac_addr2> <mac_addr3>
      M   27227    0     0    <mac_addr1> <mac_addr2> <mac_addr3>
      M   27220    0     0    <mac_addr1> <mac_addr2> <mac_addr3>
      M   2722e    0     0    <mac_addr1> <mac_addr2> <mac_addr3>
      M   27223    0     0    <mac_addr1> <mac_addr2> <mac_addr3>
```


Interpreting Statistics from Neighboring Stations

Running the `pf_stat` utility *with* the `-m` option displays information about the neighboring stations attached to the local interface `pf<inst>`.

Phy (Physical Connection)

PHY shows the type of physical connection to the FDDI network. This heading is modified to indicate the type of port being managed:

- `PhyS`: single-attached station, port S
- `PhyA`: dual-attached station, port A
- `PhyB`: dual-attached station, port B

The following states may be returned by `pf_stat` under the `Phy` heading:

State	Meaning
NONE	Port disconnected
M	Port connected to Port M on a concentrator
S	Port connected to Port S on a single-attached station
A	Port connected to Port A on a dual-attached station
B	Port connected to Port B on a dual-attached station

Frame (Frames Received)

Running `pf_stat` *without* an interval and count displays the total number of SMT frames received since the interface was activated. Running `pf_stat` *with* an interval and count displays the number of SMT frames received since the last interrogation.

More detailed information about the SMT frames can be recovered using the `pf_smtmon(1M)` utility described in “Monitoring SMT Frames (`pf_smtmon`)” on page 98.

Error (Error Frames)

Running `pf_stat` *without* an interval and count displays the total number of error frames received since the interface was activated. Running `pf_stat` *with* an interval and count displays the number of error frames received since the last interrogation. An error frame is defined as an SMT frame whose `E` (error) bit is set, and whose `E` bit is first detected by the local station. It does

not indicate the location of the cause of the error. Frequent `error` frames can indicate a noise problem on the network, either dirt (optical fiber) or electrical interference (UTP).

Lost (Lost Frames)

Running `pf_stat` *without* an interval and count displays the total number of `lost` frames since the interface was activated. Running `pf_stat` *with* an interval and count displays the number of `lost` frames since the last interrogation. A `lost` frame is defined as an SMT frame whose reception is aborted by the `local` station. It does not indicate the location of the cause of the error. A large number of `lost` frames can indicate a noise problem on the network, either dirt (optical fiber) or electrical interference (UTP).

SA (Station Address)

Displays the `MAC` address for the local station.

UNA (Upstream Neighbor Address)

Displays the `MAC` address for the neighboring station, connected upstream on the ring from the local station.

DNA (Downstream Neighbor Address)

Displays the `MAC` address for the neighboring station, connected downstream on the ring from the local station.

Example Neighbor Statistics

The following output was recovered from a single-attached station using the command shown. A temporary fault condition was simulated by disconnecting the FDDI cable from the SunFDDI/P card and then reconnecting it.

```

Start pf_stat -----> # <basedir>/pf_stat -m pf0 1 20
Accumulated statistics -> PhyS      Frame  Error  Lost    SA          UNA          DNA
                          M  c45d5463  1     1b  <mac_addr1> <mac_addr2> <mac_addr3>
                          M    27437    0     0    <mac_addr1> <mac_addr2> <mac_addr3>
                          M    27427    0     0    <mac_addr1> <mac_addr2> <mac_addr3>
                          M    27435    0     0    <mac_addr1> <mac_addr2> <mac_addr3>
Path disconnected -----> NONE     182f1    0     0    <mac_addr1> <mac_addr2> <mac_addr3>
                          NONE      0        0     0    <mac_addr1> <mac_addr2> <mac_addr3>
                          NONE      0        0     0    <mac_addr1> <mac_addr2> <mac_addr3>
Path re-established -----> M        d432    0     7    <mac_addr1> <mac_addr2> <mac_addr3>
(Lost frames indicated)  M    2707e    0     0    <mac_addr1> <mac_addr2> <mac_addr3>

Accumulated statistics -> PhyS      Frame  Error  Lost    SA          UNA          DNA
                          M  c46e5ce7  1    22  <mac_addr1> <mac_addr2> <mac_addr3>
                          M    27228    0     0    <mac_addr1> <mac_addr2> <mac_addr3>
                          M    27230    0     0    <mac_addr1> <mac_addr2> <mac_addr3>
                          M    27227    0     0    <mac_addr1> <mac_addr2> <mac_addr3>
                          M    2722e    0     0    <mac_addr1> <mac_addr2> <mac_addr3>
                          M    2722c    0     0    <mac_addr1> <mac_addr2> <mac_addr3>
                          M    27228    0     0    <mac_addr1> <mac_addr2> <mac_addr3>
                          M    27231    0     0    <mac_addr1> <mac_addr2> <mac_addr3>
                          M    2722b    0     0    <mac_addr1> <mac_addr2> <mac_addr3>
Finish pf_stat -----> M    27227    0     0    <mac_addr1> <mac_addr2> <mac_addr3>

```

Note the following observations regarding this example:

- Accumulated statistics are displayed automatically after every tenth interrogation.
- The combination of PhyS=NONE and the loss of frame activity indicates that the station is disconnected from the network.

Monitoring SMT Frames (pf_smtmon)

The `pf_smtmon(1M)` utility is an active monitor that displays the SMT frames received by the local station. It is particularly useful for diagnosing communication problems with the SunNet Manager proxy agent.

This command must be executed as `root` (or `superuser`) and has the general form:

```
# <basedir>/pf_smtmon [-i pf<inst>] [-x] [-h] [<frameclass>]
```

`-i pf<inst>` specifies the SunFDDI/P interface

`-x` displays the received SMT frames in hexadecimal

`-h` displays help information, including a list of valid frame classes

`<frameclass>` specifies one or more SMT frame classes (used to filter output)

If you do not specify an interface, `pf_smtmon` returns the SMT frames received by `pf0`. If you do not specify a frame type, `pf_smtmon` displays all the SMT frames that it receives. Use `Ctrl-C` to stop `pf_smtmon`.

To display the encoded SMT frames received by interface `pf1`, type:

```
# <basedir>/pf_smtmon -i pf1
pf1: nif_request v=0x1 t=0xfc03e781 s=10-0-4-48-6f-a5 i=0x28
pf1: nif_response v=0x1 t=0xfc03e781 s=10-0-4-8-24-5c i=0x28
pf1: nif_request v=0x1 t=0xfc00dec6 s=10-0-4-b8-6e-ab i=0x28
pf1: nif_request v=0x1 t=0xfc03e787 s=10-0-4-48-6f-a5 i=0x28
pf1: nif_response v=0x1 t=0xfc03e787 s=10-0-4-8-24-5c i=0x28
```

The elements of the SMT frames are defined as follows:

<code>class_type</code>	Identifies the SMT frame class and type (see page 100)
<code>v</code>	Version ID; identifies the structure of the SMT information field
<code>t</code>	Transaction ID; used to pair SMT response and request frames
<code>s</code>	Station ID; uniquely identifies the station transmitting the frame
<code>i</code>	InfoField Length; defines the length of the SMT information field

◆ To display the SMT frames received by interface `pf1` in hexadecimal format, type:

```
# <basedir>/pf_smtmon -i pf1 -x
pf1: nif_request v=0x1 t=0x170 s=10-0-4-8-24-5c i=0x28
004DC000 0000004F FFFFFFFF FFFF1000 0408245C 01020001 00000170 00001000
0408245C 00000028 00010008 00001000 04B86EAB 00020004 00010100 00030004
00002100 200B0008 00000001 00000001 76C467A0

pf1: nif_request v=0x1 t=0x5e0f s=10-0-d4-78-42-4d i=0x28
004D0000 0000004F FFFFFFFF FFFF1000 D478424D 01020001 00005E0F 00001000
D478424D 00000028 00010008 00001000 0408245C 00020004 01010208 00030004
00001200 200B0008 0000000B 00000002 A522BBA1

pf1: nif_response v=0x1 t=0xfc00d94a s=10-0-4-8-24-5c i=0x28
004D0000 00000041 100004B8 6EAB1000 0408245C 01030001 FC00D94A 00001000
0408245C 00000028 00010008 00001000 04B86EAB 00020004 00010100 00030004
00002100 200B0008 00000001 00000001 865549E2 0049C020 F0154E4F FFFFFFFF
FFFF1000 04B86EAB 01020001 FC00D94A 00001000 04B86EAB 00000028 00010008
00001000 D478424D 00020004 00010100 00030004 00002000 200B0008 00000001
00000001

pf1: nif_request v=0x1 t=0x5e13 s=10-0-d4-78-42-4d i=0x28
004D0000 0000004F FFFFFFFF FFFF1000 D478424D 01020001 00005E13 00001000
D478424D 00000028 00010008 00001000 0408245C 00020004 01010208 00030004
00001200 200B0008 0000000B 00000002 4AD75A79

pf1: nif_request v=0x1 t=0x5e17 s=10-0-d4-78-42-4d i=0x28
004D0000 0000004F FFFFFFFF FFFF1000 D478424D 01020001 00005E17 00001000
D478424D 00000028 00010008 00001000 0408245C 00020004 01010208 00030004
00001200 200B0008 0000000B 00000002 DCEBADA2

pf1: nif_request v=0x1 t=0x171 s=10-0-4-8-24-5c i=0x28
004DC000 0000004F FFFFFFFF FFFF1000 0408245C 01020001 00000171 00001000
0408245C 00000028 00010008 00001000 04B86EAB 00020004 00010100 00030004
00002100 200B0008 00000001 00000001 127B1D3B

pf1: nif_request v=0x1 t=0x5e1b s=10-0-d4-78-42-4d i=0x28
004D0000 0000004F FFFFFFFF FFFF1000 D478424D 01020001 00005E1B 00001000
D478424D 00000028 00010008 00001000 0408245C 00020004 01010208 00030004
00001200 200B0008 0000000B 00000002 626FA878
```

SMT Frame Classes and Types

SMT frames are used for peer-to-peer (station-to-station) management. They are divided into *classes*, which define the function of the frame. Each class is then divided into up to three *types*, which define whether the frame is an announcement (information only), a request for service, or a response to a request. Refer to the *ANSI/FDDI Station Management (SMT) X3.299 R7.3 Specification* for a detailed description of SMT frames and their functions.

The `pf_smtmon` utility is used to monitor the following SMT frame classes:

NIF (Neighbor Information Frames)

These are the most common frames displayed when you run `pf_smtmon`. As the name suggests, they carry information about a neighboring station (for example, address, description, state, MAC capabilities) and are used as keep-alive notifications that a station is still attached to the ring and functioning. An NIF frame can be an announcement, a request, or a response.

SIF (Status Information Frames)

These frames carry more detailed information about a station. SIF configuration frames describe the station configuration (for example, number of ports, number of MAC entities, connection policy); SIF operation frames describe the current status of the station. A SIF frame can be either a request or a response.

ECF (Echo Frames)

These frames are equivalent to ICMP `ping` packets and are used to test connectivity between stations. An ECF frame can be either a request or a response.

RDF (Request Denied Frame)

These frames are used to indicate that the request is rejected. If an SMT agent (such as the SunNet Manager proxy agent delivered with SunFDDI/P) receives an unsupported or unrecognized request, it issues an RDF frame to indicate that the request is rejected. An RDF frame is always a response.

ESF (Extended Service Frame)

These frames are implementation dependent. An ESF frame can be an announcement, a request, or a response.

PMF (Parameter Management Frame)

These frames are used to access remote station attributes. The Parameter Management Protocol supports both `get` (display) and `set` (modify) functions; however, the `pf_smtmon` utility can display only `PMF_get` frames. A `PMF_get` frame can be either a request or a response.

Filtering Output from `pf_smtmon`

By default, `pf_smtmon` displays all of the SMT frames received by the local station. You can filter the output generated by `pf_smtmon` by specifying one or more frame classes on the command-line: `nif`, `sif_config`, `sif_operat`, `ecf`, `rdf`, `esf`, `pmf_get`.

For example:

- ◆ **To display only the RDF frames received by interface `pf0`, type:**

```
# <basedir>/pf_smtmon rdf
```


- ◆ **To display the SIF configuration and SIF operation frames received by interface `pf1`, type:**

```
# <basedir>/pf_smtmon -i pf1 sif_config sif_operat
```

- ◆ **To display the NIF frames, SIF configuration frames, and SIF operation frames received by interface `pf1`, type:**

```
# <basedir>/pf_smtmon -i pf1 nif sif_config sif_operat
```


Managing FDDI Stations Using SunNet Manager

9 

<i>Installing the SunNet Manager Agents</i>	<i>page 104</i>
<i>Managing FDDI Networks with SunNet Manager</i>	<i>page 109</i>

This chapter describes how to configure and start the SunNet Manager agents for SunFDDI/P on your machine. It also tells you how to use a SunNet Manager console to recover FDDI statistics from the agents.

To use the SunNet Manager agents for SunFDDI/P you must run SunNet Manager 2.2 or later, and you must install a SunNet Manager console (or management application) on your network. For instructions on how to install and configure SunNet Manager on the management station, refer to the *SunNet Manager Installation Guide*.

All Sun-4d, Sun-4m and Sun-4u systems (and most other machines that connect to TCP/IP networks) support protocols that enable a SunNet Manager console to recover basic information and network statistics. In addition, SunNet Manager provides a number of agents that are used to recover more detailed information from the network:

- SunNet Manager SNMP (simple network management protocol) agents are used to gather information and to set parameters on remote machines.
- SunNet Manager proxy agents are installed locally and used to recover information as if they were located on a remote machine.

Installing the SunNet Manager Agents

You must install a SunNet Manager console (or management application) on your network before installing and configuring the SunNet Manager agents for SunFDDI/P.

Two SunNet Manager agents are provided for SunFDDI/P.

- The **SunFDDI/P local agent** (`pf_fddi`) responds to SNM requests from the console and returns statistics gathered on the client machine. These statistics are similar to those returned by `pf_stat(1M)`. See “Displaying SunFDDI/P Statistics (`pf_stat`)” on page 87 for detailed information.
- The **SunFDDI/P proxy agent** (`pf_fddismt7_2`) responds to SNM requests from the console and issues SMT Request Frames (SRF) to a specified target station. It then returns the information received from the target station in response to the request. This information is similar to that returned by `pf_smtmon(1M)`. See “Monitoring SMT Frames (`pf_smtmon`)” on page 98 for detailed information.

▼ To Install the Standard SNM Agents on a Client

Although you can recover basic information from your FDDI station without installing the standard agents supplied with SunNet Manager, you can manage your network more completely if you install these agents on each machine. You must install the standard agents *before* installing the agents for SunFDDI/P.

There are two ways to install the SunNet Manager agents on the client machine:

- ◆ **Install the package that contains the SunNet Manager agents and libraries on the client machine.**
This method ensures that the system information is updated correctly and makes it easier to remove the agents and libraries, if required.
- ◆ **Mount the SunNet Manager directory tree from the console machine and run `getagents(8)` to install the agents in a local directory.**
This is a quick way to distribute the standard agents across a large network.

Both methods are described in the *SunNet Manager Installation Guide*.

▼ To Install the FDDI Schemas on a Client

Use `pf_install_agents` to install the SunNet Manager agents for SunFDDI/P. This script copies the FDDI schema files to the directory in which the standard agents are installed and updates the configuration files for SunNet Manager. Run `pf_install_agents` on each SunNet Manager console and each SunFDDI/P station you want to manage using SunNet Manager.

1. **Log in as root or become superuser.**

2. **Start `pf_install_agents` by typing:**

```
# /opt/SUNWconn/bin/pf_install_agents
```

3. **Enter the host names of console machine(s) that you want to receive SRF trap reports from the local station.**

You can send SRF trap reports to more than one console. Each console must have the schemas for SunFDDI/P installed.

Host name of console ----->
Host name of console ----->
Press Return to terminate ---->

```
Enter hostnames (if any) to send SRF Trap reports to:
(enter blank to terminate)
      enter hostname:opus
      enter hostname:rigel
      enter hostname:
```

4. **Specify the destination directory for the SunFDDI/P schemas.**

SunNet Manager expects to find the schemas in the directory `<basedir>/SUNWconn/snm/agents`. The variable `<basedir>` is set to `/opt` by default. If you installed the SunNet Manager agents and libraries on the FDDI station, or if you used `getagents` to recover the agents from the console, then this directory exists.

a. Enter the base directory under which you installed the SunNet Manager agents as the destination directory for the SunFDDI/P schemas:

```
What is the name of the root directory under which the SunNet
Manager is installed? [/opt]: [?]
```

5. Specify the current location of the SunFDDI/P schemas.

This is the base directory under which you installed the SunFDDI/P software package (SUNWpf) using `pkgadd(1M)`.

```
What is the name of the root directory under which the pf FDDI
software is installed? [/opt]: [?]
```

The SunFDDI/P schemas are copied to the destination directory and the local configuration files are modified for SunNet Manager.

6. Start the SunNet Manager agents for SunFDDI/P by typing:

```
# /sbin/pf_snmd
```

Note – If you typed `n` in response to the question, “Do you want to start the SunNet Manager daemons for SunFDDI at boot time?”, when you installed the SunFDDI/P software, you must start the SunNet manager agents for SunFDDI/P manually each time you reboot your machine. Alternatively, you can install the start-up script into the `/etc/rc2.d` directory:

```
# cd /opt/SUNWconn/SUNWpf/utilities
# cp pf_fddidaemon /etc/rc2.d/S98pf_fddidaemon
```

Installing the FDDI Schemas on the Console

If you install the SunFDDI/P software on the console machine, you can run `pf_install_agents` locally. See “To Install the FDDI Schemas on a Client” on page 105 for detailed instructions.

If you do not have a SunFDDI/P SBus card installed in the console machine, you must either copy the script and SunFDDI/P schemas onto the console machine or NFS mount the `pf_install_agents` script from a client machine and execute it remotely.

▼ To Share `pf_install_agents` From a Client

1. On the client machine, ensure that the NFS and mount daemons are running:

```
# /usr/lib/nfs/nfs 8
# /usr/lib/nfs/mountd
```

2. On the client machine, share the `/opt` directory under which you installed SunFDDI/P.

```
# share -F nfs -o ro /opt
```

3. On the console machine, NFS mount the `/opt` directory:

```
# mount <client>:/opt /<mnt_dr>
```

4. On the console machine, run `pf_install_agents`:

```
# /mnt/SUNWconn/bin/pf_install_agents
```

Follow the directions in “To Install the FDDI Schemas on a Client” on page 105, and specify `<mnt_dir>` as the current location of the SunFDDI/P schemas.

- The SunFDDI/P schemas are located in the directory `/opt/SUNWconn/SUNWpf/agents` and the `pf_install_agents` script is located in the directory `/opt/SUNWconn/SUNWpf/utilities`.

- On the console machine, the schemas must be installed in the directory /opt/SUNWconn/snm/agents. The pf_install_agents script modifies the file /opt/SUNWconn/snm/struct/elements.schema to add the definition for the host name or MAC address of an FDDI station. If you do not run pf_install_agents on the console machine, you must modify this file to add the following line to each component (workstation, server, or router) definition:

```
string[40]      Name/Mac-Address
```

For example, if you want to specify the host name or MAC address of an FDDI interface installed in an Ultra system, add the following line to elements.schema:

```
#
# Sun workstations
#

record component.ultra1 (                                # ultrasparcstation1
    string[64]      Name
    string[40]      IP_Address
    string[40]      User
    string[40]      Location
    string[80]      Description
    string[40]      SNMP_RdCommunity
    string[40]      SNMP_WrCommunity
    string[64]      SNMP_Vendor_Proxy
    int             SNMP_Timeout
    string[40]      Name/Mac-Address
```

Add this line ----->

Managing FDDI Networks with SunNet Manager

For detailed information on how to use SunNet Manager to manage the various components of your network, refer to the *SunNet Manager User's Guide*.

The SunNet Manager agents for SunFDDI/P are used to recover FDDI statistics and SMT frame information. For detailed instructions on how to interpret this data, refer to the ANSI/FDDI standards listed in “FDDI Conformance” on page 2.

▼ To Set Up the SunNet Manager Console

1. **Log in as root or become superuser.**
2. **Start the SunNet Manager console, and initialize the runtime management database (MDB).**

```
# /opt/SUNWconn/bin/snm -i
```

3. **Use HeadStart to begin discovering the nodes in your network, or use BasicStart to build your own database with the console's graphical editor.**
Note that the Discover facility will only fully recognize the Ethernet (1e) interfaces in your network.
4. **Modify the component properties of each SunFDDI/P station in your network.**
Use the scroll bar to show all of the properties associated with the station. You must enter either the host name or the MAC address associated with the SunFDDI/P interface in the last field in the list of properties. If you enter a host name, it must be registered in the NIS/NIS+ map or in the `/etc/hosts` file on each machine in your network. If you enter a MAC address, it must be entered in ring format.
5. **Click SELECT on the check box associated with each SunNet Manager agent that you want to communicate with the station.**
If you enable the SunFDDI/P proxy agent (`pf_fddismt7_2`), you must also specify the host name (or MAC address) of the target FDDI station from which the proxy agent will request information.

Using the SunFDDI/P Local Agent (pf_fddi)

The SunFDDI/P local agent (pf_fddi) responds to SNMP requests from the console and returns statistics gathered on the client machine. These statistics are similar to those returned by pf_stat(1M). See “Displaying SunFDDI/P Statistics (pf_stat)” on page 87 for detailed information.

▼ To Gather Local Station Statistics

Use the SunFDDI/P local agent to recover a snapshot of the current state of the local FDDI station:

1. Click **SELECT** on the icon for the FDDI station that you want to interrogate.
2. Press **MENU** on the Requests menu button and release **MENU** on the Quick Dump item.
3. Choose pf_fddi as the Agent Schema and if (interface) as the Group.
4. Click **SELECT** on Apply to display local station statistics similar to those shown in Figure 9-1.

```

SunNet Manager Console: Quick Dump Report - arbois
Mon Apr 18 21:51:35 1994 [ arbois ] : Quick Dump: fddi.if

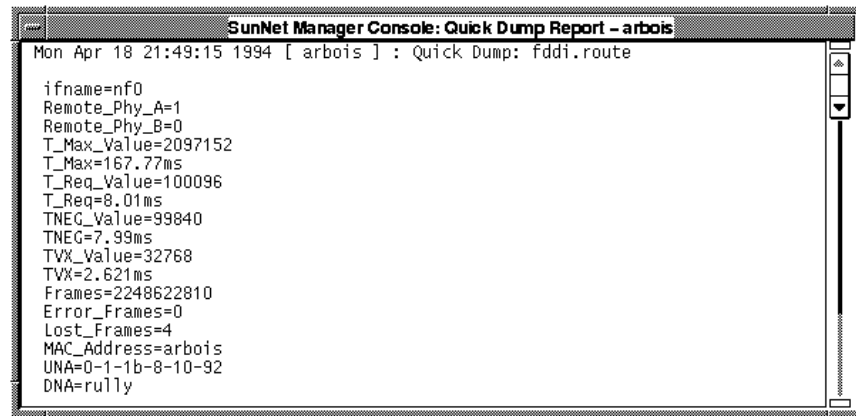
ifname=nf0
Ring_State=1
Ring=Up
ECM_State=1
ECM=In
ECM_Status=Connected
OBP_State=0
OBP=None
RMT_State=2
RMT=Ring_Op
CFM_State=12
CFM=Thru
PCM_A_State=8
PCM_A=Active
PCM_B_State=8
PCM_B=Active
Recovery=35
XmitP=3267652
XmitB=2595625496
RecvP=3057228
RecvB=946139565
RFull=1
RErr=1
XAbort=0
  
```

Figure 9-1 Quick Dump (Local Statistics)

▼ To Gather Neighbor Station Statistics

Use the **SunFDDI/P local agent** to recover a snapshot of the current state of the neighboring FDDI stations:

1. Click **SELECT** on the icon for the FDDI station that you want to interrogate.
2. Press **MENU** on the Requests menu button and release **MENU** on the **Quick Dump** item.
3. Choose `pf_fddi` as the Agent Schema and `route` as the Group.
4. Click **SELECT** on **Apply** to display neighbor station statistics similar to those shown in Figure 9-2.



```

SunNet Manager Console: Quick Dump Report - arbois
Mon Apr 18 21:49:15 1994 [ arbois ] : Quick Dump: fddi.route

ifname=nf0
Remote_Phys_A=1
Remote_Phys_B=0
T_Max_Value=2097152
T_Max=167.77ms
T_Req_Value=100096
T_Req=8.01ms
TNEG_Value=99840
TNEG=7.99ms
TVX_Value=32768
TVX=2.621ms
Frames=2248622810
Error_Frames=0
Lost_Frames=4
MAC_Address=arbois
UNA=0-1-1b-8-10-92
DNA=rully

```

Figure 9-2 Quick Dump (Neighbor Statistics)

▼ To Request Information from the Local Agent

- ◆ Use the console to create a request for data that is sent periodically to the SunFDDI/P local agent.

This is similar to running the `pf_stat(1M)` utility and specifying an `<interval>` and `<count>` as described in “Displaying SunFDDI/P Statistics (`pf_stat`)” on page 87.

The local agent responds with the requested information, which can be saved to file or displayed using the SunNet Manager Grapher.

▼ To Use the SunFDDI/P Proxy Agent (pf_fddismt7_2)

The **SunFDDI/P proxy agent** (pf_fddismt7_2) responds to SNM requests from the console and issues SMT Request Frames (SRF) to a specified target station. It then returns the information received from the target station in response to the request. The basic operation of the SunFDDI/P proxy agent is shown in Figure 9-3.




Figure 9-3 SunFDDI/P Proxy Agent

▼ To Specify the Target Station

1. Specify the target machine when you enable the SunFDDI/P proxy agent in the properties window for the proxy station.
2. Use the scroll bar to display the two SunFDDI/P agents as shown in Figure 9-4 and enter the host name of the target station on the input line.


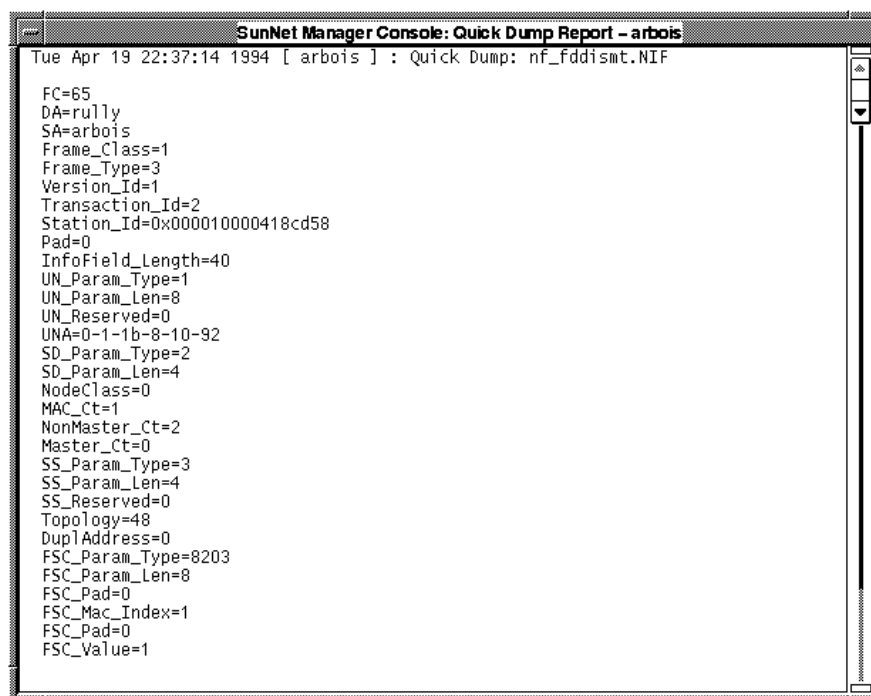


Figure 9-4 Specifying the Target Station

▼ To Gather Information from the Proxy Agent

Use the SunFDDI/P proxy agent to recover SMT information from one of the neighboring FDDI stations:

1. Click **SELECT** on the icon for the FDDI proxy station that you want to interrogate.
2. Press **MENU** on the Requests menu button and release **MENU** on the **Quick Dump** item.
3. Choose `pf_fddismt7_2` as the Agent Schema and one of the SMT frame types or FDDI MIBs as the Group.
4. Click **SELECT** on **Apply** to display SMT information similar to that shown in Figure 9-5.



```

SunNet Manager Console: Quick Dump Report - arbois
Tue Apr 19 22:37:14 1994 [ arbois ] : Quick Dump: nf_fddismt.NIF
FC=65
DA=rully
SA=arbois
Frame_Class=1
Frame_Type=3
Version_Id=1
Transaction_Id=2
Station_Id=0x000010000418cd58
Pad=0
InfoField_Length=40
UN_Param_Type=1
UN_Param_Len=8
UN_Reserved=0
UNA=0-1-1b-8-10-92
SD_Param_Type=2
SD_Param_Len=4
NodeClass=0
MAC_Ct=1
NonMaster_Ct=2
Master_Ct=0
SS_Param_Type=3
SS_Param_Len=4
SS_Reserved=0
Topology=48
DuplAddress=0
FSC_Param_Type=8203
FSC_Param_Len=8
FSC_Pad=0
FSC_Mac_Index=1
FSC_Pad=0
FSC_Value=1

```

Figure 9-5 Quick Dump (SMT Frames)

Interpreting Traps and Errors

One of the most important functions of the SunNet Manager agent for SunFDDI/P is to signal the occurrence of faults or errors on the network. Traps are sent to the consoles that you specified when you ran the `pf_install_agents` script to install the agents on the client machine. See “To Install the FDDI Schemas on a Client” on page 105 for detailed instructions.

When a cable or station fault occurs, it is signaled by the neighboring station and not by the station implicated by the fault; however, the information provided in the trap report sent to the console allows you to identify the probable source of the problem. Faults often generate a series of trap events that are caused when the ring reconfigures itself automatically.

The reason for repeated errors and lost frames is more difficult to identify because the fault is signaled by the first system to receive the error frame, which is not necessarily the station closest to the problem.

The trap report in Figure 9-6 was generated by the upstream neighbor after the cable between a single-attached station and a concentrator was disconnected. All MAC addresses are displayed in *ring* format.

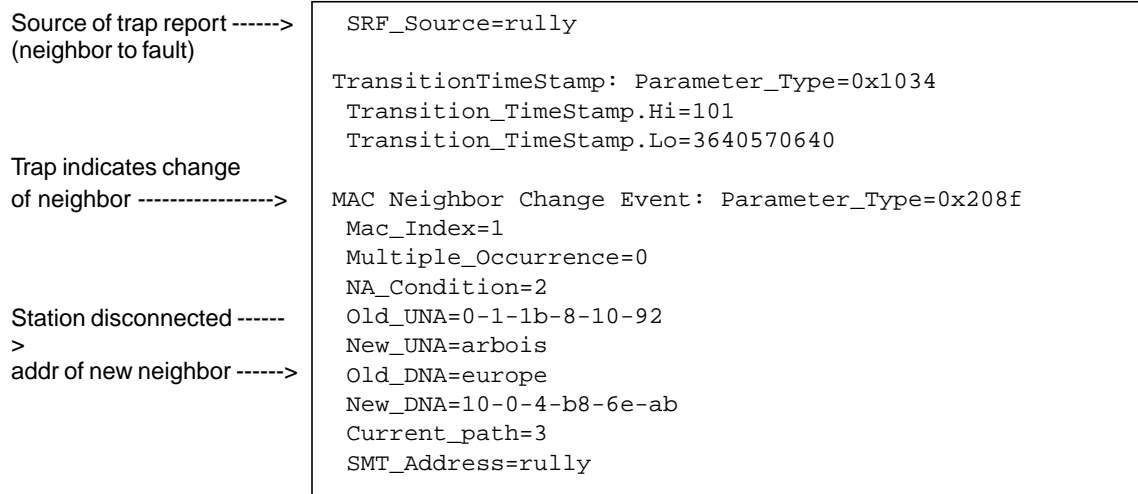


Figure 9-6 Single-Attached Station Disconnected

The trap reports in Figure 9-7 were generated when a single-attached station was reconnected to the ring.

Trap report generated by the upstream station

Source of trap report -----> (neighbor to fault)	SRF_Source=rully
Trap indicates change of neighbor ----->	TransitionTimeStamp: Parameter_Type=0x1034 Transition_TimeStamp.Hi=102 Transition_TimeStamp.Lo=845593304
Current neighbor -----> Old neighbor returns ----->	MAC Neighbor Change Event: Parameter_Type=0x208f Mac_Index=1 Multiple_Occurrence=0 NA_Condition=2 Old_UNA=0-1-1b-8-10-92 New_UNA=arbois Old_DNA=10-0-4-b8-6e-ab New_DNA=europe Current_path=3 SMT_Address=rully

Trap report generated by the reconnected station

Source of trap report -----> (reconnected station)	SRF_Source=europe
Trap indicates error condition ----->	TransitionTimeStamp: Parameter_Type=0x1034 Transition_TimeStamp.Hi=1 Transition_TimeStamp.Lo=2482471160
Frames lost during fault condition ----->	MAC Frame Error Condition: Parameter_Type=0x208d Mac_Index=1 Condition_State=0 Frame_Ct=101367280 Error_Ct=0 Lost_Ct=8 Frame_Error_Ratio=0

Figure 9-7 Single-Attached Station Reconnected

The trap report in Figure 9-8 was generated by the upstream neighbor when a dual-attached station connected directly to the FDDI ring was halted. The neighbor wraps itself around the fault automatically.

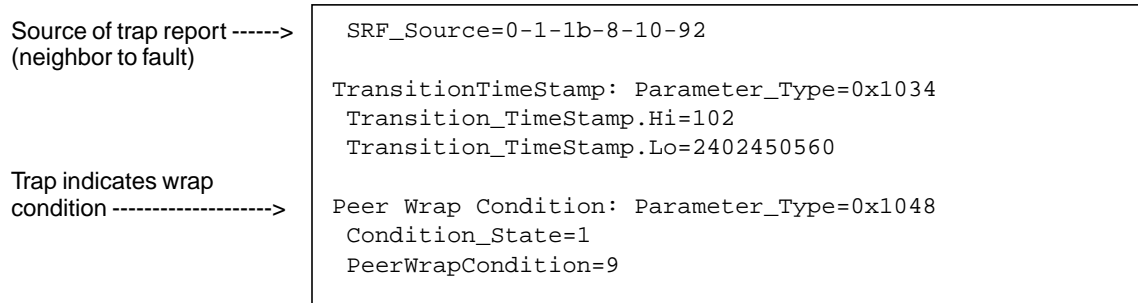


Figure 9-8 Dual-Attached Station Halted

The trap report in Figure 9-9 was generated by the upstream neighbor when a dual-attached station connected directly to the FDDI ring was restarted. It shows a change of neighbor event that occurs as the wrapped station rediscovers the neighbors to which it was connected before the fault.

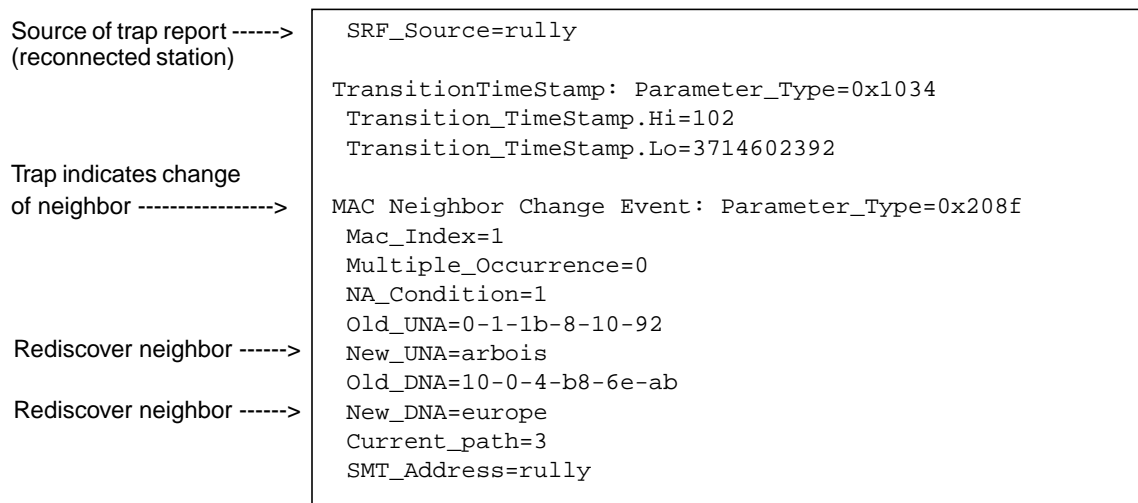


Figure 9-9 Dual-Attached Station Restarted

MAC Address Formats

The MAC address can be represented in one of two formats:

- Canonical format
- Ring format

The MAC address is always displayed in canonical format by the SunFDDI/P network utilities and hardware self-test program. It is also displayed in canonical format when you reboot your machine.

The SMT information displayed by SunNet Manager always shows the MAC address in ring format. If you specify a target station using a MAC address (instead of a host name), you must enter it in ring format.

Developing Applications that Run over SunFDDI/P

10 

Developing Applications with DLPI

page 120

This chapter describes how to develop your own applications running over SunFDDI/P networks. It assumes that you are an experienced programmer who is already familiar with the DLPI Version 2 interface.

This chapter provides a brief overview of each of the programming environments and concentrates on network addressing and other issues that are specific to the SunFDDI/P implementation of these interfaces. It does not provide a detailed specification of the interfaces.

DLPI is used to access the services of the SunFDDI/P link layer of the SunFDDI/P device driver in a Solaris 2.5.1 Hardware: 4/97 environment without the need for specialized knowledge of the underlying protocol. You can obtain a copy of the *STREAMS DLPI Specification* from Sun Microsystems, Inc. (part number 800-6915-10).

Developing Applications with DLPI

The Data Link Provider Interface (DLPI) version 2 enables any data link user to access the services of a data link provider without any special knowledge about the underlying protocols. In addition to FDDI services, DLPI supports access to X.25, LAPB, LAPD, ISDN, Ethernet, Token Ring, Token Bus, SDLC, BISYNC, and other data link protocols.

The DLPI header file is `<sys/dlpi.h>`. A detailed description of the functions contained in the header is available in the *STREAMS DLPI Specification*.

DLPI defines a STREAMS message set, state table, and conventions that support interaction between a data link service user (for example, IP or ARP) and a data link service provider (for example, the pf FDDI driver). Figure 10-1 illustrates the relationship between DLPI and the layers of the FDDI protocol.

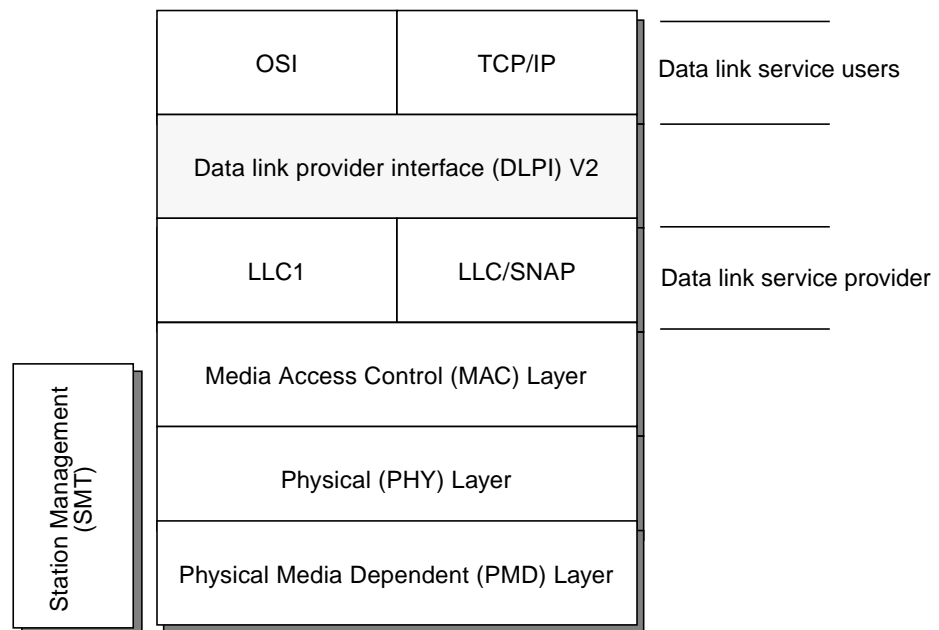


Figure 10-1 DLPI Access to FDDI

DLPI primitives are defined in terms of STREAMS messages. All DLPI messages are either type `M_PROTO` or type `M_PCPROTO` and contain the appropriate DLPI message structure. DLPI primitives are exchanged by *confirmed* negotiation—request primitives are issued by the user, which then waits to receive an acknowledge (or error) response primitive from the provider. This type of negotiation is illustrated in Figure 10-2.

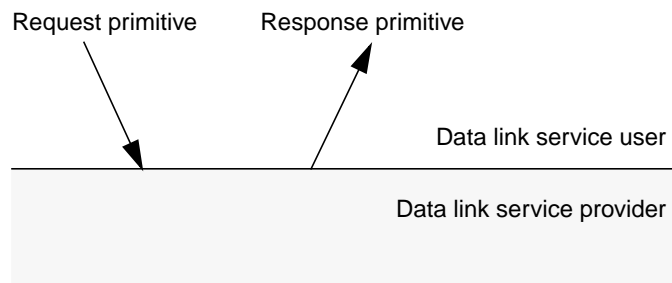


Figure 10-2 DLPI Request/Response Negotiation

DLPI Connectionless Service Primitives

Applications developed to run over FDDI use the DLPI connectionless service primitives, which support the following set of operations:

- Attaching a STREAM to a given addressable device
- Getting and setting the physical address of the attached device
- Binding a data link service access point (SAP) to the STREAM
- Transmitting and receiving data link frames
- Transmitting and receiving LLC XID/TEST messages
- Transmitting and receiving multicast addresses
- Enabling physical and data link promiscuous mode

DLPI Addressing

DLPI addressing is implemented as a DLSAP (data link service access point) address. Each DLSAP consists of a physical address plus a service access point (SAP). The format of the DLSAP is provider-dependent. Figure 10-3 shows the general format of the SunFDDI/P DLSAP.

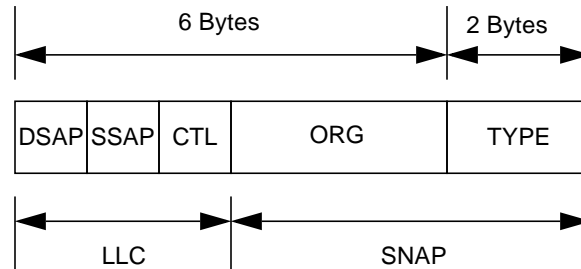


Figure 10-3 SunFDDI/P DLSAP Format

Information returned by the `DL_INFO_ACK` primitive (in response to a `DL_INFO_REQ`) is used to specify the encoding and decoding rules to insert and extract the DLSAP for a given data link provider. The information returned by the SunFDDI/P device driver is shown in Table 10-1.

Table 10-1 `DL_INFO_REQ` Frame

Element	Value	Description
<code>dl_max_sdu</code>	4352 bytes	Maximum MTU
<code>dl_min_sdu</code>	0 bytes	Minimum MTU
<code>dl_addr_length</code>	8 bytes	Length of DLSAP
<code>dl_mac_type</code>	<code>DL_FDDI</code>	FDDI MAC layer
<code>dl_reserved</code>	0	Unused
<code>dl_current_state</code>	—	Current state of data link provider
<code>dl_sap_length</code>	-2 bytes	Length of DLSAP
<code>dl_service_mode</code>	<code>DL_CLDLS</code>	Connectionless service
<code>dl_qos_length</code>	0	Not supported

Table 10-1 DL_INFO_REQ Frame

Element	Value	Description
dl_qos_offset	0	Not supported
dl_range_length	0	Not supported
dl_range_offset	0	Not supported
dl_provider_style	DL_STYLE2	Can be cloned
dl_addr_offset	76	Offset to address location
dl_version	2	DLPI Version 2
dl_brdcst_addr_length	6	Broadcast address length
dl_brdcst_addr_offset	84	Offset to broadcast address location

Note that the value `dl_sap_length` is set to minus 2, which indicates that the SAP is 2 bytes beyond the physical address.

Setting Up Servers and Diskless Clients

11 

<i>Setting Up a Server and Diskless Clients</i>	<i>page 126</i>
<i>Booting a Diskless Client Over SunFDDI/P</i>	<i>page 133</i>

This chapter describes how to set up a server so that you can boot and run diskless clients across an FDDI network. It assumes that the server and clients are attached to the same FDDI ring.

Setting up a server for a diskless client running over FDDI involves these steps:

1. Setting up the server so that it can export an operating system for a diskless client and installing the SunFDDI/P device driver and utilities
2. Introducing the diskless client to the network
3. Installing the SunFDDI/P device driver in the `/export/root/<client>` file system for the diskless client
4. Booting the diskless client over the network

Setting Up a Server and Diskless Clients

When you set up a Solaris 2.5.1 Hardware: 4/97 server so that it can export its operating system, you need to know the number of diskless clients that it will serve. This tells you the amount of disk space you need to reserve in the `/export` file system.

♦ Use `df(1M)` to check the amount of space available, for example, type:

```
% /usr/bin/df -k /export
Filesystem      kbytes  used  avail  capacity  Mounted on
/dev/dsk/c0t3d0s0 24143  17222  4511   79%       /
```

The figure in the `avail` column tells you how much space is available in the file system.

- For each Solaris 2.5.1 Hardware: 4/97 client, you will require:
 - 20 Mbytes of disk space under `/export/root`
 - 24 Mbytes (by default) of disk space under `/export/swap`

In addition, you require 10 Mbytes under `/export/exec` for the Solaris 2.5.1 Hardware: 4/97 template files (Sun-4u) plus an additional 15 Mbytes under `/export/exec`.

For example,

If you have five Sun-4u diskless clients running Solaris 2.5.1 Hardware: 4/97, you will require:

- $(5 \times 20) = 100$ Mbytes of disk space under `/export/root`
- $(5 \times 24) = 120$ Mbytes of disk space under `/export/swap`
- $(10 + 15) = 25$ Mbytes of disk space under `/export/exec`

▼ To Set Up a Server

Follow these steps to set up a server to boot diskless clients across an FDDI connection:

- 1. Install and configure the relevant SunOS operating system on the server so that it can be exported.**
See the *Solaris System Configuration and Installation Guide* for detailed instructions on how to set up an operating system for diskless clients.
- 2. Check that you have sufficient space in the `/export/root`, `/export/swap`, and `/export/exec` file systems.**
- 3. Install and configure the SunFDDI/P driver and utilities on the server.**
- 4. Ensure that you have an entry for the `sysadmin` group (with a GID of 14) in the group table `/etc/group`.**
You need this entry to use `admintool` to administer remote machines. See *Solaris 2.5 Setting up User Accounts, Printers, and Mail*.

Note – If you are running NIS on your server, you can use `admintool` to update the local `/etc/hosts` and `/etc/ethers` files. You must merge these files manually into the NIS maps afterwards.

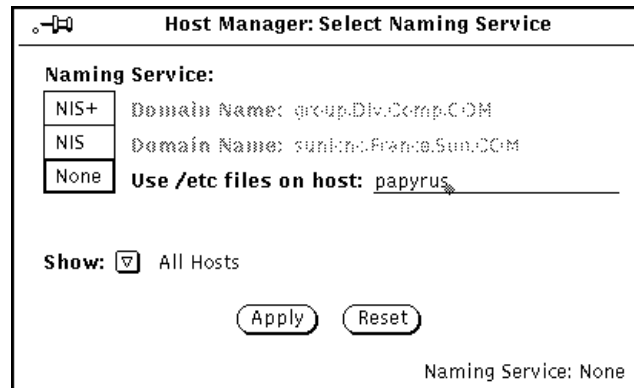
▼ To Introduce Diskless Clients to a Server

Before you can introduce a diskless client to the network you need to know its host name, IP address and Ethernet address. The six-byte Ethernet address is displayed when the machine is switched on. (Note that the Ethernet address of the client is the MAC address stored in NVRAM on its motherboard.)

- 1. Log in as a member of the `sysadmin` group.**
You must log in as a member of the `sysadmin` group to administer remote machines. Do *not* log in as `root` or become `superuser`.
- 2. Start `admintool` by typing:**

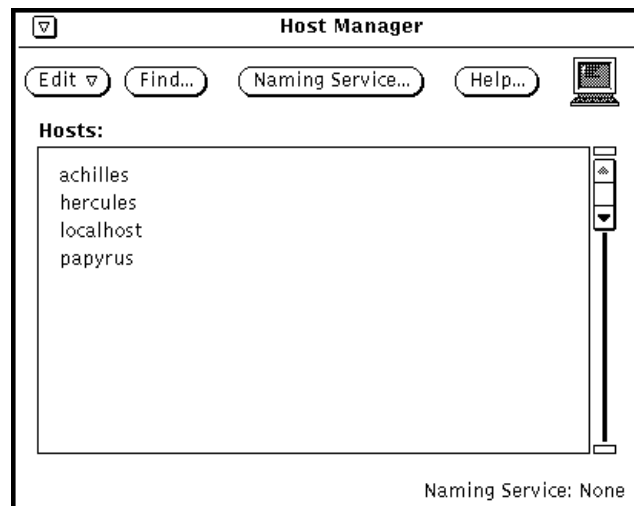
```
% /usr/bin/admintool &
```

3. Click **SELECT** on the **Host Manager** icon to display the **Select Naming Service** window.




4. Click **SELECT** on the button corresponding to the naming service used by the server to display list of available hosts.
- If the server is running NIS+, click **SELECT** on NIS+ and type your domain name.
 - If the server is running NIS or is not using a naming service, click **SELECT** on None and type the name of the host on which the files `/etc/hosts` and `/etc/ethers` are located.

5. Click **SELECT** on **Apply** to display the **Host Manager** window.



6. Press **MENU** on the **Edit** menu button and choose the **Add Host...** item to activate the **Add Host** window.



7. Press MENU on the Client Type menu button, drag and release MENU to display the defaults for diskless clients.

Host Manager: Add Host

Client Type:

Host Name:

IP Address:

Ethernet Address:

Timezone Region:

Timezone:

File Server:

OS Release:

Root Path:

Swap Path:

Swap Size:

Terminal Type:

Naming Service: None

8. Enter the host name, IP address, and Ethernet address of the diskless client.
This information uniquely identifies the diskless client on the network. The six-byte Ethernet (MAC) address has the form X : X : X : X : X : X. It is stored in NVRAM on the motherboard and is displayed when the system is booted.
9. Use MENU to set the Timezone Region and Timezone that are relevant for your location.
10. Press MENU on the File Server menu button to display a list of available servers. Drag and release MENU to specify a file server.
The file server is the machine from which the diskless client will receive its operating system. If you are running admintool on the server, it will be shown as the default.

11. Specify the root and swap paths.

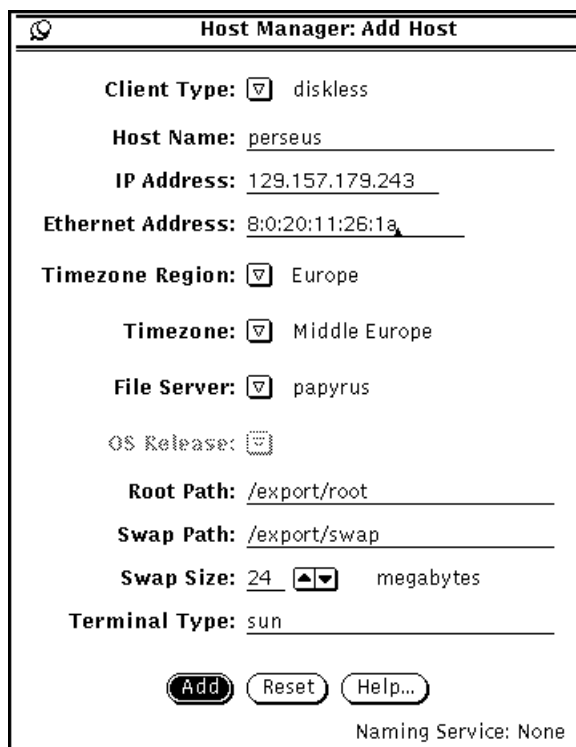
By default, these paths are set to `/export/root` and `/export/swap`. If you change these paths, you must ensure that you have sufficient disk space in the file systems you specify.

12. Specify the size of the swap space reserved for the diskless client.

This is normally set to 24 Mbytes. If you increase the amount of swap space reserved for the diskless client, you must ensure that you have sufficient space under `/export/swap` on the server.

13. Click SELECT on the Add button to update the list of recognized hosts.

This creates the export file systems for the diskless client on the server and updates the relevant files for the naming service you are using. The operation will take a few minutes. The message “operation completed” appears when the list has been updated and the name of the diskless client is added to the scrolling-list on the Host Manager window.



Host Manager: Add Host

Client Type:

Host Name:

IP Address:

Ethernet Address:

Timezone Region:

Timezone:

File Server:

OS Release:

Root Path:

Swap Path:

Swap Size:

Terminal Type:

Naming Service: None

▼ To Update the NIS Server

If you are running NIS, you must update the NIS map on the NIS server manually. If you are running NIS+ or you are not using a naming service, proceed to “Installing the SunFDDI/P Driver for a Diskless Client” on page 133.

1. On the diskless client’s server (on which you just ran `admintool`), note the new entries created in the following files:

- `/etc/hosts`
- `/etc/ethers`
- `/etc/timezone`
- `/etc/bootparams`

2. On the NIS server, update the four files with the information taken from the diskless client’s server.

3. On the NIS server, remake the NIS map by typing:

```
# cd /var/yp
# make
```

4. Change directories to the diskless client’s `root` file system.

```
# cd /export/root/<client>
```

5. Update the `yp` map for the diskless client:

```
# mv var/yp- var/yp
```

6. Replace the `noname` entry in the file `/etc/defaultdomain` with the your local NIS domain name.

Installing the SunFDDI/P Driver for a Diskless Client

After you have created the export file systems under `/export/root/<client>` and `/export/swap/<client>` on the server, you must install the SunFDDI/P device driver for the diskless client.

▼ To Install on a Client

1. Use `pkgadd` to install the Solaris 2.5.1 version of SunFDDI/P so that it is exported to the diskless client.

By default, the SunFDDI/P software is installed under `/opt`.

2. To redirect the installation so that the package is installed in the root partition for the diskless client, type:

```
# pkgadd -R /export/root/<client> -d /cdrom/sun_fddip_1_0
```

Follow the installation instructions contained in the section “Installing and Configuring SunFDDI/P” on page 13.

Booting a Diskless Client Over SunFDDI/P

You need to specify the SunFDDI/P interface when you boot your diskless client over the FDDI network.

▼ To Boot a Diskless Client

1. Check that the server is running and that its FDDI interface is up.
2. Switch on the diskless client and press `L1(Stop)-a` to halt the boot procedure.

3. At the boot prompt, type the `show-nets` command to get a list of possible interfaces.

```
ok show-nets
```

A listing similar to the following example will be displayed:

```
a) /pci@b,2000/pf@2
b) /pci@a,2000/pf@2
c) /pci@b,4000/SUNW,hme@1,1
d) /pci@l,2000/pf@2
e) /pci@0,2000/pf@2
f) /pci@0,4000/SUNW,hme@1,1
g) /sbus@f,0/sbf@),7ff0
h) /sbus@f,0/SUNW,hme@3,8c00000
i) /sbus@e,0/nf@1,7ff0
j) /sbus@3,0/nf@0,7ff0
m) MORE SELECTIONS
q) NO SELECTION
Enter Selection, q to quit
```

4. Type the letter of your selection.

Your selection will not appear on the screen. It will be copied into the buffer.

5. At the prompt, type `boot` then press CTL-y to paste your selection.

Warning – Once a diskless client has been booted over the SunFDDI/P interface, do not move the SunFDDI/P card to a different slot.

Frequently Asked Questions About SunFDDI



This appendix contains the responses to some of the most frequently asked questions about SunFDDI/P. It also contains references to more detailed information.

- **The network performance is much lower than I expected. Why can't I transfer data at 100 Mbps over my FDDI connection?**
If the applications running over FDDI don't use the available bandwidth efficiently you will see much lower performance than expected. See Chapter 4, "Improving Network Performance" for advice on how to get the best performance from your SunFDDI/P station.
- **Can I connect a dual-attached FDDI station to two different subnetworks?**
No. A dual-attached station must be either connected to the primary and secondary rings in a dual-ring configuration, or connected in a dual-homing configuration. See "Dual-Attached Station (DAS)" on page 57 for more information.
- **Can I use a dual-attached station in single-attached mode?**
Yes. Simply connect one port to the network. No overhead is incurred when you operate a dual-attached station in this way.

- **Can I use SunFDDI/P 1.0 in a redundant single-attached station (RSAS) configuration?**

No. This configuration, which was supported by SunFDDI 5.0 SBus, provides a redundant connection that is used in the event of a link failure. The equivalent configuration for SunFDDI/P 1.0 is provided by a dual-attached station connected in dual-homing mode. See “Concentrators with Dual-Homing” on page 74 for more information.

- **Can I use the SunFDDI/P 1.0 software with a hardware interface from a third-party vendor?**

No. The device driver does not support any third-party hardware products.

- **Can I connect two single-attached stations “back-to-back”, without connecting them to a concentrator?**

Yes. This configuration is supported by the FDDI specifications. You can also connect two single-attached stations and one dual-attached station directly in a ring, but this configuration is not supported by the FDDI specifications and is not recommended.

- **What is dual-homing?**

Dual-homing is a special configuration that connects a single dual-attached station to two concentrators so that there is a redundant path, which is used in the event of a link failure. See “Concentrators with Dual-Homing” on page 74 for more information.

- **What is the MTBF for the various SunFDDI/P options?**

The calculated mean time before fail (MTBF) for each of the SunFDDI/P interface cards is 1231176 hours

This number was obtained using the standard method for calculating MTBF.

- **Why can't I connect port A to port A, or port B to port B?**

This type of configuration is forbidden by the FDDI specifications. It is called “twisting” or “crossing” the ring. If you connect your FDDI station in this configuration, the link status indicator will remain blinking green to indicate an invalid connection. See Chapter 2, “Hardware Installation,” for a description of the different diagnostic LEDs and Chapter 7, “FDDI Network Topologies,” for a description of the valid network configurations.

- **Why do I get the error message “duplicate address” or “duplicate macid”?**
By default, the first SunFDDI/P interface installed in the machine adopts the host-resident MAC address stored on the motherboard. Each subsequent SunFDDI/P interface takes the card-resident MAC address stored in its IDPROM.
- **Why is performance poor when I use NFS over FDDI?**
For a single file transfer, the maximum data transfer rate is limited to 0.5 Mbytes/second when writing to an NFS mounted file system and 1.5 Mbytes/second when reading from an NFS mounted file system. This performance limitation, which affects Ethernet, FDDI, and any other fast network, is due to the NFS version 2 protocol. It will be removed by NFS version 3.
- Note that when several NFS operations are performed in parallel (several reads and writes in parallel), the aggregate throughput is much higher with FDDI than it is with Ethernet.

6. How do I disable the SunFDDI/P interface?

Use `ifconfig(1M)` to disable the SunFDDI/P interface temporarily. It will be reconfigured when you reboot your machine.

```
# ifconfig pf<inst> down
```

To disable the SunFDDI/P interface so that it is not reconfigured when you reboot the machine, change the name of the `/etc/hostname.pf<inst>` file.

```
# mv /etc/hostname.pf<inst> /etc/oldhostname.pf<inst>
```

• How do I disable the Ethernet interface?

Use `ifconfig(1M)` to disable the Ethernet interface temporarily. It will be reconfigured when you reboot your machine.

```
# ifconfig le<inst> down
```

To disable the Ethernet interface so that it is not reconfigured when you reboot the machine, change the name of the `/etc/hostname.le<inst>` file.

```
# mv /etc/hostname.le<inst> /etc/oldhostname.le<inst>
```

- **How do I use the host name assigned to the SunFDDI/P interface as the primary host name for my machine?**

In a Solaris 2.x environment, the primary host name is defined in the `/etc/nodename` file. Normally, this file contains the host name assigned to the Ethernet interface built in to your machine. Edit `/etc/nodename` to enter the host name assigned to the SunFDDI/P interface.

- **Can I associate several IP addresses with a single SunFDDI/P interface?**
Yes. To associate multiple IP addresses with a single interface, type:

```
# ifconfig pf0 <ip_address_1> netmask <mask> up
# ifconfig pf0:1 <ip_address_2> netmask <mask> up
# ifconfig pf0:2 <ip_address_3> netmask <mask> up
.
.
.
```

- **How many SunFDDI/P 1.0 cards can I install in one machine?**
You can install up to eight SunFDDI/P 1.0 cards in one machine.
- **What kind of connector can be used to connect SunFDDI/P 1.0 to optical fiber networks?**
The connector type has been changed for SunFDDI 5.0 to an SC-type connector. See Chapter 2, “Hardware Installation” for a description of FDDI cables and connectors.
- **What kind of optical fiber can be used to link SunFDDI/P 1.0?**
SunFDDI/P uses a 62.5/125 μm standard twin optical fiber. A compliant link design will have no more than 11 db of end-to-end loss. This includes loss incurred in the fiber (typically 1.2 db/km), the connectors (typically 0.5 db/connector), and any bypass switches (typically 2.5 db/switch). It also allows loss margins for source aging (typically 1 to 3 db), safety (typically 1 to 2 db), and temperature variation (typically 1 db).

-
- **Does SunFDDI/P Dual incorporate an optical bypass?**
No. You must order an optical bypass switch from a third-party vendor.
 - **Why does my optical bypass switch not work when I plug it in?**
Optical bypass switches require a power cycle to be initialized. Plug in the optical bypass switch, switch off the FDDI station, and then reboot to initialize the station. See “Optical Bypass Switches” on page 62 for a more detailed description.
 - **How can I increase the distance between two SunFDDI/P stations?**
SunFDDI/P supports multimode fiber connections only. This limits the maximum distance between stations to two kilometers (1.24 miles). To increase the distance between stations, either:
 - Use converters from multi- to single-mode fiber, and a single-mode connection between converters.
 - Add more nodes between the stations to act as repeaters.

Glossary

This glossary defines FDDI-related acronyms, abbreviations, and terms.

ARP

(Address Resolution Protocol). A protocol used within local networks that support broadcast (such as Ethernet or FDDI) that binds IP addresses to MAC addresses. View the current contents of a Sun machine's ARP table with the SunOS command, `arp -a`; specified in RFC 826.

baud

A measure of the rate of data transfer, based on the number of times per second the signal can change states on the transmission line. The transmission line ordinarily uses just two signal states to indicate data transfer, making the baud rate equal the number of bits per second that can be transferred. See also **bps**, **Bps**, **kbps**, and **Mbps**.

bps

(bits per second). A measure of the rate of data transmission. See also, **baud**, **Bps**, **kbps**, and **mbps**.

Bps

(bytes per second). A measure of the rate of data transmission. See also, **baud**, **bps**, **kbps**, and **mbps**.

CFM

(Configuration Management). The portion of CMT that manages the internal configuration of the station's MAC and PHY entities.

CMT

(connection management). The portion of the FDDI SMT that manages the operation of physical layer insertion and removal, and the logical connection of PHY entities to the MAC entities.

DAS

(dual attachment station). An FDDI station that connects to both primary and secondary rings; a Class A station, capable of dynamic reconfiguration to accommodate certain levels of failure. Two rings are used to interconnect the stations so that the failure of any station or link causes the remaining stations to reconfigure the network to maintain connectivity: wrapping.

dotted decimal notation

A way of representing the 32-bit Internet address: four decimal integers separated by decimal points, where each integer is the decimal value of one octet of the Internet address. For example, the 32-bit Internet address (binary) 11000000 00001001 11001001 00000111 is represented as 192.9.201.7 (decimal).

ECM

(Entity Coordination Management). The portion of CMT that manages the media interface to the FDDI network, which includes coordinating the activity of all the PHYs associated with that physical attachment and controlling the optional optical bypass function within the station.

FDDI

(Fiber Distributed Data Interface). A standard specified by the American National Standards Institute for fiber optic links of up to 100 Mbits/s.

gateway

A dedicated computer that connects to two or more networks and routes packets from one network “community” to the other. Gateways route packets to other gateways until the packets can be delivered directly to their specified destination across one physical network. In general usage, the term identifies any computer that transfers information between networks.

kbps

(kilobits per second). A measure of the rate of data transmission. See also, **baud**, **bps**, **Bps**, and **Mbps**.

LLC

(Logical Link Control). Local area network protocol functions corresponding to the OSI data link layer (layer 2), but providing only service access points (SAP).

MAC

(Media Access Control). Local area network protocol functions corresponding to the OSI data link layer (layer 2), but providing the following: 1) on transmission, assembles data into a frame with address and error-detection fields; 2) on reception, disassembles frame, performs address recognition and error detection; 3) links communications management.

Mbps

(millions of bits per second). A measure of the rate of data transmission. See also, **baud**, **bps**, **Bps**, and **kbps**.

NRZI-4B/5B

(nonreturn to zero inverted). A differential encoding technique that encodes four-bit data groups as five bits (data plus a signal element); satisfies the minimum signal transition requirements to transmit synchronization (clock) within the data stream.

OSI

(Open System Interconnection). The OSI model is a standard for computer-communications protocols and the implementation of these protocols. The model is a product of the *International Standards Organization* (ISO), and specifies a seven-layer architecture with one or more protocol standards at each layer.

PCM

(Physical Connection Management). The portion of CMT that manages a physical connection between the PHY being managed and another PHY, likely in an adjacent station on the FDDI network.

PHY

(Physical Layer Protocol). Performs the clock recovery and serial-to-parallel (on receive) or parallel-to-serial (on transmit) of data between the transmission medium and the MAC entity.

PMD

(physical medium dependent). Specifies the optical-to-electrical conversion mechanism to conform to FDDI.

RARP

(Reverse Address Resolution Protocol). A protocol used by diskless machines and their servers within local networks such as Ethernet or FDDI. In a RARP exchange, a diskless machine broadcasts its MAC address. A boot server responds to that client with the client's IP address, enabling the client to begin the boot process.

RMT

(Ring Management). The portion of CMT that monitors the MAC. It provides a trace mechanism to handle the stuck beaconing condition, and also detects duplicate addresses on the ring.

RSAS

(redundant single-attach station). An FDDI station in which single-attach interfaces are paired, with one active and the other a hot standby. In an RSAS with connections to two different dual-attach concentrators, you enjoy the network availability afforded by concentrators with single-attach connections, plus, for the RSAS itself, the station availability of a dual-attach station.

SAP

(service access point). A logical interface between two adjacent protocol layers (such as LLC and IP).

SAS

(single attachment station). An FDDI station that attaches to a single FDDI ring, as distinguished from a dual-attached station, which attaches to both rings in a dual ring.

SMT

(Station Management). The complete standard of PHY layer and frame-based protocols to be supported by conforming FDDI stations.

THT

(Token-Holding Timer). A timer in the FDDI MAC protocol that controls the length of time a station can hold a token for asynchronous transmission.

timing jitter

The deviation of clock recovery that can occur when the receiver attempts to recover both clocking and data from the received signal. The PHY specification describes a phase lock loop (PLL) and elasticity buffer to eliminate jitter.

TRT

(Token Rotation Timer). A timer in the FDDI MAC protocol that measures the time between receipt of tokens. It expires when it reaches a value that exceeds the Target Token Rotation Time (TTRT).

TTRT

(Target Token Rotation Time). A parameter in the FDDI MAC protocol that is negotiated by all of the stations on a ring at the time a ring is brought up. FDDI uses the TTRT in conjunction with the TRT and THT to allocate bandwidth to individual stations on the ring. The FDDI MAC protocol guarantees that the worst-case delay a token will experience while circulating the ring will never exceed twice the negotiated TTRT.

Index

Symbols

//www.sun.com, xxiv
/cdrom/sun_fddip_1_0 directory, 13
/etc/hosts file, 15, 22
/etc/rc2.d directory, 21
/etc/system file, 30

Numerics

802.3 networks, 79
802.5 networks, 79

A

access delay, 24, 26
add_drv command, 39
aggregate attenuation, 63
ANSI/FDDI standards, 50
A-ports, 57
asynchronous
 service, 65
 traffic, 65
attenuation, 51, 62

B

backup, 70
bandwidth, 1, 24, 65, 66, 69

base directory
 installing SUNWnfr, 14
 installing SUNWnft, 16
basic dual ring, 70
battery, lithium, 3
bidding, 65
bit stream, 53
boot test, 44
bridges, 25, 42
buffers, 25
bulk data transfers, 24
bus topology, 79
bypass switch, 62

C

cable
 failure, 55, 61
 keys, 34
 specification, 2, 51
card-resident MAC address, 35
checking the installation, 19
checklist, troubleshooting, 34
claim process, 26, 31, 65
classes, 100
cleaning
 fibre optic cable, 7

commands
 add_drv, 39
 getsockopt, 29
 ifconfig, 18, 36
 mount_hfs, 13
 ndd, 27, 28
 netstat, 18
 ping, 18
 pkgadd, 13, 14, 16
 pkgrm, 22
 setsockopt, 29
 umount, 18
common problems, 33, 39
concentrators, 58
configurations
 dual-ring, 60
 ring of trees, 59
connectionless service primitives, 121
connectors, 51, 57
 fiber optic, 7, 34
console, 104, 106
copper twisted-pair, 50
counter-rotating ring, 55, 70
CPU speed, 24
crossing the ring, 34
cumulative length, 76

D
DAC, dual-attached concentrator, 56, 58, 59, 69
daemons
 in.routed, 37
 vold, 18
DAS, dual-attached station, 56, 57, 61, 69
data link
 provider interface, 120
data queues, 27
data transfer, 50
 rate, 2, 25
declaring IP addresses, 12, 15
default
 buffer size, 25
 IP addresses, 15
 MAC address, 35
 target token rotation time, 26
developing applications, 119
device
 drivers, 34
 names, 4
diagnostic LED, 9, 16, 34
dialogue-based traffic, 52
directories
 /cdrom/sun_fddip_1_0, 13
 /etc/rc2.d, 21
disabling MTU path discovery, 31
disk space, 12
diskless clients, 125
displaying SunFDDI/P statistics, 87
distributed clocking scheme, 51
distribution frame, 78
 DL_INFO_ACK, 122
 DL_INFO_REQ, 122
 dl_sap_length, 123
DLPI, 119
 addressing, 122
 header file, 120
 primitives, 121
 specification, xxiii
DLSAP, 122
DNA, 96
document structure, xix
don't fragment bit, 30
dotted decimal notation, 142
downstream
 neighbor address, 96
dual ring, 70
dual-attached
 concentrator, 72
 concentrator (DAC), 56, 58, 59
 station, 51, 74
 station (DAS), 1, 34, 56, 57, 61
dual-homing configuration, 74
dual-ring configuration, 60
dynamic routing protocol, 42

E

- ECF, 100
- echo frames, 100
- ECM, 89
- efficiency, 24
- efficient use of bandwidth, 67
- ejecting the CD-ROM, 17
- electrical-bypass, 73
- electromagnetic radiation, 51
- elements.schema, 108
- encoded digital data, 51
- entity coordination management, 89
- error frames, 95
- ESF, 101
- examples, 29
- extended service frame, 101

F

- failure mechanisms, 60
- fault recovery, 72
- faults
 - cable, 61
 - station, 62
- FDDI conformance, 2
- FDDI performance, 66
- fiber optic
 - components, 5
- fibres optic cable
 - cleaning, 7
- Filename | Command > pf_
 smtmonutility, 98
- files
 - /etc/hosts, 15, 22
 - /etc/system, 30
- finishing the installation, 17
- frame classes, 98
- frames, 52, 65

G

- gated daemon, 83

- gateway, 142
- getagents script, 21, 104
- getsockopt command, 29
- guaranteed traffic, 52

H

- handling
 - fiber optic cable, 5
 - Sbus cards, 5
- hardware
 - options, 1
 - self-test, 44
- hierarchical topology, 75
- high priority, 66
- high water marks, 27
- hop count, 81, 83
- horizontal distribution frame, 78
- host name, 12, 15, 108
- host-resident MAC address, 35
- hubs, 58

I

- IDPROM, 35
- ifconfig command, 18, 36
- improving
 - ring latency, 24
- in.routed daemon, 37
- increased transmission time, 66
- independent data paths, 74
- inherent redundancy, 78
- installation
 - checking the, 19
 - requirements, 12
 - scripts, 12, 16, 17
- installing
 - multiple cards, 4
 - SNM agents, 20 to 21
 - the hardware, 5 to ??
 - the software, 13 to 18
- instance
 - names, *See* device names

- numbers, 4
- interference, 51
- intermediate distribution frame, 78
- IP
 - addresses, 12, 15, 36
 - routing, 79

J

- jumpers, 5

K

- kbps, 142
- kernel space, 25
- keying, 57
- keys, 34

L

- LAN interfaces, 12
- Late_Ct parameter, 66
- latency, 24
- LED, diagnostic, 9, 16, 34
- link status indicator, 9, 16, 34, 40
- list of packages, 14
- lithium battery, 3
- LLC, logical link control, 143
- local agent, 104, 110
- lost frames, 96
- low water mark, 27

M

- MAC, 81
 - address, 35
- MAC, medium access control, 49, 50, 52
- main distribution frame, 78
- major numbers, 39
- master SBus slot, 4
- maximizing data throughput, 25
- maximum
 - bandwidth, 24

- length of ring, 24
- number of bypass switches, 63
- number of cards, 2, 4
- number of stations, 24
- throughput, 24
- transfer unit (MTU), 25, 83

- Mbps, 143

- media dependent, 49, 50, 51
- medium access control (MAC), 2, 49, 50, 52

- memory, 12

- microprocessor speed, 24
- minimum requirements, 12
- mixed FDDI/Ethernet networks, 79
- MMF, multimode fiber, 51
- monitoring SMT frames, 98
- mount_hsf command, 13
- moving PCI cards, 4

- MTU

- negotiation, 30
 - path discovery, 30, 31
- multicast addresses, 121
- multimode fiber, 51
- multiplexer, 58

N

- ndd command, 27, 28
- negotiated MTU, 30
- neighbor information frames, 100
- netstat command, 18, 36

- network

- architecture, 55
 - efficiency, 24
 - recovery, 62
 - reliability, 60, 69
 - size, 69
 - topologies, 59
 - traffic, 69

- NFS, 106

- NIF, 100

- NIS/NIS+, 109

- NIS/NIS+ servers, 12

noise, causes of, 7
NRZI-4B/5B, 143
number of cards supported, 2, 4
NVRAM—TOD chip, 3

O

one-way ring, 55, 60, 71
optical
 bypass switch, 62, 63, 73
 fiber, 50
 power budget, 62, 63
optimum
 performance, 24
 route, 81
OSI, 143
overall system performance, 66, 83

P

packages
 SUNWnf, 39
packet
 fragmentation, 26
 size, 25, 30, 79
parallel routes, 82
parameter management frame, 101
parameters
 SO_RCVBUF, 29
 SO_SNDBUF, 29
 T_opr, 26
 T_Req, 31
 T_req, 26, 31
 tcp_recv_hiwat, 28
 tcp_xmit_hiwat, 28
 tcp_xmit_lowat, 29
 udp_recv_hiwat, 28
 udp_xmit_hiwat, 28
 udp_xmit_lowat, 29
patch revisions, xxiv
patches, xxiv
path discovery, MTU, 30, 31, 42
PCM, 91
peer-to-peer management, 100

permission token, 24, 65
pf_fddi agent, 104, 110
pf_fddismt7_2 agent, 104, 109, 112
pf_install_agents, 106
pf_install_agents script, 20
pf_macid utility, 86
pf_smtmon command, 38
pf_smtmonutility, 101
pf_stat utility, 110
pf_stat utility, 87
PHY, physical layer, 49, 50, 51, 143
physical
 address, 122
 connection, 95
 connection management, 91
 delay, 24
 layer, 49, 50, 51
 media dependent, 50, 51
 media dependent layer, 49
 medium, 51
physical layer (PHY), 2
physical media dependent (PMD), 2
Pin, 56, 57
ping command, 18
pkgadd command, 13, 14, 16
pkgrm command, 22
placement of stations, 69
PMD, 51
PMD, physical media dependent, 49, 50,
 143
PMF, 101
PMF_get, 101
port A, 136
port B, 136
potential performance, 24, 66
Pout, 56, 57
power budget, 62
power-up self test, 44
primary
 components, 44
 input, 56, 57
 output, 56, 57

- path, 74
- ring, 55, 57, 60
- priority-based traffic, 52
- product list, 1
- programming interface, 119
- proxy agent, 43, 103, 104, 109, 112

Q

- queues, 27
- quick dump, 43

R

- R7.2.99, 2, 49
- RARP, reverse address resolution protocol, 144
- RDF, 100
- RDF, request denied frame, 38, 43
- rebooting the machine, 18
- receive buffers, 25
- receive packets, 93
- receivers, 51
- recovery, 55, 72
- recovery mechanisms, 60, 62
- RecvP, 93
- reduced transmission time, 66
- relative hop count, 81
- reliability, 60, 69
- relocating PCI cards, 4
- removing the software, 22
- repeater, 51
- request denied frame, 100
- request denied frame (RDF), 38, 43
- required bandwidth, 69
- restricted
 - token exchange, 65
 - traffic, 65
- RFC 1390, 25
- right to transmit, 65
- ring, 52
 - architecture, 72

- failure, 70
- format, 114
- length, 24
- ring management, 90
- ring of trees, 73, 77, 81
- ring of trees configuration, 59
- ring utilization, 24
- Ring_OP, 92, 94
- RMT, 90
- root
 - concentrator, 72, 77
 - of the tree, 75
- rotation time, 65
- routed, 81
- routing metrics, 83
- routing tables, 37
- RSAS, redundant single-attached station, 144

S

- SAC, single-attached concentrator, 56, 58, 69
- safety notice, 3, 5
- SAP, 122
- SAS, single-attached station, 56, 69
- schemas, 20, 108
- scripts
 - examples, 29
 - getagents, 21
 - installation, 12, 16, 17
 - pf_install_agents, 20
 - start-up script for SNM agents, 21
 - to modify water marks, 29
- secondary
 - input, 57
 - output, 57
 - ring, 55, 57, 60
- send buffers, 25
- service access point, 122
- setsockopt command, 29
- SIF, 100
- signal layer, 51

- simple network management
 - protocol, 103
- sin, 57
- single-attached
 - concentrator (SAC), 56, 58
 - station (SAS), 1, 34, 56
- single-mode fiber, 51
- SMF, single-mode fiber, 51
- SMT request frame (SRF), 20, 38, 112
- SMT, station management, 49, 60
- SMT, station mangement, 50
- SNMP, simple network management
 - protocol, 103
- SO_RCVBUF parameter, 29
- SO_SNDBUF parameter, 29
- Sockets, 25, 29
- solving common problems, 39
- Sout, 57
- SRF, SMT request frame, 20, 38, 112
- stability, 70
- stable network, 70
- stand-alone
 - concentrator, 72
- start-up script for SNM agents, 21
- static electricity, 5
- static routes, 42
- station
 - failure, 55, 60
 - faults, 62
 - initialization, 55
 - management, 49, 50
 - upstream, 56
 - wrapping, 55, 62
- station management (SMT), 2, 60
- status information frames, 100
- STREAMS
 - DLPI specification, xxiii
 - modules, 27
- subnetwork number, 36
- Sun Web site, xxiv
- SunFDDI features, 1
- SunNet Manager, 103
 - console, 104, 106
 - daemons, 43
 - proxy agent, 103
 - schemas, 20, 108
- SunOS IP routing, 79
- SUNWnf package, 39
- SUNWpf package, 21
- switches, 5
- symbols, 51, 53
- synchronization, 51
- sys/dlpi.h file, 120
- system performance, 66

T

- T_max parameter, 65
- T_opr parameter, 26, 65
- T_Req parameter, 31
- T_req parameter, 26, 31, 65
- target
 - station, 104, 112
 - token rotation time, 26, 31, 65
- TCP
 - applications, 27
 - high water marks, 27
 - low water mark, 28
- TCP/IP networks, 103
- tcp_recv_hiwat parameter, 28
- tcp_xmit_hiwat parameter, 28
- tcp_xmit_lowat parameter, 29
- technical assistance, xxiv
- test, boot, 44
- theoretical
 - maximum bandwidth, 24
 - maximum number of stations, 24
- throughput, 24
- THT, 66, 144
- timed token
 - rotation protocol, 52, 65
- timing jitter, 144
- token, 65
 - ring, 79

- rotation timer, 66
- Token Ring, 65
- total ring length, 24
- traffic, 52
- transceiver unit, 6
- transfer rate, 66
- transmission time, 66
- transmit packets, 92
- transmitters, 51
- trap reports, 20, 114
- tree of concentrators, 75, 77, 80
- trees, 72
- troubleshooting checklist, 34
- TRT, 66, 145
- TTRT, 26, 31, 65, 145
- tuning
 - high performance networks, 24
 - the high water marks, 27
 - the low water marks, 28
 - the maximum transfer unit, 30
 - the socket options, 29
 - the target token rotation time, 31
 - the water marks at boot, 29
- twisting the ring, 34
- types, 100
- types of traffic, 26

U

- UDP
 - applications, 27
 - high water marks, 27
 - low water mark, 28
- udp_recv_hiwat parameter, 28
- udp_xmit_hiwat parameter, 28
- udp_xmit_lowat parameter, 29
- umount command, 18
- UNA, 96
- unrestricted traffic, 65
- unshielded twisted-pair, 50
- upstream station, 56
- user space, 25

- user-configurable hardware, 5
- utilities
 - pf_macid, 86
 - pf_smtmon, 38
 - pf_stat, 110
- UTP, 50

V

- video applications, 24
- voice applications, 24
- volume manager (vold), 13, 18

W

- wall boxes, 78
- water marks, 27
- Web site
 - Sun, xxiv
- work groups, 72
- wrapping, 55, 62

X

- X.3.139-1987, 2, 49
- X3.148-1988, 2, 49
- X3.166-1990, 2, 49
- X3.299 R7.3, 2, 49
- xmitP, 92

Reader Comment Card

Your comments and suggestions are important to us. Please let us know what you think about the *SunFDDI/P 1.0 Adapter User's Guide*, part number 802-6674-10.

1. Were the procedures well documented? Yes No

Please explain: _____

2. Were the tasks easy to follow? Yes No

Please explain: _____

3. Were the illustrations clear? Yes No

Please explain: _____

4. Was the information complete and easy to find? Yes No

Please explain: _____

5. Do you have additional comments about the *SunFDDI/P 1.0 Adapter User's Guide*?

You can send detailed comments via email to smcc-docs@sun.com, or send a fax to *SMCC Doc Feedback* at (415) 786-6443.

Your Name: _____

Title: _____

Company Name: _____

Address: _____

City: _____ State/Province: _____

Country: _____ Zip/Postal Code: _____

Email Address: _____

Telephone: _____

Thank you.

**AIR MAIL
PAR AVION**

IBRS/CCRI No. 808



INTERNATIONAL BUSINESS REPLY MAIL/REPOSE PAYEE

PERMIT NO. 808 MOUNTAIN VIEW CA

POSTAGE WILL BE PAID BY ADDRESSEE

**INFORMATION PRODUCTS
M/S MPK14-108
SUN MICROSYSTEMS INC
2550 GARCIA AVE
MOUNTAIN VIEW CA 94043-9551
UNITED STATES OF AMERICA**

NE PAS AFFRANCHIR

**NO POSTAGE
NECESSARY
IF MAILED
TO THE
UNITED STATES**

