

AIX® Version 4.1.5r3

Release Notes

FXRNA/IS3

Third Edition (April 1998)

This edition of *Release Notes* applies to AIX 4.1.5r3 and to all subsequent releases of this product until otherwise indicated in new releases or technical newsletters.

THIS MANUAL IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. The laws of some states and countries do not allow the disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you. As such, the above warranty disclaimer shall only apply to the extent permitted by law.

While reasonable efforts have been made to assure the accuracy of this document, Motorola, Inc. assumes no liability resulting from any omissions in this document, or from the use of the information obtained therein. It is not warranted that the contents of this publication or the accompanying source code examples, whether individually or as one or more groups, will meet your requirements or that the publication or the accompanying source code examples are error-free. This publication could include technical inaccuracies or typographical errors.

Motorola reserves the right to revise this document and to make changes from time to time in the content hereof without obligation of Motorola to notify any person of such revision or changes

It is possible that this publication may contain references to, or information about Motorola products (machines and programs), programming or services that are not announced or released in your country or region. Such references or information may not be construed for any purpose to mean that Motorola intends to announce or release such products, programming or services in your country or region.

Any reference to a licensed program in this publication is not intended to state or imply that you can use only that licensed program. You can use any functionally equivalent program instead.

No part of this material may be reproduced or copied in any tangible medium, or stored in a retrieval system, or transmitted in any form, or by any means, radio, electronic, mechanical, photocopying, recording or facsimile, or otherwise, without the prior written permission of Motorola, Inc.

If the documentation contained herein is supplied, directly or indirectly, to the U.S. Government, the documentation is provided only as "restricted computer software" as defined in FAR 52.227-19 Commercial Computer Software--Restricted Rights.

© COPYRIGHT MOTOROLA, INC. 1998. ALL RIGHTS RESERVED. Printed in the United States of America.

© Copyright International Business Machines Corporation 1994. All rights reserved.

Contents

Chapter 1 Introduction

Purpose	1-1
Scope	1-2
Prerequisites.....	1-3
Overview of Contents.....	1-4
Related Information.....	1-5
Optional Software Information	1-5
README Files	1-5
Softcopy Documentation	1-5
Hardcopy Documentation.....	1-6
Getting Help for System Problems	1-8

Chapter 2 General AIX Issues

Overview	2-1
What's New in this Release	2-2
Value-Added Features	2-2
xdiff Utility Available	2-2
Environmental Monitor Commands.....	2-3
Environmental Monitor Special Files (Devices)	2-3
Installation.....	2-4
Installation Resources	2-4
AIX 4.1.5 Shipped on Multiple CD-ROMS	2-4
Installing AIX 4.1.5 on an Existing AIX System	2-4
bos.content_list Fileset	2-4
Trusted Computing Base (TCB).....	2-5
Updating to AIX 4.1.5. in a NIM Environment	2-5
Do Not Install devices.xrft with Network Installation Management	2-5
Migration of Diskless/Dataless Clients and NIM SPOTs.....	2-5
Installing NetWare.....	2-6
Node-locked Licenses Unavailable for Some Products	2-6
Checking the License Server	2-7
Base Operating System.....	2-8
New splitlvcopy Command.....	2-8
New pppstat Command	2-9
National Language Support (NLS).....	2-10
Korean, Chinese (Traditional), and Latin American Spanish Keyboards	2-10

Chinese (Simplified) Keyboard Translate Table	2-10
Chinese (Simplified) Input Method Support.....	2-11
Error Logging with errpt Command	2-11
Known Problems on All Systems	2-12
lscfg and lsdev Commands Do Not Return Memory Size	2-12
SMIT Options Incorrect for SNA Protocol	2-13
getattr() Routine Does Not Release Memory	2-14
InfoExplorer May Hang with more than 64 Databases	2-15
Known AIX 4.1 Problems Filed with the Vendor	2-16

Chapter 3 FX Series Systems

Overview	3-1
FX Series Operating System	3-2
Split Mode Functionality	3-2
Managing CPU Module Firmware Levels Across Systems.....	3-5
Paging Space Values.....	3-6
Default Dump Devices	3-6
System Reboots after a System Dump Taken	3-6
Network Installation Management Not Supported	3-6
Tape Devices Diagnostic Not Supported	3-7
Data Integrity Checker.....	3-7
On-Demand Diagnostics	3-7
New rpthours utility available for FX Series Systems	3-10
AIX commands added for FX Functionality	3-11
AIX commands added for Split Mode Functionality	3-12
General Hardware Notes	3-13
Setting up Telco Alarms.....	3-13
Use ESD Mitigating Procedures	3-13
Offlining Personality Modules Not Recommended.....	3-13
Hot Pull Support.....	3-13
CPU Module Subsystem	3-14
Set Time on New CPU Modules.....	3-14
CPU Module Status after a Reboot	3-15
Do Not Use fctl Program to Online and Offline CPU Modules	3-15
OS Watchdog.....	3-15
ECC Parity Handling	3-16
MFIO Module	3-17
Upgrading Disks from 2Gb to 4Gb.....	3-17

Contents

Interconnect Module (ICM)	3-24
Any Interconnect Module Failure in a Domain Takes Down the Entire Domain.....	3-24
CPU Modules May Go Offline When an Interconnect Module Fails...	3-24
Do Not Remove the Interconnect Module without Disconnecting the Power Cord	3-25
Power Module and Fantrays	3-26
Power Cycle the System Using Power Module Latches	3-26
The Power Daemon Takes Time to Complete Its Shutdown.....	3-26
Do Not Boot System with the Fantrays Removed	3-26
Boot System Carefully after Failure of All Power Modules	3-27
Boot System Carefully after Failure of All Fantrays	3-28
Known Problems on FX Systems	3-29
Overview.....	3-29
Creating Swap Partitions While Volume Groups are Nonftvg Interferes with Disk Mirroring	3-30
Time Zone not Set on SYSNEW when using Split Mode to do a BOS-NEW Installation	3-32
Loading and Unloading Kernel Extensions Leaks Memory	3-33
Cannot Create split-notify.-%-.tmp File without Write Access to Current Directory	3-34
Message Indicates a Failure to Mirror Rootvg when doing a Preservation Install.....	3-35
System may Acquire Wrong Status for Asynch and Ethernet Personality Modules at Reboot.....	3-36
fxnvtcl does not Remove Lock File when Interrupted.....	3-37
File Systems Named “jfslog” are not Included in File System Listings	3-38
Cron Jobs may Generate Unnecessary E-mail against rotlog.....	3-39
Secondary May Hang During Boot Cycle.....	3-40
Possible Bootlist Error if Rootvg includes Multiple Hdisks.....	3-41
Hours in Service in Module EEPROMs may be Counted Twice	3-42
CPU Reintegration Has Temporary Critical Fault Under Heavy Load	3-43
Making Split Mode State Changes while Booting Secondary System can Corrupt the Secondary Boot Record	3-44
Upgrading from AIX 4.1.4r6 to AIX 4.1.5 Causes Problems with ASYNC Logical Devices	3-45

Chapter 4 Documentation Notes

Overview 4-1
FX Series Administering Your Fault Tolerant System..... 4-2
Installing FX manpages..... 4-3

Table 2-1. Keyboard Configuration Information	2-10
Table 2-2. TARs Filed with IBM.....	2-16
Table 3-1. AIX Commands added for FX Series Systems	3-11
Table 3-2. AIX Commands Added for Split Mode Functionality	3-12

Introduction 1

Purpose

The *AIX 4.1.5 Software Release Notes* contain information on new, changed, or removed functionality of the AIX operating system that was not available at the time of publication of the AIX documentation, and contains additional information not documented in other publications. The *Release Notes* also contain information on the functionality of FX Series systems.

Scope

AIX 4.1 .5r3 Release Notes contain information on:

- general notes about AIX 4.1, including new or changed features
- known problems in this release of the AIX 4.1 operating system
- notes concerning FX Series systems
- changes and corrections to documents in the AIX 4.1 product library

Prerequisites

Some instructions in this guide assume that you are familiar with:

- AIX system administration commands
- The System Management Interface Tool (SMIT)
- FX Series systems
- software installation operations, such as apply and commit

For general information about AIX, refer to the *AIX Version 4.1 System Management Guide: Operating System and Devices*.

For FX-specific information, refer to *Administering Your Fault Tolerant System* and *Managing System Storage*, in conjunction with the *FX Series Hardware Installation* and *FX Series Operating System Installation* guides.

Overview of Contents

The *Release Notes* are organized as described in the table below.

This Chapter...	Discusses...
Chapter 2, "General AIX Issues"	general issues pertaining the the AIX operating system
Chapter 3, "FX Series Systems"	software and hardware issues relating to the FX Series
Chapter 4, "Documentation Notes"	changes and corrections to the AIX and the FX Series libraries

Related Information

Optional Software Information

Licensed programs include AIX 4.1 and optionally purchased products. Each program can have one or more separately installable filesets. The Package Installation Database for Current Media is an InfoExplorer database that contains information about the contents of these separately installable filesets.

README Files

For information on how to view the README files for the base operating system (BOS) and optional software products, refer to the *AIX 4.1 Installation Guide*.

Some README files are shipped in separate directories that are not accessible if you follow the procedure described in the *AIX 4.1 Installation Guide*. To find the README files installed on your machine, use this command:

```
find /usr -name "*README" -print
```

Softcopy Documentation

- The FX Series titles listed below are available in pdf format under MCG Systems, at:
<http://www-public.phx.mcd.mot.com/ext/literature/PDFLibrary/>
- The majority of the IBM AIX documentation library is shipped electronically with the BOS CDs with the InfoExplorer help system. Refer to the *Installation Guide* for directions on installing InfoExplorer. Once installed, InfoExplorer can be invoked by typing **info** at the command line.

Hardcopy Documentation

Custom hardcopy documentation for general AIX and for FX Series systems can be ordered by calling the Motorola Computer Group Literature Center at 888-432-1877 or at 602-804-7378.

AIX Version 4.1 Documentation for all Systems

The following hardcopy AIX documentation ships with your system (non FX):

- *AIX 4.1 Release Notes*
- *AIX 4.1 Installation Guide*
- *AIX System Diagnostics (ASD) Quick Reference*
- *Installing and Configuring NCDware*

The following documents may be especially useful to novice AIX users:

- *Guide to System Information*
- *Getting Started*
- *Quick Reference*
- *iFOR/LS System Management Guide*
- *iFOR/LS Tips and Techniques*
- *Problem Solving Guide*
- *Messages Guide and Reference*
- *Performance Tuning Guide*

Other MCG titles for generic AIX include:

- *Installation Checklist*
- *Quick Installation Guide*
- *Network Installation Management Guide and Reference*

-
- *Network Installation Management Tips and Techniques*
 - *VMEbus Device Driver Reference*
 - *SVR4 Porting Guide*
 - *Managing System Storage*
 - *Configuring and Maintaining the System*

FX Series Documentation

Additional hardcopy FX documentation includes:

- *Administering your Fault Tolerant System*
- *Managing System Storage*
- *Motorola FX Series Fault Tolerant Architecture Overview*
- *Configuring and Maintaining the System*
- *Operating System Installation Guide*
- *Operating System Installation Troubleshooting*
- *FX Series Diagnostics and Troubleshooting*
- *FX Series LED and Alarm Quick Reference*
- *Making and Using Backups*
- *FX Series System Hardware Installation Guide*

Note *Managing System Storage* and *Configuring and Maintaining the System* have been listed under both AIX documentation and FX Series documentation, because while they treat generic AIX, they also have sections which are specific to the FX Series.

Getting Help for System Problems

If you encounter difficulties contact your Motorola Computer Group Sales office or Motorola Computer Group's customer support group at:

- U.S.A. 1-800-551-1017
- Canada 1-800-387-2416
- Maidenhead, U.K. 44-1628-39121
- Paris, France 33-1-467-43560
- Duesseldorf, Germany 49-211-65899-55

When you call, please be prepared to provide the following information:

- the type and configuration of your FX Series system
- the level of AIX that you are running
- your system serial number
- the name of your company, your name, and a telephone number
- a brief description of the problem, including the severity of its impact on your ongoing efforts

This information will be forwarded to the appropriate technical engineering contact, who will return your call promptly.

General AIX Issues **2**

Overview

This chapter includes information on:

- new generic features of AIX 4.1.5r3
- value-added features specific to Motorola's AIX 4.1.5r3 product
- general guidelines about AIX 4.1.5r3
- known problems along with workarounds where possible

What's New in this Release

The following features have been implemented in AIX 4.1.5r3:

- support for IBM's version of AIX 4.1.5
- ability to upgrade the system from AIX 4.1.4r6 or from AIX 4.1.5r2.
- Netscape Navigator value pack
- Year 2000 support

Value-Added Features

The AIX 4.1 product is based on IBM's AIX 4.1 operating system. Motorola has made changes to this product to support Motorola-specific hardware platforms, while retaining the functionality and command structure of IBM's AIX 4.1 operating system. This section describes value-added features that apply to all systems running Motorola's version of the AIX 4.1 operating system.

xdiff Utility Available

The `xdiff` program is available in AIX 4.1 Service Level AOS1.2 and later releases. `xdiff` is a graphical interface to the `diff` command that allows side-by-side viewing of differences between files and automated merging of selected differences to create a new version of a file.

`xdiff` is automatically installed with the Personal Productivity bundle. Alternatively, you can install it directly from the `X11.apps.motorola` fileset.

Environmental Monitor Commands

These new commands are used with the environmental monitor and are supported on all systems except RISC PC (and RISC PC Plus) and FX systems:

- envctl
- envdiag
- rc.motpowerfail

These commands are documented in the manual pages in section 1M. You must use the `man` command to access these manual pages; they are not available through InfoExplorer.

Environmental Monitor Special Files (Devices)

These pages describe aspects of environmental monitor special files (devices) and are supported on all systems except RISC PC, RISC PC Plus, PowerStack II, and FX systems:

- envmon
- envmon_msg
- envmon_odm

You can find these manual pages in section 7. You must use the `man` command to access these manual pages; they are not available through InfoExplorer.

Installation

Installation Resources

For more detailed descriptions of general AIX installation procedures and troubleshooting help, see the *Installation Guide* and the *Quick Installation Guide*. For FX Series Systems, see the *FX Series Operating System Installation Guide*, the *FX Series Operating System Installation Troubleshooting Guide*, and the *FX Series System Hardware Installation Guide*.

AIX 4.1.5 Shipped on Multiple CD-ROMS

AIX 4.1.5 is shipped on multiple media. The BOS and some additional software is installed on Volume 1. The DCE and additional software are installed on Volume 2. There are also two value added CDs.

Installing AIX 4.1.5 on an Existing AIX System

To avoid problems during installation, place the entire system in a dormant state by:

- making sure that all users are logged off
- stopping all applications

If applications are left running, unpredictable results may occur and some updates may not be applied.

bos.content_list Fileset

AIX 4.1.5 includes a new, optionally installable fileset, `bos.contents_list` that lists all installable files and the filesets/packages that contain them. The fileset includes one file, `/usr/lpp/bos/AIX_file_list`.

Trusted Computing Base (TCB)

AIX 4.1.5 contains some changes to permissions of files for security reasons. For example, the `/etc/passwd` file has changed permissions from 664 to 644. Because of a problem with reject and Trusted Computing Base (TCB), we recommend that you commit the 4.1.5 updates if you have the TCB enabled on your system.

To determine whether TCB is enabled or not, run the `tcbck` command as root. If TCB is enabled, `tcbck` displays a syntax message. If TCB is not enabled, an error message displays.

Updating to AIX 4.1.5. in a NIM Environment

If you are updating a NIM environment from a version of AIX 4.1 earlier than Service Level AOS1.2, you should first update the NIM software on the master, clients, and SPOTs. As a result, you will have access to a function that permits an easier update using NIM. The enhanced update facilities are described in the *Network Installation Guide and Reference*.

Do Not Install `devices.xrft` with Network Installation Management

If you are using Network Installation Management (NIM) on your systems, do not install the fileset(s) `devices.xrft` into a NIM “spot” or an `lpp_source` resource. This fileset is incompatible with the NIM environment and is not used on most systems.

Migration of Diskless/Dataless Clients and NIM SPOTs

Migration to a new level of AIX currently is not supported for diskless and dataless clients. Also, migration of a SPOT that is not a converted `/usr` file system is not supported.

Installing NetWare

When installing NetWare for AIX on your system using a language other than English, do the following:

1. At the system prompt, type
`locale`
2. Note the value of the `LC_ALL` variable.
3. Set the `LC_ALL` variable to `en_US` using the command:
`export LC_ALL=en_US`
4. Install NetWare.
5. After installing NetWare, reset the `LC_ALL` variable to the original value noted in step 2.

Node-locked Licenses Unavailable for Some Products

The *iFOR/LS System Management Guide* and *iFOR/LS Tips and Techniques* include information about and references to node-locked licenses. However, node-locked licenses are not available for some products that are NetLS license-enabled. Only concurrent access licenses are generated for these products.

Motorola products that this affects are:

- C for AIX
- C Set ++
- XL FORTRAN
- XL Pascal
- IBM COBOL Set for AIX
- Configuration Management Version Control (CMVC)

You must use a concurrent access license with the above products.

Checking the License Server

In order to install software licenses on an AIX system, the license server must be running. Before installing software licenses, check to make sure the license server is available using SMIT:

Processes & sSbsystems->Subsystems->List all Subsystems

Find the entry containing “netlsd” and confirm that it is active:

```
netlsd          netlsd          11846  active
```

If netlsd is not active, use SMIT to start the license server:

Processes & Subsystems->Subsystems->Start a Subsystem

And choose netlsd from the list of subsystems.

Base Operating System

New `splitlvcopy` Command

The `splitlvcopy` command has been added for AIX 4.1.5. This command splits copies from one logical volume and creates a new logical volume for them.

To use this command, you must either have root user authority or be a member of the system group.



Although the `splitlvcopy` command can split logical volumes that are open, including logical volumes containing mounted file systems, this is not recommended. You may lose consistency between `LogicalVolume` and `NewLogicalVolume` if the logical volume is accessed by multiple processes simultaneously. When splitting an open logical volume, you implicitly accept the risk of potential data loss and data corruption associated with this action. To avoid the potential corruption, close logical volumes and unmount file systems before splitting.

You cannot use SMIT to run this command. Message catalogs are not supported for this command and therefore the error messages are provided in English only with no message catalog numbers. Documentation for `splitlvcopy` consists of the man page. No additional information for `splitlvcopy` exists in InfoExplorer.

New pppstat Command

2

The pppstat command has been added for AIX 4.1.5. This command extracts and displays the RAS (Reliability, Availability, and Serviceability) information of the PPP (Point to Point Protocol) subsystem. This command only applies to AIX 4.1.5.

National Language Support (NLS)

Korean, Chinese (Traditional), and Latin American Spanish Keyboards

The mappings and input methods for Korean, traditional Chinese, and Latin American Spanish locales have been enhanced to support the new PS/2 keyboards. However, because these keyboard options do not appear under the Manage Language Environments menus in SMIT, you need the following information to configure the system with these keyboards:

Table 2-1. Keyboard Configuration Information

Language	Command
Korean	<code>chkbd /usr/lib/nls/loc/KW@alt.lftkeymap < /dev/lft0</code>
Traditional Chinese	<code>chkbd /usr/lib/nls/loc/zh_TW@alt.lftkeymap < /dev/lft0</code>
LA Spanish	<code>chkbd /usr/lib/nls/loc/es_LA.lftketmap < /dev/lft0</code>

Chinese (Simplified) Keyboard Translate Table

In AOS1.3 and later releases of AIX 4.1, you can view the keyboard information for the Simplified Chinese Keyboard (Code Set IBM-eucUCS) in the file:

```
/usr/lpp/X11/defaults/xmodmap/zh_CN/keyboard
```

If you are running with the ZH_CN locale (UTF-8 codeset), this keyboard information is located in the file:

```
/usr/lpp/X11/defaults/xmodmap/ZH_CN/keyboard
```

Chinese (Simplified) Input Method Support

As of the AOS 1.3 release of AIX 4.1, support for the 5-stroke input method has been removed. Pressing Shift-F5 and Shift-F6 no longer invokes the 5-stroke input method under the simplified Chinese (zh_CN) locale.

Error Logging with errpt Command

When you run the `errpt` command under a non-English language environment, in some cases the system may return a 4-digit hexadecimal value instead of the error message text. This indicates that the system cannot find a message number for that message in your language's codepoint catalog.

Change the language environment setting to English with the following command:

```
LANG=C
```

and then run the `errpt` command. The English text for the error message is displayed.

Known Problems on All Systems

This section contains information on known problems in this release of AIX 4.1. Workarounds have been provided whenever they are available.

For additional troubleshooting information, refer to *Installation Troubleshooting* and *Problem Solving Guide and Reference*.

lscfg and lsdev Commands Do Not Return Memory Size

Tracking Number

TAR 1-39744 and TAR 1-39647

Problem

As of AIX 4.1 Service Level AOS1.2, the `lscfg` and `lsdev` commands no longer return the system memory size.

Workaround

To obtain the total size of memory available on your system, use any of the following commands:

- `lsattr -E -l sys0 -a realmem`
- `lsattr -E -l mem0`
- `bootinfo -r`

Note You must have root access to use the `bootinfo -r` command.

SMIT Options Incorrect for SNA Protocol

2

Tracking Number

TAR 1-46856

Problem

The SMIT interface for selecting the protocol to be run on the integrated SCC ports has changed, causing a minor change to the SMIT procedure for configuring these ports for Express SNA communications only. The SMIT default value “hdle, sna” is not valid.

Workaround

Manually type in the protocol you want, i.e. “sna” or “hdlc.”

getattr() Routine Does Not Release Memory

Tracking Number

TAR 1-43187

Problem

The `getattr()` routine allocates memory for the `CuAt` structures that are returned. This memory is automatically freed when your application exits. However, you need to free this memory in your application if it invokes `getattr()` routine several times and runs for a long time.

Workaround

None.

InfoExplorer May Hang with more than 64 Databases

2

Problem

If you try to read a CD with more than 64 InfoExplorer databases, InfoExplorer may core dump.

Workaround

To prevent the core dump, you should run the following command before invoking InfoExplorer:

```
export MALLOCTYPE=3.1
```

Known AIX 4.1 Problems Filed with the Vendor

Motorola's AIX 4.1 operating system release is based on IBM's AIX 4.1 release.

The following problems have been filed with IBM. Resolutions of these problems are pending fixes from IBM.

Table 2-2. TARs Filed with IBM

TAR Number	Description
1-43408	makestack has hard-coded list of tape devices.
1-37149	Insufficient physically contiguous memory is available.
1-39425	You cannot print more than 999 files correctly.
1-39587	The network boot hard codes device uniquetypes.
1-39647	Refer to "lscfg and lsdev Commands Do Not Return Memory Size" on page 2-12 for details.
1-39768, 1-38533	Refer to "InfoExplorer May Hang with more than 64 Databases" on page 2-15 for details.
1-42956	The stress tests crash the system.

FX Series Systems 3

Overview

AIX 4.1.5r3 contains special functionality for FX Series System Release 1.2. Motorola has made changes to IBM's AIX 4.1.5 OS in order to support Motorola-specific hardware platforms and to add fault tolerance, while retaining the functionality and command structure of IBM's AIX 4.1 operating system.

This chapter describes information about and guidelines for FX Series systems.

FX Series Operating System

The following issues are related to FX Series operating system.

Split Mode Functionality

In previous releases of the FX Series, it has been necessary to take the system off-line for an extended period of time in order to upgrade CPU modules, or to upgrade firmware or software.

Split Mode functionality has been added as part of the FX 1.2 release in order to minimize the disruption of service during such upgrades. Split Mode allows you to take a fully redundant fault tolerant FX Series system and divide it into two simplex systems. One side is primary and provides active service, while the other is upgraded.

Note While Split Mode allows you to boot the secondary system at will, booting the primary/active system during Split Mode will cause the system to become unavailable.

Because Split Mode requires the system to be made simplex, you should not use it to replace existing hardware or to upgrade modules for which other means of upgrade exist. Split Mode is intended to be used only during CPU module upgrades, CPU module firmware upgrades, and software upgrades.

In order to use Split Mode, your system must be running AIX 4.1.5r3 and it must have version 1.3 of the FXBug firmware.

Detailed information on Split Mode is available in the *Split Mode Overview and User's Guide*.

Updating to FXBug Version 1.3 Firmware

An **fwupdate** command has been added to the AIX 4.1.5r3 OS that automates part of the process of updating to the 1.3 version of the firmware. The script also gives appropriate guidance concerning the remaining steps. You must have at least two CPU modules active on the system in order to use the **fwupdate** command.

The following procedure details the steps for upgrading firmware via the **fwupdate** command.

Note This procedure requires that you reboot the system.

1. Install AIX 4.1.5r3 using standard procedures. For more information, refer to the Operating System Installation Guide for FX Series.
2. Execute the **fwupdate** command from the command line:

```
fwupdate
```
3. Copy down instructions that are output to the screen.
4. Disable autoboot on the master CPU:

```
autoboot -c X -o ab=n -o vb=n -o rb=n -o nb=n
```

where *X* is the number of the master CPU module.
5. Following the instructions that you have written down, reboot the system.
6. Type the remaining instructions at the `FX-Bug>` prompt one at a time.
7. Boot the system using the `pboot` command.
8. When the system comes back online, verify that it is running "FX Series Debugger release version 1.3." If not, begin this procedure again.
9. Bring the checker CPU module(s) back on line:

```
modchange -r -l CPU-X
```

Where *X* is the number of the off-lined CPU module.

Do not Boot the Primary/Active System

Split Mode allows you to upgrade a passive/secondary system while the active/primary system provides service. Split Mode is designed to allow unlimited reboots on the passive system during the upgrade process. However, it does not allow for reinstalling or rebooting the primary system; in fact it is designed to obviate the need for a primary reboot. **Rebooting the primary system will cause the entire FX System to become unavailable** and will require system-level recovery procedures.

Use the pboot Command with Upgraded CPU Modules, rather than the gevboot Command

The gevboot command may be unable to access the boot device after a CPU module has been upgraded. This problem is caused because gevboot relies on the NVRAM variable fw-boot-path (the "bootlist"), and the contents of the new CPU's bootlist are indeterminate until one of the following three things has been done:

- the operating system is (re)installed,
- the bootlist command is run manually on SYSNEW, or
- split mode operation is complete

Use the pboot command instead of the gevboot command with newly upgraded CPU modules.

LEDs and Console Messages Reflect Module Errors even when that is the Expected Module State.

When mbus restrictions are imposed during Split Mode, module LEDs reflect that the modules are out of service, and console messages appear stating that the module and its LED's are inaccessible. This is the expected state of the modules and the messages may be safely ignored.

Switched Ethernet Hubs May Cause a Networking Delay

After SYSOLD and SYSNEW trade IP addresses during Split Mode, there may be a several minute delay as switching hubs try to recognize the new configuration.

Nonswitched hubs are not effected.

SMPE Utility Does not accept * Metacharacter

The Split Mode smpe utility does not accept the wildcard (*) metacharacter. It is necessary to include the entire file name when using smpe.

Managing CPU Module Firmware Levels Across Systems

In order to support Split Mode, the CPU module firmware for FX Series Release 1.2 has been revised to “FX Series Debugger release version 1.3,” or FX-Bug 1.3. Previous FX Series systems were released with version 1.2 of the FX-Bug firmware.

In order to run synchronously, all CPU modules in a system must be running the same version of firmware. If you place a CPU module with the 1.2 version of the firmware into a system running AIX 4.1.5r3, the module’s firmware will be automatically flashed to version 1.3 before it is integrated. This will allow you to use existing CPU modules in an AIX 4.1.5r3 system.

Current AIX 4.1.4r6 and AIX 4.1.5r2 systems, however, are not able to integrate a CPU module with version 1.3 of the firmware. Contact your service provider to see when a PTF will be available which will enable AIX 4.1.4r6 and AIX 4.1.5r2 systems to flash CPU module firmware to the 1.2 version.

Paging Space Values

When AIX reaches or exceeds the paging space warning level, it may kill processes, including all the system daemons. Make sure that paging space is set to the recommended value in SMIT. The SMIT screen shows you the current value and the recommended value.

FX Series daemons do not handle the `SIGDANGER` signal and will be killed like any other AIX process.

Note The behavior of paging space on an FX Series system is no different than on any other AIX 4.1 system.

Refer to the *Managing System Storage* guide for additional information on paging space.

Default Dump Devices

The default primary dump device for FX Series systems is in the `/dev/sysdump0` logical volume. This default is different from the default of the other Motorola AIX 4.1 systems.

To recover the dump, use the standard dump recovery procedures from a logical volume.

System Reboots after a System Dump Taken

The default value of the `sys0` attribute, `autorestart` is `true`. This causes the FX Series system to reboot once a system dump is taken.

Note This behavior is different from that of our other AIX 4.1 systems.

Network Installation Management Not Supported

Support for Network Installation Management is not available on FX Series systems.

Tape Devices Diagnostic Not Supported

AIX System Diagnostics are supported. However, the diagnostic for tape devices is not supported.

Data Integrity Checker

Future releases of the AIX operating system may include support for the Data Integrity Checker (DIC). The DIC will compute checksums on all fixed disk transfers and generate errors on checksum errors. Approximately one percent of the disk space on every fixed disk attached to the system is appropriated for this anticipated functionality.

On-Demand Diagnostics

On-Demand Diagnostics have been added to AIX 4.1.5 as a new functionality on FX Series systems. This feature allows system administrators to do bug-level diagnostics on individual modules from the AIX command line using either the **offtest** or the **vterm** utilities. On-Demand Diagnostics are currently available for CPU modules, Asynchronous I/O Modules, MFIO modules, and Ethernet controller modules.

Administrators may want to use **offtest** to diagnose modules which have been off-lined by the Configuration Management System (CMS) and to decide whether to try reintegrating the modules or to replace them. The **offtest** utility can also be used to periodically test modules in order to detect any potential problems before the modules fail. Modules must be taken off line prior to running the On-Demand Diagnostics.

Once invoked, either from the command line or from an automated script such as a **cron** job, **offtest** accesses a set of bug-level diagnostic commands and runs them sequentially. If the module passes all tests, **offtest** returns a value of 0; a non-zero return value indicates that the module failed one or more of the diagnostics.

The default test scripts, which are located in the `/etc/ft` directory, include a combination of bug diagnostic commands and higher level test commands, which are preceded by '!'. In general, the scripts:

- turn the module's power off,
- turn the power back on after a brief delay,
- set the bug's MPCR register offset,
- set the bug's FIFO register offset,
- establish a virtual terminal connection to the module,
- mimic the appropriate bug prompts on the console,
- run maintenance bus commands,
- run bug diagnostic commands, and
- log messages.

Meanwhile, the **offtest** utility analyzes the response from the bug testing, terminates the testing, and reports the overall result to the log file or to standard output.

The scripts, which contain information about the purpose and the duration of each test, may be easily modified using any standard text editor, or you can instruct the **offtest** utility to use a separate script altogether. An example of where this may be useful is if you would like to run abbreviated scripts on a regular basis, but also want recourse to the entire suite of tests available in the default scripts.

When run in verbose mode, **offtest** provides a list of the bug-level tests that have been run and the results of each test. Operators are able to use a `-L` option to redirect output to a log file. Log file data will be similar to what appears on the console. The following example is a section of the log file from a CPU test:

```
Fri Jun 27 10:15:00 1997: Testing 01-W1566B03D,  
serial number FP999, in Slot c1-f10
```

```
Fri Jun 27 10:15:00 1997: Using: test script
    /etc/ft/CPU.offtest
Fri Jun 27 10:15:00 1997: TEST STARTED
Fri Jun 27 10:15:00 1997: Ensure power is off
Fri Jun 27 10:15:14 1997: Turning power on
Fri Jun 27 10:15:48 1997: Opening virtual
    terminal connection
Fri Jun 27 10:15:51 1997: Virtual Terminal
    Opened for 0xCF100000
Fri Jun 27 10:15:58 1997: --Break Detected--
Fri Jun 27 10:15:58 1997: FX-Bug>sd
Fri Jun 27 10:16:00 1997: sd
Fri Jun 27 10:16:02 1997: ver
Fri Jun 27 10:16:02 1997: Debugger/Diagnostics
    Type/Revision.....=FX Series/1.2
Fri Jun 27 10:16:02 1997: Debugger/Diagnostics
    Revision Date.....=05/05/97 (RM01)
Fri Jun 27 10:16:02 1997: MicroProcessor
    Version/Revision.....=0004/0305
Fri Jun 27 10:16:02 1997: CPU Module Baseboard
    Artwork Revision.....=6
Fri Jun 27 10:16:02 1997: CPU Module Baseboard
    PLD Major/Minor Revision.....=02/08
Fri Jun 27 10:16:02 1997: Checking the BUG Version
    by looking for "RM01"
Fri Jun 27 10:16:02 1997: Checking the Memory Size
Fri Jun 27 10:16:04 1997: md 0x8e230000:1;b
Fri Jun 27 10:16:04 1997: 8E230000 3A :
Fri Jun 27 10:16:04 1997: Memory Size from CPU:
    0x8000000 (134217728) bank 0 0x8000000 bank 1 0
Fri Jun 27 10:16:04 1997: ** Switching to the
    FX-Bug Diagnostics **
Fri Jun 27 10:16:06 1997: sd
Fri Jun 27 10:16:08 1997: cf
```

```
Fri Jun 27 10:16:08 1997: RAM Configuration Data:
Fri Jun 27 10:16:10 1997: N?
Fri Jun 27 10:16:10 1997: y
Fri Jun 27 10:16:12 1997: 00000000?
Fri Jun 27 10:16:12 1997: 00004000
Fri Jun 27 10:16:14 1997: 01000000?
Fri Jun 27 10:16:14 1997:.
Fri Jun 27 10:16:16 1997: onb dirty
Fri Jun 27 10:16:16 1997: ONB DIRTY:
    Dirty RAM..... Running ---> PASSED
Fri Jun 27 10:16:18 1997: onb edc_cb
Fri Jun 27 10:16:18 1997: ONB EDC_CB: EDC Check
    Bit Test..... Running ---> PASSED
```

The **vterm** utility provides direct access to the bug, so that individual diagnostics can be run directly from the command line. This may be useful, for example, when the error log shows that the module has failed a single test, and you would like to run the single test again. Some bug-level tests, including “onb dirty,” may return false results under the **vterm** utility because of conflicts over system resources.

For more information on On-Demand Diagnostics, see the **offtest** and the **vterm** manpages.

New rpthours utility available for FX Series Systems

A **rpthours** utility has been added to AIX 4.1.5 to provide information on the availability of FX system modules. The script is intended to be run by the system administrator on a monthly basis.

AIX commands added for FX Functionality

The following commands, and their manpages, have been added to standard AIX in order to support FX Series functions. For directions on how to load these manpages, see “Installing FX manpages” on page 4-3.

Table 3-1. AIX Commands added for FX Series Systems

autoboot	autopair_gen	cfgftmod
changed	chftvg	chftvgdev
chgftldev	chgftmod	cmsnotify
defcms	eebuild	eecpufix
eeprogram	eeeprom	errtrans
ftbugver	ftcmsuser	ftcpuclean
ftdecode	ftgather	ftreport
fttree	ftupdateVPD	ftvalue
fwupdate	fxnvcv	fxnvctl
lightsd	lsdomain	lsmatch
lsmod	lsmodparent	lsstate
mkftvg	modcfgmgr	modchange
modconvert	modtest	modundefine
offtest	phonehome	pmfstat
powerd	restart	rotlog
rpthours	settod	ucfgftdev
ucfgftldev	ucfgftmod	vterm

AIX commands added for Split Mode Functionality

The following commands, and their manpages, have been added to standard AIX in order to support Split Mode functions. For directions on how to load these manpages, see “Installing FX manpages” on page 4-3.

Table 3-2. AIX Commands Added for Split Mode Functionality

cfgprid	flash_ext	ftctl
ISC_Abort	ISC_Conn	ISC_Dereg
ISC_Disc	ISC_Exec	ISC_ExecTerm
ISC_LogMSG	ISC_Read	ISC_Recv
ISC_Reg	ISC_RegRecv	ISC_Send
ISC_Status	ISC_Write	ISCU_broadcast_message
ISC_log_message	ISCU_send_message	ISCU_smpe
ISCU_transfer_file	mb_split	pflash
pridd_ext	sm_apps	sm_notify
smft	smmt	smpe
splitstate	spope	spstart
spstop		

General Hardware Notes

The following issues are related to FX Series system hardware.

Setting up Telco Alarms

You can classify the severity of a fault by noting the Telco status LED that is lit. These LEDs are found on the front of the Fan Trays. In addition you can hook up the corresponding alarm relay found on the back of the InterConnect Module to the Central Office Alarm bus. If you have a Model 626, do NOT connect the alarm relays on the back of InterConnect Module in the Expansion Chassis. Use the alarm relays on the back of the InterConnect Module in the Main Chassis only.

Use ESD Mitigating Procedures

ESD mitigating procedures should be used at all times when working around the FX series. It is especially important to use this protection when inserting tapes into the tape drive.

Offlining Personality Modules Not Recommended

Do not offline a personality module directly. Instead offline the personality module's associated module. This offlines both the controller and the personality module in a safe and consistent manner.

Hot Pull Support

You can hot pull, hot insert, and hot integrate most modules on FX Series systems. To hot pull, hot insert, and hot integrate means that you can remove, insert, and integrate any of these modules without power cycling, rebooting, or otherwise stopping the system. We do however, recommend that you offline the module before you remove the module. Refer to *Hardware Installation* for more information. Do not hot pull modules during Split Mode.

CPU Module Subsystem

The following issues are related to the CPU module subsystem.

Set Time on New CPU Modules

If you bring the system down to replace any of the CPU modules with a new CPU module, make sure that the time on the clocks in the new CPU modules is set to the current time.

If a new CPU module is added to a running system, then we highly recommend that the Real Time Clock be initialized previously using some other chassis. There is a minor but potential risk that this problem can cause future problems even if the CPU module is replaced when AIX is running. Refer to “CPU Reintegration Has Temporary Critical Fault Under Heavy Load” on page 3-43 for more information.

To check and set the clock on a CPU module, follow these steps:

1. Stop the reboot in the bug if it is set to autoboot on power on.
2. At the bug prompt check that the time on the master CPU is correct:

```
FX-Bug>time
```

3. If the time is not correct, set it using the `set` command:

```
FX-Bug>set mmddyyhhmm
```

where the format for the time is “month month day day year year hour hour minute minute.” So that the time comes up correctly once the system boots, you must set the time in Greenwich Mean Time (GMT), not local time. For example, Mountain Standard Time (MST) is 7 hours earlier than GMT.

Once the system boots, make sure all the CPU modules are in sync so they all have their clocks set to the same time.

If you replace only one of the CPU modules while the system is down, make sure it is not the master during boot by leaving it unlatched. Reintegrate the module after boot and its clock will be set to the current time.

CPU Module Status after a Reboot

The AIX4.1.5 installation process is designed to automatically bring online and integrate all CPU modules found in the system. CPU modules retain their state across reboots of the system. CPU modules that are online and integrated before a reboot are automatically brought back online and integrated after the reboot and CPU modules which are offline before a reboot remain offline after the reboot, except in the case where booting manually was forced off of a previously off-lined CPU module.

The process of onlining and integrating all CPU modules will not complete until several minutes after the login prompt is displayed.

Do Not Use `ftctl` Program to Online and Offline CPU Modules

Use the SMIT interfaces or if desired the `modchange` command to online, offline, and acknowledge faults in the CPU modules. Do not use the `ftctl` program included in the FX Series to online and offline CPU modules. `ftctl` may leave your machine in an undesirable state that is difficult to recover from without rebooting. The `ftctl` command can still be used to view CPU module information (that is, `ftctl -status`).

Also, do not use the `ftctl -split`, `ftctl -unsplit`, and `ftctl -splitswitch` during Split Mode. These are low level commands that are only partially responsible for system transitions.

OS Watchdog

The OS watchdog timer is currently set to 1 second The operating system; however, refreshes the timer every .5 second.

ECC Parity Handling

An ECC hard error encountered on a simplex system results in a panic. A CPU module reporting an ECC hard error in a DMR or TMR system is taken offline and powered down.

MFIO Module

The following issues are related to the MFIO module.

Upgrading Disks from 2Gb to 4Gb

Overview

MFIO modules are now available which contain 4Gb hard drives. It is possible to upgrade existing MFIOs, which contain 2Gb hard drives, to the newer module without interrupting service.

The following procedures show two options for upgrading an existing MFIO to a new MFIO with a 4Gb hard drive without rebooting the system.

The key step is to replace the old MFIO definition in the ODM database with the new one for the MFIO containing the larger drive. In the first procedure below, this includes converting a fault tolerant volume group to a non-fault tolerant volume group, and later converting it back into a fault tolerant volume group. In the second procedure, the ODM attributes are changed directly from the command line. Because the first procedure may involve taking multiple MFIO modules off-line, it may take longer than the second. However, it may be a better procedure for those who are unfamiliar with ODM.

The general steps for both procedures are to:

- make sure all the functionality provided by the module is mirrored before you start
- take the module (or modules) offline
- remove the devices on the module from their logical devices
- remove the module from the database
- replace the module, and add it as a new module
- add the newly created devices to the logical devices

Procedure 1--Upgrading Disks Using a Volume Group Type Conversion

Note For this procedure to work, the mirrored volume group must be made simplex. This means that all mirrored drives in the same I/O domain and same volume group as the MFIO module you are upgrading must be offlined.

1. Make sure that /tmp has at least 10 MB of free space by running

```
df /tmp
```

This space is needed during step 11.

2. Gather disk information for the volume group by running

```
lsscsi -v NAMEvg
```

where *NAME* is a variable for the name you have given the volume group. Make sure that the volume group is properly mirrored as a fault tolerant volume group by verifying that the CMS state is “ft_full.”

3. Take the MFIO module offline

- Check to see what modules are included in the volume group by running

```
lsvg -p NAMEvg
```

For each *hdisk* in the list, run the following commands:

```
lsstate -s | grep hdisk
```

where *hdisk* is the name of the hard disk

```
lsstate -s | grep Loc | grep I/O
```

where *Loc* is taken from the location field in the previous output.

- Take the appropriate modules offline using SMIT.
smit cms => Modules => Manage Modules => Take Module Offline
Select appropriate MFIO module.

- Confirm that the module has been taken offline and that the volume group is “ft-partial” by running

```
lsscsi -v NAMEvg
```

It is essential that you do not proceed until `autolvfix` has completed and the `lsvg` command shows that all appropriate disks have been offlined. You can use “`ps -deaf`” to track the `autolvfix` processes.

4. Determine whether the module contains an ethernet device by running

a. `lsent`

b. `lsstate -s | grep pent`

where *pent* is the name of each device output by the `lsent` command

c. `lsstate -s | grep Loc | grep I/O`

where *loc* is the location field output by the previous command

5. If the module does contain an ethernet controller which is configured as part of a logical ethernet device, remove the physical device from the logical ethernet device using SMIT.

```
smit enet => Logical Ethernet Device => Change/Show a Logical Ethernet Device
```

Select the appropriate logical device and then delete the appropriate physical device.

6. Convert rootvg to non FTVG type using SMIT.

```
smit lvm => Volume Groups => Set Characteristics of a Volume Group => Change a Volume Group
```

Enter “*NAMEvg*” for “VOLUME GROUP name,” and change “Fault Tolerant Volume Group (FTVG)?” to no. (Use F4 to select options.)

7. Remove the MFIO from the ODM database using SMIT.

smit cms => Modules => Manage Modules => Remove Module from Database

Select appropriate MFIO module.

Then, run `lsmmod -Dl Moduleid` to verify that the module has been removed from ODM.

8. Physically replace the existing MFIO module with the new module containing the 4Gb hard drive.

9. Configure the new MFIO device using SMIT.

smit cms => Modules => Manage Modules => Add Module

Select MFIO1 as the template and specify the module location or Slotid, for example c1-f13.

Then, verify that the module has been integrated by running

```
lsdev -C | grep Slotid
```

where *Slotid* is the number of the slot containing the newly integrated module.

10. Add the hard drive to the root volume group using SMIT and the hard drive identifier (such as “hdisk1”) from step 9.

smit lvm => Volume Groups => Set Characteristics of a Volume Group => Add a Physical Volume to a Volume Group

11. Convert the volume group back to type FTVG using SMIT.

smit lvm => Volume Groups => Set Characteristics of a Volume Group => Change a Volume Group

Enter “*NAMEvg*” for “VOLUME GROUP name” and change “Fault Tolerant Volume Group (FTVG)?” to yes. (Use F4 to select options.)

Note It may take more than 30 minutes per disk to mirror data. If the console terminal is needed, do this step from a shell script running in the background.

12. If applicable, add the physical ethernet device to the logical ethernet device.

13. Verify the result by running

```
lsscsi -v NAMEvg
```

14. Repeat the process in order to upgrade the hard drive in the other I/O domain.

Procedure 2--Upgrade Disks by Changing ODM Attributes Directly

In this second procedure, it is not necessary to take all the disks in one side of the volume group out of the FTVG. Therefore, this procedure may be faster, because it is not necessary to reestablish mirroring on multiple disks.

1-5. The first five steps are identical to those in procedure 1.

6. Fix the logical volume dependency on the disk or disks you are upgrading (hdiskx, where x stands for the appropriate number) in ODM.

```
> odmget -q "dependency=hdiskx" CuDep
```

CuDep:

```
name = "l_NAMEvg" dependency = "hdiskx"
```

```
> odmdelete -q "dependency=hdisk1" -o CuDep
```

Note Step number 6 needs to be performed on all hard disks on the MFIO module which you are removing. For example, if I/O-1 contains hdisk1 and hdisk2, then you need to perform the procedure on both hdisk1 and hdisk2.

7-9. Steps 7, 8, and 9 are identical to those in procedure 1.

10. Add the new disk(s) to the Volume Group.

smit lvm => Set Characteristics of a Volume Group => Add a Physical Volume to a Volume Group

Add the appropriate hard disk to *NAMEvg*.

11. Resync the Volume Group by taking it offline and then bringing it back online.

smit cms => Modules => Manage Modules => Take Module Offline

Select appropriate module.

smit cms => Modules => Manage Modules => Bring Module
Online

Select appropriate module.

Note This final step may take more than 30 minutes. If the console terminal is needed, do this step from a shell script running in the background.

12-14. Steps 12, 13, and 14 are identical to those in procedure 1.

Interconnect Module (ICM)

The following issues are related to the Interconnect module.

Any Interconnect Module Failure in a Domain Takes Down the Entire Domain

In systems with expansion chassis, the failure of an ICM module in either the main or expansion chassis takes down that entire IO domain. Because the IOBus does not function properly if any of the ICM modules providing power to terminate the bus fails, the system powers off all the modules in that domain.

The main chassis ICM module is listed as the module at fault even if the expansion chassis ICM module was the one that failed. Look at the failure reports in the system status log,

`/var/ft/changed_log`, to determine which module actually failed.

CPU Modules May Go Offline When an Interconnect Module Fails

If an ICM module fails due to power loss, the loss of power usually sends all significant CPU modules into E-state. The hardware must do this to prevent a possible IOBus lockup when the power loss occurs. If the system is running with:

- two CPU modules one of them reports itself offline
- three CPU modules then either:
 - one module is out-of-sync (to be automatically onlined) and one module is offline OR
 - two modules are offline

If CPU modules fail when an ICM module fails, they may be reintegrated immediately, without waiting for the ICM module to be replaced or reconnected to the power. We recommend first bringing the CPU modules back online before reintegrating the

ICM module. CPU modules take less time to integrate than the disk mirroring operations that start when the ICM module and its dependent modules in that domain are reintegrated.

Do Not Remove the Interconnect Module without Disconnecting the Power Cord

ICMs can be pulled out with the system still running (that is hot pulled), but you must remove the power cord first. The ICM modules are the only hardware modules whose latches do not disconnect the modules from the power.



Caution

Pulling the module (or putting it back in) with the power connected causes arcing on the voltage pins. Repetition of this would eventually cause pitting of the contacts, leading to unreliable operation.

Power Module and Fantrays

The following issues are related to the Power module and Fantrays.

Power Cycle the System Using Power Module Latches

The FX Series system has no on/off switch. If the system is powered using Power modules, their latches serve as the means to power cycle the system. After the system has shut itself down and powered itself off, restart it by opening and closing ALL power module latches.

The Power Daemon Takes Time to Complete Its Shutdown

When the power daemon (`powerd`) is stopped using `stopsrc -g ft` or `stopsrc -s powerd`, the SRC command returns immediately after telling the daemon to stop. The daemon process takes 10-15 seconds to shut down completely.

If you try to start the daemon immediately using `startsrc -g ft` or `startsrc -s powerd`, you see an error message:

```
Can't start multiple instances of powerd
```

Wait 10-15 seconds, or use `lssrc -s powerd` to verify that `powerd` has terminated before trying to restart the daemon.

Do Not Boot System with the Fantrays Removed

The FX Series system shuts itself down if both fantrays in a main or expansion chassis are removed. Do not reboot the system without replacing the fans. The system will shut itself down again after booting without the fantrays, however the system can be damaged by operating at too high of a temperature during the boot.

Boot System Carefully after Failure of All Power Modules

The FX Series system shuts itself down if all the power modules in a power chassis fail. Because it is unlikely for all modules to fail at the same time, you may wish to reboot the system to check out the state of the modules before replacing them.

Normally the power module batteries stay charged during and after system shutdown. If there is any chance they may not be fully charged, allow the power modules to recharge for 40 minutes before booting the system. The system brings all the power modules online and checks to see if they are working properly. If they are not working properly, the system should shut itself down again.

You must have at least one fully functional power module for the system to operate safely. The system may be able to boot with partially functioning power, but it is not safe to operate the system that way. Check the power modules by running `pmftstat`. The output looks like the following example:

```
>>>> Power Module (POWER-0 p1-f1 Available) [ 5.3 ] <<<<
Module LED: Inservice 'ON' Outofservice 'OFF'
Alarms: Rack Alarm 'OFF'
Analog Switches: CurrA 'CLOSED' CurrB 'CLOSED'
Battery Load 'OFF' Battery Charger 'ON' Convertor 'ON'
Digital Inputs: AC 'OK' Temp 'OK' CurrA 'OK' CurrB 'OK'
Temp '48' C CurrA '4758' mA CurrB '4524' mA
Volts '55' V Batt '52' V
```

Make sure the following is true in your output:

- The two Analog Switches must be CLOSED.
- The Digital Inputs must all show OK.
- The Temp must be below 80C.
- The Volts must be greater than 52 V.

Boot System Carefully after Failure of All Fantrays

The FX Series system shuts itself down if both fantrays in a main or expansion chassis fail due to low fan speed or inability to maintain a low enough system temperature. Because it is unlikely for both fantrays to have a mechanical failure at the same time, you may wish to reboot the system to check out the state of the modules before replacing them. Make sure any overheating in areas in which the system operates is corrected first.

You must have one fantray working without fault for the system to run safely. The system should shut itself down after boot if both fantrays still have faults. If it does not, and the fans are not operating properly, shut the system down immediately and replace the fantrays.

Check the fan speeds and temperatures by running `pmftstat`. The output looks like the following:

```
>>>> Fan Tray Module (FANS-0 c1-t0 Available) [ 5.4 ] <<<< Module
LED: InService 'ON' OutofService 'OFF' System LED: InService 'ON'
OutofService 'OFF' Telco LED: Critical 'OFF' Major 'OFF' Minor
'OFF' Fan Speed: Fan#1 '2119' RPM Fan#2 '2119' RPM Fan#3 '2200' RPM
Fan Temp: '33' C
```

Normal fan speeds are in the range 1000-3300 RPM. The temperature must be below 60 C.

Known Problems on FX Systems

Overview

This section contains information on known problems with FX systems in this release of AIX 4.1. Workarounds have been provided whenever they are available.

Creating Swap Partitions While Volume Groups are Nonftvg Interferes with Disk Mirroring

Tracking Number

TAR 1-10199

Problem

If you use the SMIT ASSIST menu to extend swap partitions on a nonftvg disk, the number of physical volumes associated with the corresponding logical volumes will be restricted to one. Thus, you will not be able to mirror the volume groups, which involves spanning the mirror copies across two physical volumes.

This situation occurs if you extend swap partitions on SYSNEW when it is the Split Mode secondary, because the root volume groups are not in the ftvg state while the system is split. In the case of Split Mode, this problem will result in a failure to re-mirror the root volume groups at the end of the Split Mode process.

Workaround

Manually change the value of the "MAXIMUM NUMBER of PHYSICAL VOLUMES" from 1 to 2 for the newly created paging logical volumes.

The following shows how to change the maximum number of physical volumes from the command line, using a volume group called paging00 as an example. Alternately, you can use the SMIT->System Management (Physical & Logical Storage)->Logical Volume Manager->Logical Volumes screen to identify and modify the volume groups.

Find the paging logical volumes (this assumes that you have placed your paging logical volumes in order to steal data volume groups):

```
# lsvg -l rootvg
```

Change the maximum number of physical volumes:

```
# chlv -u 2 paging00
```

Check that the changes are correct:

```
# lslv paging00
```

Repeat with any other volume groups where applicable.

Note During Split Mode, these changes should be done before transitioning forward from the SPLIT state. Otherwise, the transition from UNSPLIT to FT_COMPLETED will fail, and the system will remain in the UNSPLIT state. It is possible to fix the logical volumes at this point and then transition to forward to FT_COMPLETED.

Time Zone not Set on SYSNEW when using Split Mode to do a BOS-NEW Installation

Tracking Number

TAR 1-10205

Problem

When a BOS-NEW install is done on SYSNEW, the time zone is reset to the default CST time zone. While the time zone is not set, the time itself is set during the Split Mode Transition from SIMPLEX to SPLIT.

Workaround

Use the `chtz` command to set the appropriate time zone.

Loading and Unloading Kernel Extensions Leaks Memory

Tracking Number

TAR 1-10213

Problem

Kernel extensions, such as the pridd, leak a memory page when they are unloaded.

Workaround

Limit loads and unloads of kernel extensions.

Cannot Create split-notify.%-%.tmp File without Write Access to Current Directory

Tracking Number

TAR 1-10255

Problem

You cannot execute the `spstart` command in a directory to which you do not have write access, because it writes a temporary file to the current directory.

Workaround

Run the command from a local directory to which you have write permissions.

Message Indicates a Failure to Mirror Rootvg when doing a Preservation Install

Tracking Number

TAR 1-10228

Problem

When doing a preservation installation in FT mode, there is a message stating that “the root volume group will not be mirrored in opposite I/O domains. Rootvg is not fault tolerant.”

Workaround

None needed. The message may be safely ignored, because while the OS is only being installed on a single disk, it is being duplicated on the mirrored disk by the logical volume manager.

System may Acquire Wrong Status for Asynch and Ethernet Personality Modules at Reboot

Tracking Number

TAR 1-10247 and TAR 1-10262

Problem

When the system is booted, the change daemon gets status directly from devices in order to determine which are offline and which are online.

This is not the case, however, with Ethernet and Asynch devices. Because of the frequency of cable or network problems, the Ethernet and Asynch subsystems manage faults by offlining the appropriate personality module. This architecture allows them to monitor the controller and to bring the link back online when the connection is reestablished. When the system is booted, the change daemon uses the last known status of the personality modules from before the reboot.

This algorithm can cause discrepancies in those situations in which the status of the link changes while the system is being rebooted. If, for example, an Ethernet cable is dropped before the system is rebooted, and is then reattached while the system is booting, it will be considered offline, even though the link is good.

Workaround

Contact your service representative to see when a PTF will be available.

fxnvctl does not Remove Lock File when Interrupted.

Tracking Number

TAR 1-11803

Problem

The `fxnvctl` utility does not remove its lock file when interrupted. The lock will interfere with the use of `restart` and `autoboot`.

Workaround

Allow `fxnvctl` to run to completion. If the utility is interrupted, remove the lock file (`/tmp/fxnvctl.lock`).

File Systems Named “jfslog” are not Included in File System Listings

Tracking Number

TAR 1-11804

Problem

If a file system is named “jfslog” it will not be included in the list of file systems created by the `splitstate` utility.

Workaround

Do not use jfslog as a file system name.

Cron Jobs may Generate Unnecessary E-mail against rotlog

Tracking Number

TAR 1-11866

Problem

Extraneous e-mail messages, containing rot.log information, may be generated by cron jobs.

Workaround

Redirect the output from the “rotlog” cron job to `/dev/null`.

Secondary May Hang During Boot Cycle

Tracking Number

TAR 1-11921

Problem

On rare occasions, mbus traffic may cause the secondary system to hang while booting in the SPLIT state.

Workaround

Transition the system backwards to the SIMPLEX state and transition forward to the SPLIT state again. Reboot the secondary system.

Note This workaround is an exception to the rule noted on 3-41. Before making the state transition, wait long enough to ensure that the boot is hung.

Possible Bootlist Error if Rootvg includes Multiple Hdisks

Tracking Number

TAR 1-11923

Problem

At the beginning of the Split Mode process, the boot list is saved. It is then restored before autoboot is reenabled at the end of the process. It is possible, however, that if the root volume group contains multiple hard disks in each I/O domain, the boot partition will be on a different disk following an OS installation than at the beginning of Split Mode.

Workaround:

Once you have completed the Split Mode procedure, identify which disk contains the boot partition:

```
lslv -l hd5
```

Check the bootlist to ensure that the appropriate hdisk is listed:

```
bootlist -m normal -r
```

If the list does not contain the appropriate hdisk, create a new bootlist:

```
bootlist -m normal devicex
```

where *devicex* is the device, or a list of devices in order of preference, that you want to boot from.

Hours in Service in Module EEPROMs may be Counted Twice

Tracking Number

TAR 1-12765

Problem

History logs are kept in module EEPROMs for FX Series modules. The hours-in-service field for each online module is updated every 25 hours by an **at** job. Occasionally, when the system is rebooted, two or more **at** jobs are started. This results in the modules being credited with double (or more) the correct in-service time.

Workaround

After the system is rebooted, check that only a single **at** job is pending to run the “`/usr/sbin/updhours -h 25`” command. You can check the list of pending jobs using the **atq** command. It provides a list of names of files in the `/var/spool/cron/atjobs` directory. If more than one of files contains the **updhours** command, you can stop the redundant jobs using the following syntax:

```
$at -r <JobName>
```

CPU Reintegration Has Temporary Critical Fault Under Heavy Load

Tracking Number

TAR 1-48836

Problem

During CPU module integration the critical Telco alarm may come on. This occurs during the final CPU integration phase or the “critical integration” time. If the system is running an application mixture that results in modifications to large numbers of DRAM pages faster than the system can copy them to the target CPU module, then the critical alarm may fire. This is a temporary alarm that the system automatically clears after the CPU module integration completes.

CPU module integration happens at boot time or as the net result of fault condition(s) that occur during system operations that eventually result in operator actions to re-integrate the CPU module(s).

Workaround

Perform CPU module integration when the system is less busy. There is no workaround if the CPU module must be integrated during a period when the application mixture is such that a large number of memory pages are modified more quickly than the system can copy them to the target CPU Module.

Making Split Mode State Changes while Booting Secondary System can Corrupt the Secondary Boot Record

Tracking Number

TAR 1-83297

Problem

Changing Split Mode states while the secondary system is booting, can cause occasional boot record corruption on the secondary. This occurs when the `savebase` utility, run automatically by AIX during system startup, is interrupted by state-change activities. Problems have also been observed when changing states during firmware flash updates. To avoid problems related to these issues, do not make state changes using `splitstate`: (1) During a secondary system boot, until the secondary system has completely booted to a logon prompt, or (2) While firmware flash updates are in progress.

Workaround

To avoid problems related to these issues, do not make state changes during a secondary system boot, until the secondary system has completely booted to a logon prompt, or while firmware flash updates are in progress.

Note If the secondary system hangs for some reason during a boot cycle, see “Secondary May Hang During Boot Cycle” on page 3-40.

Upgrading from AIX 4.1.4r6 to AIX 4.1.5 Causes Problems with ASYNC Logical Devices

Problem

Due to modifications in the structure of the ASYNC logical device between AIX 4.1.4R6 and AIX 4.1.5, ASYNC logical devices will be "inaccessible" after an OS upgrade.

Workaround

It is necessary to redefine the logical devices by re-entering the member0 and member1 information. You may reenter the information from the command line or by using SMIT.

The following steps are appropriate for a logical device lfxas0, with devices fxas0 and fxas1. If your devices have different names, you should substitute them accordingly.

1. Identify which ASYNC devices belong to which ASYNCH logical devices:

```
lsattr -E -l lfxas0
```

2. While logged in as root, redefine the devices

```
chdev -a member0=fxas0 -a member1=fxas1 -l lfxas0
```

At the end of the procedure, the logical device will be "available" and "on_line_both" (provided that both members are "available" and "on_line", and that the connection between the ASYNCPM and the patch panel is working).

Documentation Notes 4

Overview

This chapter documents changes and corrections to various books and InfoExplorer documents in the AIX 4.1 documentation set.

FX Series Administering Your Fault Tolerant System

The discussion in “How CMS Uses the AIX Error Logging Facility with the Change Daemon” on page 2-15 of this guide does not apply to either AIX 4.1.5r2 or to AIX 4.1.5r3. The AIX error logging facility is no longer used as the communication path between the device drivers and the change daemon. Instead, messages are now logged to and retrieved from a circular buffer in the kernel.

Installing FX manpages

On page 5-13 of the *Operating System Installation Guide*, there is a list of filesets that need to be manually installed. That list is no longer accurate. The appropriate filesets for the FX Series manpages are:

- devices.man.en_US.sys.ft.data
- devices.man.en_US.sys.iobus.data
- devices.man.en_US.sys.mbus.data
- devices.man.en_US.sys.fxnvrn.data
- devices.man.en_US.xrft.lvm.data
- devices.man.en_US.xrft.base.data
- devices.man.en_US.xrft.cms.rte.data
- devices.man.en_US.xrft.cms.presaddr.adt.data
- devices.man.en_US.xrft.flash.data
- devices.man.en_US.xrft.ondd.data
- devices.man.en_US.xrft.pmft.rte.data
- devices.man.en_US.xrft.iscs.adt.data
- devices.man.en_US.xrft.iscs.rte.data
- devices.man.en_US.xrft.smmgt.rte.data

4

In order to install the filesets:

- Choose
->Software Installation
From the main SMIT menu.
- Next, choose:
->Custom Install
- Finally, enter:

devices.man.en_US
for the fileset.

4

AIX® Version 4.1.5r3

Release Notes

FXRNA/IS3

Release Notes

102 - 180 pages
3/8" - 1/2" spine