

SunFDDI™ 6.0 on the Sun Enterprise™ 10000 Server



THE NETWORK IS THE COMPUTER™

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900 USA
650 960-1300 Fax 650 969-9131

Part No.: 806-3610-11
November 1999, Revision A

Send comments about this document to: docfeedback@sun.com

Copyright 1999 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, California 94303-4900 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. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd. For Netscape Communicator™, the following notice applies: Copyright 1995 Netscape Communications Corporation. All rights reserved.

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, Sun Enterprise, SunFDDI, SunNet Manager, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive 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.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 1999 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, Californie 94303-4900 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, la distribution, et la décompilation. Aucune partie de ce produit ou document 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. 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.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd. La notice suivante est applicable à Netscape Communicator™ : Copyright 1995 Netscape Communications Corporation. All rights reserved.

Sun, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, Sun Enterprise, SunFDDI, SunNet Manager, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées 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.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox 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 Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



Adobe PostScript

Contents

Dynamic Reconfiguration (DR) Issues	1
Alternate Pathing (AP) Issues	2
SunFDDI SAS and DAS Adapters	3
SunFDDI SAS Adapter Configurations	3
Point to Point	3
FDDI Concentrator	4
SunFDDI DAS Adapter Configuration	4
Ring Using No FDDI Concentrator	4
Dual Homing	5
SunFDDI DAS and SAS Adapter Configuration	6
Point to Point	6
FDDI Concentrator	6
SunFDDI SAS and DAS Adapter Issues	8
Optical Bypass	8
Broken Pipes	8

SunFDDI 6.0 on the Sun Enterprise 10000 Server

SunFDDI 6.0 on the Sun Enterprise 10000 Server explains the SunFDDI™ 6.0 software and SunFDDI SBus adapters (SAS/DAS) in the Sun Enterprise™ 10000 environment.

This document is intended for the Sun Enterprise system administrator, who has a working knowledge of UNIX® systems, particularly those based on the Solaris™ operating environment. If you do not have such knowledge, read the Solaris User and System Administrator AnswerBook™ documentation provided with this system, and consider UNIX system administration training.

Dynamic Reconfiguration (DR) Issues

When detaching a system board, the SNMP daemons must be killed and the SMT driver must be unloaded from the kernel when using Solaris 2.5.1 and Solaris 2.6 operating environments prior to detaching the system board.

This can be done by executing the following command from the host window command line:

```
# /etc/opt/SUNWconn/nf/bin/nf_snmd_kill
```

The Station Management (SMT) daemon starts two agents that collect and return FDDI statistics to a SunNet Manager™ (SNM) console. Without the SNMP daemons, network data will not be available for monitoring.

After the system board has been detached and at least one FDDI adapter is plumbed and up, you can restart the SNMP daemons and reload the SMT driver into the kernel. Type the following command from the host window command line:

```
#/etc/opt/SUNWconn/bin/nf_fddidaemon start
```

When you remove a system board that is part of a FDDI configuration, no data will travel through the FDDI adapters once the detach is completed.

To restore the configuration, attach the system board and plumb (`ifconfig`) the FDDI adapters. This allows the data to travel through the FDDI adapters.

The FDDI adapter must have the `inet` (IP address), `netmask`, and `broadcast` values set to be in an active state. The implications of having a FDDI adapter that does not pass data will vary with your specific topology.

Note – Check the SunSolve database for any patches that may be issued against DR and FDDI 6.0.

Alternate Pathing (AP) Issues

When using AP with FDDI, you must specify an unique MACID for the metanetwork. To ensure that the MACID is set properly at boot time, place `ifconfig` commands in the `/etc/rcS.d/S30rootusr.sh` startup script. For more information, refer to the *Alternate Pathing User Guide* for your version of the AP software.

Before detaching a system board that has an active metanetwork, the FDDI path group must be switched or DR will not allow the system board to be detached.

By switching the path group, the system board is detachable and the data being transmitted/received by the metanetwork continues uninterrupted. All AP devices, whether primary or secondary, are part of the ring. Detaching any FDDI device may constitute breaking the FDDI ring, depending on your specific topology.

SunFDDI SAS and DAS Adapters

The SunFDDI SBus adapters (SAS and DAS) with SunFDDI 6.0 software were verified on the Sun Enterprise 10000 against Solaris 2.5.1, Solaris 2.6 and Solaris 7 operating environments. This verification was designed to establish the maximum configuration; in other words, how many SunFDDI adapters can be supported per domain on the Sun Enterprise 10000.

Operating Environments	Maximum SAS Adapter Configurations (Sbus Only)	Maximum DAS Adapter Configurations (Sbus Only)
Solaris 2.5.1	7	7
Solaris 2.6	7	7
Solaris 7(64 bit)	9	9
Solaris 7(32 bit)	9	9

SunFDDI SAS Adapter Configurations

The following section describes the configurations used in testing SunFDDI SAS adapters.

Point to Point

FDDI SAS adapters in one domain are directly connected to FDDI SAS adapters in another domain and must have the same subnet address. For example, the FDDI SAS adapter in domain A with an IP address of 10.10.1.80 and the FDDI SAS adapter in domain B with an IP address of 10.10.1.81 constitute a valid configuration because both have the same subnet address (10.10.1).

- Advantages

- No FDDI concentrator needed.

- Disadvantages

- A maximum of two domains can be part of this configuration.

- Only two FDDI SAS adapters can communicate with each other. If one fails then the connection is lost.

Does not support alternate pathing (AP).

FDDI Concentrator

The FDDI concentrator connects one or more FDDI SAS adapters from multiple domains to the FDDI concentrator, via the M port(s). The FDDI concentrator must have multiple M ports as well as one A and one B port. The FDDI Concentrator used must be smart enough to know the subnet address for each FDDI SAS adapters.

As with Point to Point, the FDDI SAS adapters must have the same subnet address to communicate to one another.

This configuration is AP capable when a domain has two or more FDDI SAS adapters connect to the FDDI concentrator.

- Advantages

- Can communicate with multiple domains and multiple FDDI SAS adapters.
 - DR/AP capable.

- Disadvantages

- The FDDI concentrator is now part of the equation and must be smart enough to know subnet addresses for each of the FDDI SAS adapters connected to the FDDI concentrator.

SunFDDI DAS Adapter Configuration

The following section describes the configurations used in testing SunFDDI DAS adapters.

Ring Using No FDDI Concentrator

The ring configuration consists of connecting the A port of one FDDI DAS adapter to the B port of another FDDI DAS adapter and continuing this process until you have created a ring that encompasses the desired number of FDDI DAS adapters.

This configuration can go across domains. As with SAS, the subnet addresses must be the same in order for FDDI DAS adapters to communicate. The ring configuration is AP capable when two or more FDDI DAS adapters are installed into the ring and are configured in the same domain.

For example, two domains, A and B, connected in a ring configuration with NO FDDI concentrator or any other devices or HUBs involved.

Domain A consists of two system boards. Each system board has one FDDI DAS adapter. System board one has FDDI DAS adapter nf0 and system board two has FDDI DAS adapter nf1.

Domain B consists of one system board with one FDDI DAS adapter nf2.

Based on this information, the cables should be connected as follows to create a ring configuration:

nf0 (A port) is connected to nf1 (B port)

nf1 (A port) is connected to nf2 (B port)

nf2 (A port) is connected to nf0 (B port)

- Advantages

- No FDDI concentrator needed.

- Tolerates single failure.

- DR/AP capable.

- Disadvantages

- FDDI does not support optical bypass. Therefore, if a FDDI DAS adapter within the ring is not configured, the data will not be allowed to pass through that FDDI DAS adapter, thus breaking that part of the ring.

- Tolerates only a single failure in ring.

Dual Homing

If you have two FDDI concentrators, you can create a dual homing topology when the B port of a FDDI DAS adapter is connected to the M port of first FDDI concentrator and the A port of the same FDDI DAS adapter is connected to the M port of the second FDDI concentrator. You can repeat this process for as many M ports as exist on the FDDI concentrators.

Do not connect A and B ports of the same FDDI DAS adapter to the same FDDI concentrator.

As with SAS, the subnets address must be the same in order for FDDI adapters to communicate.

This configuration is AP capable when a domain has two or more FDDI DAS adapters connect to the FDDI concentrators.

- Advantages

- Two FDDI rings - active and standby.

- Tolerates multiple points of failure.

- DR/AP capable.

- Disadvantages

- The FDDI concentrator is now part of the equation and must be smart enough to know the subnet addresses for each of the FDDI DAS adapters connected to the FDDI concentrator.

SunFDDI DAS and SAS Adapter Configuration

The following section describes how the DAS and SAS FDDI adapters can be configured together.

Point to Point

Basically follow the same process as with SAS Point to Point, having the FDDI DAS adapter act like a FDDI SAS adapter by using only one port (A or B). Thus you connect the port of the FDDI SAS adapter to the selected port of the FDDI DAS adapter.

- Advantages

- No FDDI concentrator needed.

- Disadvantages

- A maximum of two domains can be part of this configuration.

- Only two FDDI SAS adapters can communicate with each other. If one fails then the connection is lost.

- Does not support Alternate Pathing (AP).

FDDI Concentrator

There are two FDDI concentrator configurations available for DAS adapters to communicate with SAS adapters; one uses a ring, the other does not. The following sections describe both.

DAS Adapters in a Ring Configuration to Communicate With SAS Adapters

First, connect the FDDI SAS adapters into the M ports of the FDDI concentrator. Second, determine where in the FDDI DAS ring configuration you want the FDDI concentrator to enter.

In the example below, the FDDI concentrator will be connected between FDDI DAS adapters nf1 and nf2. The A port of FDDI DAS adapter nf1 is connected into the B port of the FDDI concentrator, then the A port of the FDDI concentrator is connected into the B port of the FDDI DAS adapter nf2 to complete the ring.

For example:

Ring Configuration before FDDI concentrator

nf0 (A port) is connected to nf1 (B port)

nf1 (A port) is connected to nf2 (B port)

nf2 (A port) is connected to nf0 (B port)

Ring Configuration with FDDI Concentrator

nf0 (A port) is connected to nf1 (B port)

nf1 (A port) is connected to FDDI Concentrator (B port)

FDDI Concentrator (A port) is connected to nf2 (B port)

nf2 (A port) is connected to nf0 (B port)

■ Advantages

Can communicate with multiple domains and multiple FDDI adapters, both DAS and SAS.

DR/AP capable.

■ Disadvantages

The FDDI concentrator is now part of the equation and must be smart enough to know what the subnet addresses are for each of the FDDI adapters.

DAS Adapters Not in a Ring Configuration to Communicate With SAS Adapters.

Treat the FDDI DAS adapter like a FDDI SAS adapter by connecting either the A or B port into the M port of the concentrator. Then, from another domain, connect the FDDI SAS adapter port into the M port of the same concentrator.

Do not connect both A and B ports of the same DAS adapter to the same concentrator.

The A and B port of a concentrator can also be used to connect multiple concentrators together.

- Advantages

- Can communicate with multiple domains and multiple FDDI adapters, Both DAS and SAS.

- DR/AP capable.

- Disadvantages

- The FDDI concentrator is now part of the equation and must be smart enough to know the subnet addresses for each of the FDDI adapters connected to the FDDI concentrator.

SunFDDI SAS and DAS Adapter Issues

The following section describes outstanding issues with SAS and DAS adapters.

Optical Bypass

This feature is not supported by FDDI DAS or SAS adapters. To bypass a shutdown DAS adapter in a ring without a concentrator, you need an optical bypass (a device that allows you to connect incoming and outgoing fibers). There is no need to bypass a station if you use an FDDI concentrator. The FDDI concentrator provides the bypass.

Broken Pipes

When running heavy I/O load between domains using Solaris 2.6 software, you may experience broken pipes. Check the value of `tcp_rexmit_interval_max`:

```
# ndd -get /dev/tcp tcp_rexmit_interval_max
```

If the return value is not 60000, use the following command to reset the parameter:

```
# ndd -set /dev/tcp tcp_rexmit_interval_max 60000
```

To avoid losing this change between system reboots, you can create a script and put it in `/etc/rc2.d` directory.

