Stratus® ftServer® System Administrator's Guide for the Linux® Operating System

Stratus Technologies
R003L-05a

Notice

The information contained in this document is subject to change without notice.

UNLESS EXPRESSLY SET FORTH IN A WRITTEN AGREEMENT SIGNED BY AN AUTHORIZED REPRESENTATIVE OF STRATUS TECHNOLOGIES, STRATUS MAKES NO WARRANTY OR REPRESENTATION OF ANY KIND WITH RESPECT TO THE INFORMATION CONTAINED HEREIN, INCLUDING WARRANTY OF MERCHANTABILITY AND FITNESS FOR A PURPOSE. Stratus Technologies assumes no responsibility or obligation of any kind for any errors contained herein or in connection with the furnishing, performance, or use of this document.

Software described in Stratus documents (a) is the property of Stratus Technologies Bermuda, Ltd. or the third party, (b) is furnished only under license, and (c) may be copied or used only as expressly permitted under the terms of the license.

Stratus documentation describes all supported features of the user interfaces and the application programming interfaces (API) developed by Stratus. Any undocumented features of these interfaces are intended solely for use by Stratus personnel and are subject to change without warning.

Stratus, the Stratus logo, ftServer, Continuum, Continuous Processing, StrataLINK, and StrataNET are registered trademarks of Stratus Technologies Bermuda, Ltd.

The Stratus Technologies logo, the ftServer logo, Stratus 24 x 7 with design, The World's Most Reliable Servers, The World's Most Reliable Server Technologies, ActiveService, ftGateway, ftMemory, ftMessaging, ftStorage, Selectable Availability, XA/R, SQL/2000, ftScalable, and The Availability Company are trademarks of Stratus Technologies Bermuda, Ltd.

The registered trademark Linux is used pursuant to a sublicense from the Linux Mark Institute, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis. NEBS is a trademark of Telcordia Technologies, Inc.

Red Hat, Red Hat Linux, and Red Hat Enterprise Linux are registered trademarks of Red Hat, Inc. in the United States and other countries. UNIX is a registered trademark of X/Open Company, Ltd., in the U.S.A. and other countries.

All other trademarks are the property of their respective owners.

Manual Name: *Stratus ftServer System Administrator's Guide for the Linux Operating System*
Part Number: R003L
Revision Number: 05a
Software Release Number: ftServer System Software for the Linux Operating System, Release 4.0.1
Publication Date: June 2007

Stratus Technologies, Inc.
111 Powdermill Road
Maynard, Massachusetts 01754-3409

Contents

Preface xi

1. Introduction to ftServer System Administration 1-1
- ftServer System Terminology 1-2
- System and Network Administration Overview 1-2
  - Installing and Updating Software on ftServer Systems 1-2
  - Updating ftServer System Firmware 1-2
  - Configuring Your ftServer System 1-3
  - Managing Data Storage Devices 1-3
  - Using the Stratus Fault-Tolerant ftSSS 1-3
  - Network Management and Reporting 1-3
  - Troubleshooting ftServer Systems 1-3
- Additional Documentation and Resources 1-3
  - Red Hat Enterprise Linux 1-4
  - Stratus ftServer System Documentation 1-4
  - Linux and UNIX Documents 1-5

2. Installing the Operating System and ftServer System Software 2-1
- Installation Overview 2-2
  - Boot Media 2-3
  - Default System Setup 2-3
    - Linux Version Information 2-4
    - Storage Default Settings 2-4
    - Ethernet Default Configuration 2-5
    - USB and RS232C Default Settings 2-5
    - Default System Initialization and Run-Level Control 2-6
    - Default User Environments, Shells, and Access Control 2-6
  - System Indicators and Switches 2-6
- Separately Released and Optional Distribution Components 2-6
- Installation Interfaces 2-7
- Supported Hardware and Firmware 2-7
- Pre-Installation Checklist 2-7

Installing the Linux Operating System and ftSSS 2-10
Booting the Operating System 2-10
Installing the Operating System 2-12
Installing ftSSS for Fault Tolerance 2-13
Reinstalling ftSSS After a Failed Installation 2-15
Booting in Linux Rescue Mode 2-15
Post-Installation Tasks and Considerations 2-16
Performing an Installation Without a Kickstart File 2-17
Additional Documentation and Resources 2-19

3. Updating ftServer System Firmware 3-1
Updating the System BIOS 3-1
Updating BMC Firmware 3-5

4. Updating the Operating System and ftServer System Software 4-1
General Upgrade Considerations 4-2
Upgrading or Restoring the Linux Operating System 4-3
Stratus Kernel Modules 4-4
Upgrading the Linux Operating System 4-5
Restoring the Linux Operating System 4-6
Upgrading or Restoring Stratus ftSSS 4-7
Creating a Backup System Disk 4-9
Recovering from a Failed Software Upgrade 4-9
Related Information and Resources 4-10

5. Setting Up the ftServer System 5-1
Setting Up Internal Disk Storage 5-2
Internal Disk Storage Overview 5-2
The Console Log and the `/var/log/messages` File 5-2
Configuring Internal Disks 5-3
Managing Partitions 5-3
Default Internal Disk Configuration for a Newly Installed System 5-6
Checking the Current State of the Internal Disk Subsystem 5-7
Storage Device Definition 5-8
Setting Up RAID Arrays on Internal Disks 5-8
RAID Array Overview 5-8
Creating a RAID-1 Array 5-9
Creating a RAID-0 Array 5-11
Creating and Mounting a File System 5-13
Checking the Current State of RAID 5-13

Removing and Replacing Internal Disks 5-14
Disk Insertion 5-15
Administering RAID Arrays on Internal Disks 5-15
To Stop a RAID Array and Move It to Another System 5-15
Errors and Faulty Mirrors 5-16
Removing a Faulty Mirror 5-16
Resynchronization 5-17
Replacing a Failed Disk 5-18
Configuring Safe Mode 5-19
Manually Creating Partitions on Blank Disks and Adding to RAID-1 Arrays 5-19
Replacing Defective Disks Interactively 5-20
Replacing Defective Disks Manually 5-20
The `duplex_blank_disk` Command 5-24
Setting Up External ftScalable Storage 5-25
System Backup and Disaster Recovery 5-25
Ethernet Devices 5-25
Physical Device Naming 5-25
Monitoring and Configuring Channel-Bonding Interfaces 5-27
Monitoring Channel-Bonding Interfaces 5-27
Configuring Channel-Bonding Interfaces 5-28
Determining Interface Device Names 5-29
MAC Addresses 5-30
Other System Configuration Information 5-31
Disabling Hyperthreading 5-31
Configuring the System Video Display 5-32
Managing the System Clock 5-33
Additional Documentation and Resources 5-33

6. Managing Data Storage Devices 6-1
CD-ROM Drives 6-1
SCSI Tape Drives 6-2
USB Storage Devices 6-2
USB Floppy Drives 6-4
USB Solid-State Devices 6-4
Additional Resources 6-5

7. Using ftServer Fault-Tolerant Utilities and Software 7-1
The `ftsmaint` Command 7-1
Device Path Enumeration 7-5
ftServer System Device Path Enumeration 7-6
`ftsmaint` Examples 7-11
Displaying System Status 7-11
Bringing System Components Down and Up 7-14
Removing a PCI Adapter From Service and Bringing It Into Service 7-14
ActiveService Network Support 7-14
Kernel Memory Dump File Management 7-16

8. Simple Network Management Using Net-SNMP and ftlSNMP 8-1
Installing and Configuring ftlSNMP 8-1
ftlSNMP Inventory 8-2
Manually Installing and Upgrading the ftlSNMP RPM 8-3
ftlSNMP Prerequisites 8-4
SNMP Configuration File Updates 8-4
The `snmpd.conf` File 8-5
The `ftlsubagent.conf` and `ftltrapsubagent.conf` Files 8-5
Configuring SNMP to Start at System Initialization 8-6
Configuring SNMP for Service Management 8-6
SNMP Foundations and Concepts 8-8
ftlSNMP Management Commands 8-8
The Basic Net-SNMP Commands 8-9
MIBs 8-11
Some Objects Defined by Standard MIBs 8-12
SNMPv3 Support 8-13
SNMP's View of a Network 8-14
Extensions and Fault-Tolerant SNMP Operation 8-15
Installing Remote Network Management Services 8-16
Configuring SNMP for Remote Service Management 8-16
Deploying SNMP Agents and Subagents 8-17
Verifying Traps 8-18
Managing SNMP 8-18
Testing Your SNMP Configuration 8-19
Managing ftServer Hardware Components 8-20
Example: Managing Hardware 8-21
Testing Ethernet Ports 8-23
Example: Testing Ethernet Ports 8-23

SNMP and MIBS 8-24
Device Enumeration 8-25
ftServer System Operation State Management 8-25
SNMP Network Management Station Considerations 8-26
Initial SNMP Testing 8-27
Initial Testing of `ftltrapsubagent` 8-28
Initial Testing of `ftlsubagent` 8-29
Removing ftlSNMP 8-30
OpState:State Definitions 8-30
OpState:Reason Definitions 8-32
GET and SET Operations for ftlSNMP MIB Objects 8-33
SRA-ftLinux-MIB OID Values and Properties 8-33
Trap Filtering 8-33
Trap-Filtering Capability 8-34
Activating and Deactivating Trap Filtering 8-34
Trap-Filtering Examples 8-35

9. Troubleshooting ftServer Systems 9-1
LED and Visual Diagnostics 9-1
System Boot Problems 9-1
Normal Boot Sequence 9-2
Possible Boot Problems 9-3
Missing Stratus Drivers Prevent Booting 9-3
GRUB Problem 9-4
RAID Problem 9-5
Automatic Reboot After Boot Monitoring Timeout 9-5
System Log Messages 9-6
Error and Log Messages Regarding Keyboard and Mouse 9-6

Appendix A. Linux Packages A-1

Index Index-1

Figures

Figure 2-1.	SATA Drive Arrangement for Installation	2-9
Figure 5-1.	CPU-I/O Enclosures: Front Panel with Drive Slots Fully Populated	5-3
Figure 7-1.	ftServer Enclosures: Locations of Major Enumerated Devices (Front View)	7-9
Figure 7-2.	ftServer Enclosures: Locations of Major Enumerated Devices (Rear View)	7-10
Figure 8-1.	AgentX-Enabled Extensions and Subagents	8-17
Figure 8-2.	Operational State Management on ftServer Systems	8-25

Tables

Table 1-1.	ftServer System-Specific Documentation	1-4
Table 2-1.	CD-ROMs Included With ftServer Systems	2-2
Table 5-1.	Default Internal Storage Allocation	5-6
Table 5-2.	Ethernet Devices in ftServer CPU-I/O Enclosures	5-26
Table 7-1.	Device Paths of ftServer Devices	7-6
Table 8-1.	Meaning of *Duplex* for ftServer System Components	8-26
Table 8-2.	Operation State Values, Names, and Definitions	8-30
Table 8-3.	Reason Codes, Names, and Definitions	8-32
Table 8-4.	Set Operations Currently Implemented in ftlSNMP	8-33

Examples

Example 5-1.	Checking the Current State of the Internal Storage Subsystem	5-8
Example 5-2.	Checking the Current State of RAID	5-14
Example 5-3.	Resynchronization	5-17
Example 5-4.	Running GRUB	5-23
Example 5-5.	Pairing a Spare Internal Disk with the Running System Disk	5-24
Example 5-6.	Default Configuration of Embedded Ethernet Devices	5-28
Example 7-1.	Displaying System Status with the `ftsmaint` Command	7-11
Example 8-1.	Traps that Can Occur for I/O Element 11 When Trap Filtering Is Off	8-35
Example 8-2.	Traps That Can Occur for I/O Element 11 When Trap Filtering Is On	8-39
Example 8-3.	Traps That Can Occur When Trap Filtering Is Off	8-41
Example 8-4.	Traps That Can Occur When Trap Filtering Is On	8-42
Example 9-1.	Possible Keyboard and Mouse Error Messages at Boot Time	9-6
Example A-1.	Displaying Information About a Stratus-Proprietary RPM	A-2
Example A-2.	Displaying Information About an Open-Source RPM	A-3

Preface

The *Stratus ftServer System Administrator's Guide for the Linux Operating System* (R003L) documents tasks and information for system administrators of Stratus systems running a supported Linux distribution and ftServer System Software for the Linux Operating System (ftSSS).

This document is intended for system and network administrators using or migrating to Stratus systems running a supported Linux distribution and ftSSS, and for system and application programmers who develop tools and scripts for use on these systems. Background knowledge of Linux or UNIX® shells, tools, and systems, and Linux or UNIX and TCP/IP network server and network administration technologies is assumed.

Revision Information

This document is a revision.

This revision incorporates the following changes:

- It documents the procedure for upgrading the system to Redhat Linux Release 4, Update 5.

The previous revision of this document incorporated the following changes:

- It documented additional installation steps for systems with the optional ftScalable™ Storage array.
- It documented new device IDs for the trays in an ftScalable Storage array.
- It documented new CD and DVD device behavior.
- It corrected CPU- and I/O-element IDs in Table 5-2.

Notation Conventions

This document uses the notation conventions described in this section.

Warnings, Cautions, and Notes

Warnings, cautions, and notes provide special information and have the following meanings:

WARNING

A warning indicates a situation where failure to take or avoid a specified action could cause bodily harm or loss of life.

CAUTION

A caution indicates a situation where failure to take or avoid a specified action could damage a hardware device, program, system, or data.

NOTE

A note provides important information about the operation of an ftServer system.

Typographical Conventions

The following typographical conventions are used in this document:

- The italic font introduces or defines new terms. For example:

  ftServer systems use replicated, *fault-tolerant* hardware to eliminate single points of failure and protect data integrity.

- The bold font emphasizes words in text. For example:

  Update the BIOS **before** you install or upgrade ftSSS.

- The monospace font represents text that would appear on your display screen. The monospace bold font represents text you must type in examples that contain both user input and system output. The monospace italic font represents terms in command lines that are to be replaced by literal values. For example:

  To display the state of a CPU enclosure, type a command in the following format:

  ```
  /opt/ft/bin/ftsmaint ls n
  ```

  If you type **`/opt/ft/bin/ftsmaint ls 0`** at the prompt, the following output appears:

  ```
  H/W Path : 0
  Description : CPU Node Assembly
  .
  .
  .
  ```

- The percent sign (%) and the number sign (#) are standard default prompt signs that have a specific meaning at a command prompt. Although a prompt is sometimes shown at the beginning of a command line as it would appear on the screen, you do not type it.
  - % indicates you are logged in to a user account and are subject to certain access limitations.
  - # indicates you are logged in to the system administrator account and have *superuser* access. Users of this account are referred to as `root`. The # prompt sign used in an example indicates the command can only be issued by `root`.

Syntax Notation

This document uses the following format conventions for documenting commands:

- Square brackets ([ ]) enclose command argument choices that are optional. For example:

  `cflow [-r] [-ix] [-i] [-d` *`num`*`] `*`files`*

- The vertical bar (|) separates mutually exclusive arguments from which you choose one. For example, the following shows two mutually exclusive, but optional, arguments:

  `command [`*`arg1`*` | `*`arg2`*`]`

  The following example shows two mutually exclusive arguments, one of which is required:

  `command `*`arg1`*` | `*`arg2`*

  In either case, you may use either `arg1` or `arg2` when you type the command.

- Ellipsis (...) indicates that you can specify the preceding argument as many times as you need to on a single command line. For example,

  `command [`*`arg1 arg2 arg3`*` ...]`

  NOTE

  Dots, brackets, and braces are not literal characters; you should **not** type them. Any list or set of arguments can contain more than two elements. Brackets and braces are sometimes nested.

Getting Help

Stratus provides complimentary access to StrataDOC, an online-documentation service that enables you to view, search, download, and print customer documentation. You can access StrataDOC at the following Web site:

http://stratadoc.stratus.com

If you have a technical question, you can find the latest technical information at the Stratus Technical Support Web site:

http://www.stratus.com/support/technics.htm

If you are unable to resolve your questions with the help available at this online site, you can contact the Stratus Customer Assistance Center (CAC) or your authorized Stratus service representative. For information about how to contact the CAC, see the following Web site:

http://www.stratus.com/support/cac

Chapter 1
Introduction to ftServer System Administration

This chapter discusses the following topics:

- "ftServer System Terminology"
- "System and Network Administration Overview"
- "Additional Documentation and Resources"

ftServer systems running a supported Linux distribution together with ftServer System Software for the Linux Operating System (ftSSS) operate as fault-tolerant servers. The supported server models are the Stratus ftServer T40, T65, 2400, 4300, and 5700 systems. Every ftServer system uses replicated, fault-tolerant hardware to eliminate single points of failure and protect data integrity in all areas of data handling, including:

- **Processing.** Replicated, fault-tolerant processing components process the same instructions at the same time. In the event of a component malfunction, the partner component acts as an active spare that continues normal operation, preventing system downtime and data loss.
- **Internal Storage.** Mirrored RAID 1 (*Redundant arrays of inexpensive disks*) arrays of Serial Advanced Technology Attachment (SATA) disks prevent single disk failures from causing data loss. Replacement disks are automatically recognized and mirrored.
- **Networking.** Duplexed network components maintain network connectivity. When the operating system detects a malfunction in the primary member of a duplexed pair, it automatically fails over to the secondary member.

ActiveService architecture built into ftServer systems supports these features with self-checking hardware and onboard diagnostics to detect, isolate, and report potential problems before they affect server operation, offering complete hardware diagnostics and alarms.

ftServer System Terminology

Each ftServer system houses two *CPU-I/O enclosures*. Each CPU-I/O enclosure includes a *CPU element* and an *I/O element*, as follows:

- CPU element 0 and I/O element 10: The top enclosure, also referred to as CPU-0, I/O-10.
- CPU element 1 and I/O element 11: The bottom enclosure, also referred to as CPU-1, I/O-11.

System and Network Administration Overview

Most examples and discussions in this guide assume that you are acting with root-user or superuser privileges. They do not always specify when you should (or should not) be acting as root.

Installing and Updating Software on ftServer Systems

If you ordered the operating system from Stratus, the supported Linux distribution and ftSSS were installed at the factory. You also received distribution CDs containing ftSSS and the operating system. These CDs are provided so that you can, if necessary, reinstall your ftServer system's software.

If you ordered a supported Linux distribution from a vendor other than Stratus, the vendor can also provide you with distribution CDs for installing or reinstalling the system software. You need to use the ftSSS CD with the CDs for the supported Linux distribution to achieve a fault-tolerant ftServer system.

Before you install your ftServer system's software, you must prepare your ftServer system by following the hardware installation instructions provided in the installation guide for your system.

See Chapter 2 and Chapter 4 for software installation and update procedures.

Updating ftServer System Firmware

Specific firmware versions are required for a given release of ftSSS. The firmware in an ftServer system delivered from the factory does not require updating. However, subsequent updates to ftSSS may require firmware updates. When updating ftSSS, use the versions of firmware and software that are supplied on that ftServer System Software for the Linux Operating System CD. Chapter 3 provides procedures for updating your system BIOS and Stratus Baseboard Management Controller (BMC) firmware.

Configuring Your ftServer System

After installing the Linux operating system and ftSSS, you must configure your system. See Chapter 5 for configuration information.

Managing Data Storage Devices

In addition to the SATA disk storage discussed in Chapter 5, your system supports CD-ROM drives, tape drives, and USB storage devices. Chapter 6 provides a discussion of these devices and the information needed to manage them.

Using the Stratus Fault-Tolerant ftSSS

While you can use standard Linux tools to perform many system administration tasks on your ftServer systems, some tasks on fault-tolerant systems require specialized supporting utilities. The ftSSS includes libraries and utilities to support fault-tolerant system administration tasks. Chapter 7 discusses management tasks and utilities to manage fault-tolerant features and supporting applications of your system. They include the ActiveService Network (ASN) package that provides support for ASN access using an attached modem. This package allows the Stratus Customer Assistance Center (CAC) or your authorized Stratus service representative to provide remote support for your system. Your system comes with ASN installed.

Network Management and Reporting

ftSSS includes optional utilities to allow remote support of your ftServer system. These include an extensible network administration framework and a server-monitoring utility that provides notification services. Chapter 8 discusses the configuration and use of the optional ftlSNMP package implementing Simple Network Management Protocol (SNMP) for managing network objects. The ftlSNMP package is typically installed with ftSSS.

Troubleshooting ftServer Systems

Problem identification, system and application diagnostics, and system configuration to resolve problems with ftServer systems are essential troubleshooting tasks. Chapter 9 discusses system features and procedures to assist you in troubleshooting ftServer systems.

Additional Documentation and Resources

The following resources provide additional information that may be helpful to you in administering your ftServer system.

Red Hat Enterprise Linux

Documentation for the Red Hat Linux operating system is available at http://www.redhat.com/docs.

Stratus ftServer System Documentation

The ftServer T Series StrataDOC CD-ROM provided with your system contains all of the system documentation for ftServer systems that run the Linux operating system. It is provided in Adobe Acrobat® Portable Document Format (PDF) and HTML format for viewing and printing.

This manual occasionally refers to other documentation that is specific to your particular ftServer system. Table 1-1 lists the system-specific documentation, all of which is available on the ftServer T Series StrataDOC CD-ROM.

Table 1-1. ftServer System-Specific Documentation *(Page 1 of 2)*

Manual	Description
Stratus ftServer 2400, 4300, 4600, and 5700 Systems: Installation Guide (R575)	Describes how to install and set up your ftServer 2400, 4300, 4600, or 5700 system. It also lists the part numbers of customer-replaceable units (CRUs), components that you can easily replace.
Stratus ftServer 2400, 4300, 4600, and 5700 Systems: Operation and Maintenance Guide (R574)	Documents how to operate, diagnose, and maintain your ftServer 2400, 4300, 4600, or 5700 system. It explains how to remove and replace the CRUs and how to interpret system operational status based on the state of the light-emitting diodes (LEDs).
Stratus ftServer T40 CO Systems: Installation Guide (R588)	Describes how to install and set up your ftServer T40 CO system. It also lists the part numbers of CRUs.
Stratus ftServer T40 AC and T65 AC Systems: Installation Guide (R596)	Describes how to install and set up your ftServer T40 AC or T65 AC system. It also lists the part numbers of CRUs.

Table 1-1. ftServer System-Specific Documentation *(Page 2 of 2)*

Manual	Description
Stratus ftServer T40 AC and T65 AC Systems: Operation and Maintenance Guide (R597)	Documents how to operate, diagnose, and maintain an ftServer T40 AC or T65 AC system. It explains how to remove and replace the CRUs and how to interpret system operational status based on the state of the LEDs.
Stratus ftServer T40 CO Systems: Operation and Maintenance Guide (R589)	Documents how to operate, diagnose, and maintain an ftServer T40 CO system. It explains how to remove and replace the CRUs and how to interpret system operational status based on the state of the LEDs.

Linux and UNIX Documents

The following sources provide further information about the Linux and UNIX operating systems.

- The Linux Documentation Project: http://www.tldp.org/

  The *System Administrator's Guide* LDP v0.8 and the *Network Administrators Guide* LDP v2.0 are two freely redistributable publications available from the Linux Documentation Project Web site that you can use to supplement this *Stratus ftServer System Administrator's Guide for the Linux Operating System* (R003L). The LDP also provides other online manuals, how-to documents, and links to additional online-accessible data.

- *Sunsite* Web page, University of North Carolina, provides well-ordered archives and links to many resources: ftp://sunsite.unc.edu/pub/Linux/.

- *Linux Administration Handbook*, by Nemeth, Snyder, and Hein, copyright 2002, Prentice Hall PTR, div. of Pearson Education, Inc.:

  http://vig.prenhall.com/catalog/academic/product/0,4096,0130084662,00.html
  This volume is a reference manual for both system and network administration of the Linux operating system. It focuses on available (at time of publication) open source tools but incorporates in-depth knowledge of UNIX administration utilities and network management practices.

- *Linux in a Nutshell--A Desktop Quick Reference*, 4th Ed., by Siever, Figgens, and Weber, copyright 2003, O'Reilly & Associates, Inc.:

  This book can help you assess available tools and assemble an effective toolkit for managing servers and networks, for example.

- *Open Source Network Administration*, by James Kretchmar, copyright 2004, Prentice Hall PTR, div. of Pearson Education, Inc.:

  http://vig.prenhall.com/catalog/academic/product/0,4096,0130462101,00.html
  This volume is a basic reference for common open source TCP/IP network administration utilities and technologies, including SNMP tools and methods.

- *UNIX Systems: Advanced Administration and Management Handbook*, Bruce H. and Karen B. Hunter, copyright 1996, Pearson Education. Although this book is an older book that covers only the UNIX operating system, the advice and wisdom packaged here for administrators of UNIX-type operating systems and TCP/IP networks is usually available only through intensive experience.

Chapter 2
Installing the Operating System and ftServer System Software

This chapter discusses the following topics:

- "Installation Overview"
- "Separately Released and Optional Distribution Components"
- "Installation Interfaces"
- "Supported Hardware and Firmware"
- "Pre-Installation Checklist"
- "Installing the Linux Operating System and ftSSS"
- "Post-Installation Tasks and Considerations"
- "Performing an Installation Without a Kickstart File"
- "Additional Documentation and Resources"

The procedures described in this chapter are for a fresh installation or complete reinstallation of a supported Linux operating system and the ftServer System Software for the Linux Operating System (ftSSS) on ftServer systems. Chapter 4 describes an upgrade procedure and system and package restoration procedures for upgrading a recent distribution to the current software distribution level.

NOTE

Be sure to read the accompanying *Release Notes: Stratus ftServer System Software for the Linux Operating System* (R005L) document before you undertake an initial installation or a reinstallation of Linux and ftSSS.

Installation Overview

An installable distribution CD-ROM (CD) set is provided. Table 2-1 lists the CDs included in this distribution.

Table 2-1. CD-ROMs Included With ftServer Systems

CD-ROM	Contents
Stratus ftServer System Software for the Linux Operating System	Stratus fault-tolerant system software
ftServer StrataDOC (Linux Version)	The system hardware and software documentation, including the present document, in Adobe Acrobat® PDF and HTML formats.
ftSSS Debug Info	Includes `debuginfo` RPMs.
Set of Red Hat® Enterprise Linux 4 operating system CDs	Red Hat Enterprise Linux 4 operating system, related packages, and documentation

If your site did not purchase the Linux operating system distribution from Stratus, you must perform the complete initial Linux operating system installation. However, if you purchased an ftServer system and the Linux operating system from Stratus, the Linux operating system and ftSSS are preinstalled. You do not need to reinstall or upgrade this software. You should familiarize yourself with this chapter, then log in.

To log in to the system

1. Log in as `root`.
2. Enter the default password, `ftServer`.

(A password must have at least six characters.) After you read and accept the necessary end-user license agreements, continue to Chapter 5 to begin configuring the ftServer system.

If you have a current installation that requires only minor upgrading but is at least as recent as Release 4.0, review your *Release Notes: Stratus ftServer System Software for the Linux Operating System* (R005L), then see Chapter 4 and, if necessary, Chapter 3.

If some time has elapsed since your Stratus ftServer system was delivered, there may be updated documentation and software that may be useful to you.

NOTE

Firmware updates may require ftSSS updates. ftSSS updates may require firmware updates. When updating ftSSS, use the versions of firmware and software that are supplied on the ftServer System Software for the Linux Operating System CD.

From time to time, Stratus may issue an update to ftSSS. See Chapter 4 for information about updating from an ftSSS update disk.

CAUTION

The procedure described in this chapter is for a full installation or reinstallation of a supported Linux operating system and ftSSS. It is assumed that no valuable data exists on the target system disks. **The installation and upgrade installation procedures will destroy existing data on the drives in the bottom slot (labeled 1) of each CPU-I/O enclosure**.

If you need to repair a corrupted system, or update the system to a new release, determine whether the upgrade procedure in Chapter 4 will meet your requirements before doing a full installation.

Boot Media

The supported Linux operating system is provided on a set of distribution CDs available from Stratus or the Linux operating system vendor. See the *Release Notes: Stratus ftServer System Software for the Linux Operating System* (R005L) for information about specific, supported Linux operating systems.

You should boot from the ftSSS CD, unless you are using a custom kickstart file.

Default System Setup

This section provides an overview of the default setup that is provided on initial installation or complete reinstallation of the Linux operating system distribution. It does not reflect installation of optional packages.

NOTE

The Linux operating system installer program does not support customer-added and unknown hardware. Any such hardware should be added, and the system configured as required to support it, only after installation procedures have been completed and the system has been determined to function as expected.

Linux Version Information

You can check the installed version of the Linux operating system on your system using the `uname` command. The `-r` specifies that the kernel release level be returned.

```
# uname -r
2.6.9-34.EL.serial.idesmp
```

To check the installed ftSSS distribution release level using the **rpm** command, enter:

```
$ rpm -q lsb-ft-eula_display
lsb-ft-eula_display-4.0-65
```

Storage Default Settings

CAUTION

Prepare for installation by installing disk drives of identical size and geometry into the internal storage in the bottom slot (the two slots labeled 1: sda and sdd) of each CPU-I/O enclosure. Remove all internal drives from the other four slots (sdb, sdc, sde, and sdf). The installation process pairs the two installed drives.

NOTES

1. **Do not** install 400GB SATA disk drives in slots sda and sdd. For performance reasons, these drives are not supported as boot disks, but can be used freely as data disks.
2. ftServer 2400, 4300, and 5700 systems may be mounted in a rack or in a pedestal. ftServer T40 CO, T40 AC, and T65 AC systems are available only in a rack. The terms *top* and *bottom* assume a rack (horizontal) installation. In pedestals, systems are rotated 90 degrees counterclockwise from their rack-mounted position, so *top* becomes *left* and *bottom* becomes *right*.

The installation process creates a disk drive RAID array, pairing sda and sdd drives as a mirrored set that holds the entire installed Linux software distribution and ftSSS. On this mirrored drive set, the GRUB bootloader on the master boot record at track 0 makes both drives in the set bootable using GRUB configuration data stored in the /boot partition. Storage is allocated as shown in Table 5-1.

- The /boot directory is installed as an ext2 file system on `/dev/sda1` and `/dev/sdd1`, on a partition of at least 256 MB. This partition contains the GRUB boot loader configuration file and GRUB restricted shell, as well as the Linux kernel and initial RAM disk files.
- A swap partition provides 2048 MB of swap space.

NOTE

Regardless of installed system memory, the current Linux kernel can only use about 2 GB of swap space per swap partition. The kernel can address swap partitions on more than one drive, so it is possible to improve swap performance on multi-drive systems by using swap partitions on each installed drive pair. However, for fault tolerance, always use mirrored swap partitions. See ‘‘Setting Up Internal Disk Storage’’ on page 5-2 for more information.

- The root partition comprises 32 GB.
- The remainder of the space is an extended partition that includes 36 GB or more of unused space. This extended partition can be further divided by logical partitions.

All internal drives used in mirrored RAID arrays in the internal hot-swap drive bays must have firmware that meets ftServer specifications. **Do not update your internal drives with firmware from sources other than Stratus.** Contact your Stratus Customer Assistance Center (CAC) or your authorized Stratus service representative for your internal SATA drive support requirements.

Ethernet Default Configuration

At installation, the 10/100/1000-Megabits per second (Mbps) embedded Ethernet adapters of each CPU-I/O enclosure are configured. For more information about configuring Ethernet interfaces, see ‘‘Ethernet Devices’’ on page 5-25.

USB and RS232C Default Settings

Only the USB keyboard and, optionally, a USB mouse should be attached to the system during an initial installation of the operating system.

In most cases, attached devices are recognized and addressable on installation (as is a standard USB keyboard, for example), although hot-plugged devices may not be. USB 2.0 interface specifications are supported.

After installation, you may need to set serial-port flow control and data-rate characteristics for attaching serial data communications equipment or data terminal equipment, such as an asynchronous terminal, a printer, or attached modem. Note that the Linux operating system and most application software treat data communications equipment somewhat differently from data terminal equipment. For information on setting up Serial Port 1 for ActiveService Network (ASN) use, see one of the following documents:

- *Stratus ftServer 2400, 4300, 4600, and 5700 Systems: Installation Guide* (R575)
- *Stratus ftServer T40 CO Systems: Installation Guide* (R588)
- *Stratus ftServer T40 AC and T65 AC Systems: Installation Guide* (R596)

Default System Initialization and Run-Level Control

Default system initialization and run-level-controlled process configuration are basic and not tailored specifically for server operations. You will need to configure the system as required for your specific application.

Default User Environments, Shells, and Access Control

The GRUB bootloader package supplied with the Linux distribution includes a restricted GRUB shell that can be entered at system boot for boot loader configuration and boot recovery operations. See *grub*(1) for a discussion of this feature. Make sure that GRUB requires root privilege, and password-protect this shell for system security.

System Indicators and Switches

See the operation and maintenance guide for your system for information about the LED indicators and switches of your ftServer system.

Separately Released and Optional Distribution Components

You can install and use provided optional tools. You can install optional packages by using the `rpm` command to select and install packages from a mounted CD-ROM drive. Appendix A describes how to use the `rpm` command to learn more about the packages provided on the ftServer System Software for the Linux Operating System CD-ROM.

CAUTION

Some installed and optional utilities may depend on specific versions of other packages required by your system for fault-tolerant operation. Always use the software packages that are provided with the ftSSS

distribution, unless you have good reason to replace a package. Before updating a distribution package, use `rpm` to check dependencies. Note that `rpm` does not always reveal specific release-level dependencies.

From the ftSSS distribution, ASN and ftISNMP packages are installed as options and require additional configuration before they can be used. See ‘‘ActiveService Network Support’’ on page 7-14 and Chapter 8 for information on configuring and using these utilities.

Installation Interfaces

The installation process has two parts. First, you install the Linux operating system, and then you perform the ftSSS installation.

You must connect a supported monitor to the VGA port on the rear of the system, and a supported USB mouse and keyboard to a USB port on the rear of the system.

The Linux operating system and ftSSS installation process use an attached LCD or SVGA-capable monitor attached to the SVGA connector at the ftServer back panel and a USB mouse and keyboard attached to a USB port.

Supported Hardware and Firmware

The Linux operating system combined with ftSSS can be installed only on supported ftServer T40 CO, T40 AC, 2400, 4300, and 5700 systems. Do not attempt to install this software combination on an ftServer system that does not support it. ftSSS interfaces with ftServer firmware that has been tuned to support fault tolerance. Specific ftSSS releases may require corresponding updates to the system BIOS or baseboard management controller (BMC) firmware.

Pre-Installation Checklist

The following checklist is provided for an initial installation or a full reinstallation of a supported Linux operating system and ftSSS. If you need to upgrade an existing ftSSS release, see Chapter 4.

❏ Check that you have current release notes and installation guides for your distribution. To restore your installation, make sure you have the appropriate version of the release level that you will restore. The *Release Notes: Stratus ftServer System Software for the Linux Operating System* (R005L) in PDF format, on the ftServer T Series StrataDOC CD, provides the information that you need.

Also check the system hardware against the release notes. The I/O elements and installed devices must comply with any identified configuration requirements and

support restrictions on hardware that apply either to this installation procedure or to the current ftSSS release generally.

- ❏ The installation CDs ask that you read and accept end user license agreements (EULAs). You should not perform the installation if you cannot accept the EULAs or are not authorized to accept them. Installation terminates without completion if you decline a required EULA. You can read the text of the required EULAs in the installation guide for your system.
- ❏ If you are reinstalling a release, back up all data files and prepare backup files to reconfigure the system for security and network operation after the installation procedure has been performed.
- ❏ Verify that your system BIOS and BMC firmware versions are compatible with the ftSSS release you are about to install. You can obtain required versions of firmware from the ftSSS CD-ROM.

  If system BIOS or BMC firmware updates are needed, you must update the firmware **before** you begin the installation or upgrade process. You can obtain the required firmware from the ftSSS CD that comes with the distribution. See Chapter 3 for details.

- ❏ Use the CD-ROM drive in the top CPU-I/O enclosure during the installation. Do not use the drive in the bottom CPU-I/O enclosure. Verify that the top enclosure is the active enclosure.
- ❏ Ensure that the keyboard, mouse, and console are attached to the system. The installer uses a graphical user interface on the ftServer system console, which consists of a monitor attached to an SVGA port and an attached USB keyboard and mouse.
- ❏ With the ftServer system halted, detach all peripheral devices from the system. This includes the following items.
  - Optional ftScalable Storage arrays (remove optical FC cables from FC PCI adapters)
  - Unsupported adapters
  - USB devices (except keyboard and mouse)
  - Other serial devices

  No USB devices other than the keyboard and mouse should be attached to the ftServer system.

- ❏ No external modem or other devices should be attached to the serial ports of the ftServer system.
- ❏ The bottom slots (slots sda and sdd) in the internal storage enclosures must contain a matched pair of supported disk drives having the same model, firmware level, and geometry. These must **not** be 400GB SATA drives. **All other drives**

must be removed from the system. See the drive arrangement in Figure 2-1. For information about the supported disk drives, see the operation and maintenance guide for your system.

During the installation, the two installed drives will be paired and configured using RAID-1 mirroring. A Linux boot partition, a swap partition, and a root partition will be installed on the paired drives.

Figure 2-1. SATA Drive Arrangement for Installation

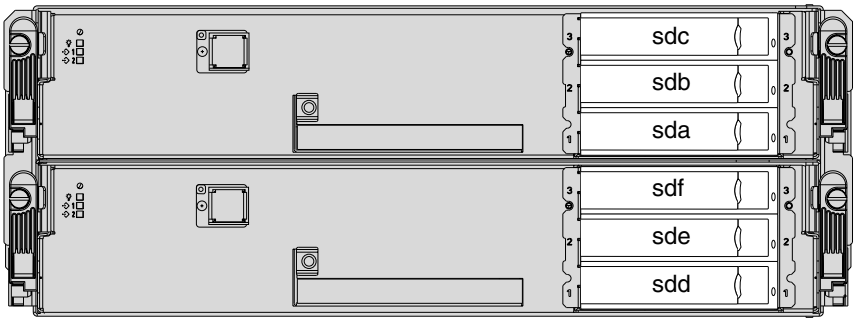

- ❏ Make sure that the system and monitor power connections are secure and firmly plugged in before beginning an installation procedure. Power cabling should be guarded against inadvertent disconnection during the installation process. The monitor may use a separate power source.
- ❏ Begin the installation process with both CPU-I/O enclosures inserted and with the ftServer system halted. A red LED is lit at the left side of each CPU-I/O enclosure. Ensure that the green power button at the right side of the top CPU-I/O enclosure is lit; the lit LED indicates that it is the active (primary) enclosure. If the top CPU-I/O enclosure is not the active enclosure, you must make it the active enclosure by performing the actions listed in step 2 of the installation procedure in "Installing the Linux Operating System and ftSSS" on page 2-10.
- ❏ For security during installation and initial configuration, isolate the ftServer system from networks and other communicating hosts.
- ❏ Be aware that a system with a newly installed operating system and ftSSS is not secure until it has been configured. In addition, you may want to change the default network settings.

Installing the Linux Operating System and ftSSS

The installation procedure requires five operating system CDs (for Red Hat Enterprise Linux 4) and one ftSSS CD. The installation process creates two bootable hard drives.

When the operating system installation is complete, you will power down the system, reboot, and install ftSSS fault-tolerant software.

CAUTION

The operating system installation process destroys all data on the two hard drives.

When the installation is complete, the two disks will be mirrored, bootable disks. Before beginning installation, make sure that you have completed the items in "Pre-Installation Checklist" on page 2-7.

Instructions for installing or reinstalling the Linux operating system and ftSSS, are in the following sections:

- "Booting the Operating System"
- "Installing the Operating System"
- "Post-Installation Tasks and Considerations"
- "Installing ftSSS for Fault Tolerance"
- "Reinstalling ftSSS After a Failed Installation"

Booting the Operating System

Perform the following steps to boot the system from the ftSSS CD and to begin the operating system installation process.

1. Ensure that the system is configured according to the instructions in "Pre-Installation Checklist" on page 2-7.

   NOTE

   Make sure that identical hard disks are inserted in the bottom slots—slots sda and sdd—of the internal storage enclosures. **Remove all other internal disks** from the system.

2. Verify that the power switch is lit (green) on the **top** CPU-I/O enclosure. Then lift the switch cover and press the switch for 2 or 3 seconds to start up the system.

   If the bottom CPU-I/O enclosure is the one whose power switch is lit, follow these steps to make the top enclosure the active enclosure.

a. Remove the power cord from the bottom enclosure for 10 seconds.

b. Lift the power-switch cover of the top CPU-I/O enclosure, whose power switch should now be lit, and press it for 2 or 3 seconds to restart the system.

3. Put the ftSSS CD into the CD-ROM drive in the top CPU-I/O enclosure. The system boots from this CD.
4. As the system begins to boot, press **F2** to enter the BIOS Setup program.
5. In the BIOS Setup program, select the following values:

   a. In the Boot menu, select **CD-ROM** as the top boot device to make it the first bootable device on the system. (To move **CD-ROM** to the top of the list, highlight it, then press **Shift**-**+** to move it up the list.) Also, no other bootable drives should precede the internal drives in the boot sequence.

   b. In the Advanced menu, click the Monitoring Configuration tab, then set Boot Monitoring to **Disabled**

      If the system contains any PCI adapters that contain option ROMs, the system boots faster if you disable Option ROM for the PCI slots that contain the adapters. To disable Option ROM, select Monitoring Configuration on the Advanced tab. Select each slot that contains an adapter and use the PLUS SIGN (+) key to change the values to Disabled.

6. **Save** the BIOS configuration (press **F10**), then **Exit** to continue booting the system.

   When the startup screens and BIOS POST messages are complete, the system boots from the CD-ROM drive.

7. A screen with the Red Hat Logo appears, prompting you to choose an installation method. To install the default Red Hat Linux packages for ftSSS, press **Enter**. Otherwise, follow the instructions on the screen.
8. When you are prompted to specify the location of the boot media, remove the ftSSS CD and replace it with Linux operating system CD number 1, the boot CD. Close the drive tray and then select **Local CDROM**.
9. Continue with “Installing the Operating System.”

Installing the Operating System

It takes about two hours to install all of the required packages of an ftServer default installation. Perform the following tasks to install the operating system from the operating system installation CDs.

1. Follow the installation-process prompts that instruct you when to eject the currently inserted Linux installation CD and insert the next one. After inserting the new CD, press **Enter** on the keyboard.

   After having put all of the Linux CDs into the drive in succession, you will be prompted to reinsert CD 1 to install a few more packages.

   After the operating system is completely installed, file-system dismount and system-shutdown messages appear on the screen.

2. When the message `You may safely reboot your system` appears, remove CD 1 from the drive.
3. Press the lit power button on the top CPU-I/O enclosure to power down the system.
4. When the LEDs on the disk drives and the CD-ROM drive are unlit, press the lit power button on the top CPU-I/O enclosure to reboot the system. The Linux operating system boots.
5. After several moments, a number of Red Hat setup screens and licensing agreements appear for your input.
   a. The Language screen appears. Select the default language for the system and click **Next**.
   b. The Welcome screen appears. Click **Next**.
   c. The License Agreement screen appears. Read the license agreement. Select the **Yes** option to accept the terms of the agreement. Click **Next**.
   d. The Keyboard screen appears. Select the type of keyboard for the system and click **Next**.
   e. The Mouse Configuration screen appears. Select the type of mouse for the system and click **Next**.
   f. The Root Password screen appears. Create a password for root. The password must contain at least six characters. Type the password again in the Confirm box, and click **Next**.
   g. The Network Setup screen appears. Select the default configuration and click **Next**.
   h. The Security Level screen appears. Enable or disable the firewall, and select trusted services and devices. Click **Next**.
   i. The Time Zone screen appears. Select the time zone for the system and whether you use local time or Coordinated Universal Time (UTC). Click **Next**.

j. The Date and Time screen appears. Select the current date and time. Click **Next**.

k. The Display screen appears. Select the display settings and click **Next**.

l. The Red Hat Login Screen appears. Enter or create a Red Hat login to activate the services included in your operating system subscription. Click **Next**.

m. The System User screen appears. Set up a non-administrative user account and click **Next**.

n. The Additional CDs screen appears. Insert a Linux Extras disk and click **Install**. Alternatively, skip this step by clicking **Next**.

o. The Finish Setup screen appears. Click **Next**.

6. The Red Hat login screen appears. Log on to the system as username `root` and with the password you created in item f in the previous step. (In the event you did not create a new password, the default password is `ftServer`.)

Installing ftSSS for Fault Tolerance

Perform the following tasks to install the ftSSS and enable fault-tolerant operation of your system. This process may take about an hour.

1. With the system booted and the top CPU-I/O enclosure the active enclosure, insert the ftServer System Software for the Linux Operating System CD in the top CPU-I/O enclosure.
2. Right click on the desktop and click **Open Terminal**.
3. If the CD-ROM drive is not already mounted, mount it by typing the following command:

   ```
   # mount /media/cdrom1
   ```

4. Run the ftSSS installation script by typing:

   ```
   # /media/cdrom1/install.sh
   ```

5. The installation program displays prompts that ask you to answer questions, accept EULAs, run system checks, and advise you of your system's fault-tolerance policy. Press **Enter** at **all of the prompts** the installation program presents **except** the final prompt to reboot your system.
6. If you do **not** have an ftScalable Storage array, press **Enter** at the prompt to reboot your system and skip ahead to step 13.

   If you **do** have an ftScalable Storage array type NO at the system boot prompt and go on to the next step.

7. Add the line starting with `vendor=Stratus` to the `/etc/scsi_id.config` file as shown in the example below:

```
# If you normally don't need scsi id's, or might be
# attaching devices of an unknown functionality, black
# list everyone. This is the default
# behaviour (if no -b or -g is specified).
#
options=-b

# Then white list devices on your system that have correct and
useful id's:

vendor=Stratus, model="AA-D91900", options=-g
```

8. Ensure that the ftScalable Storage array is powered on and ready, then reconnect the FC cables between the ftServer system and the array's controller tray. See the *ftScalable Storage: Getting Started Guide* (R601) for details.
9. Type the following commands to start the `multipathd` daemon

```
# multipath -v2 -d
# chkconfig --add multipathd
# chkconfig multipathd on
# multipath
# /etc/init.d/multipathd start
```

   The daemon will start automatically on subsequent reboots.

10. Verify that the multipath arrays are now visible by viewing the `/dev/mapper` file. By default, each array has a persistent name based on its world-wide ID.
11. Create a file system and mount the array.
12. Manually reboot the ftServer system by typing the following command.

```
# shutdown -r now
```

13. After the system reboots, when the ftServer startup screen appears again, press **F2** to enter the BIOS configuration menu.
14. In the Advanced menu, click the Monitoring Configuration tab and set OS Boot Monitoring to **Enabled**.
15. In the Boot menu, set the first boot device to be the **Hard Drive**.
16. **Save** the BIOS configuration (press **F10**), then **Exit** to continue booting the system.
17. After the system boots, remove the CD from the CD-ROM drive. The installation of the operating system and ftSSS is now complete.

To quickly verify the installation, issue the following command:

```
# /opt/ft/bin/ftsmaint --nocheck ls
```

The command should output an inventory of system components and their operational status. Your system software installation is complete, but you still need to configure the network and perform a few other tasks. See ‘‘Post-Installation Tasks and Considerations.”

Reinstalling ftSSS After a Failed Installation

If the attempt to install ftSSS fails, before again trying to install it, perform the following steps to uninstall the requisite files and attempt to reinstall ftSSS in the correct sequence.

1. Remove all Stratus packages except `eula_display` with this command:

   ```
   rpm -e --nodeps --allmatches `rpm -qa | grep lsb-ft | grep
   -v eula_display`
   ```

2. Remove `eula_display` with this command:

   ```
   rpm -e --nodeps lsb-ft-eula_display
   ```

3. Reinstall ftSSS by mounting the CD and issuing the `install.sh` command. If the system does not then reboot automatically, reboot the system manually.

Booting in Linux Rescue Mode

If the OS installation fails, whether or not ftSSS installed successfully, attempt to boot the system in rescue mode. Use the Red Hat Enterprise Linux CD-ROM #1 to boot in rescue mode.

To boot in rescue mode

1. Disconnect any floppy disk drive attached to the system’s USB port.

   NOTE

   If a floppy drive is connected when you boot in rescue mode, the system will be unable to find the internal storage drives.

2. Insert Red Hat Enterprise Linux CD-ROM #1 into the CD-ROM drive in the top CPU-I/O enclosure. The system boots from this CD.
3. Complete steps 4 through 6 of the procedure in ‘‘Booting the Operating System’’ on page 2-10.

4. After the system boots from the CD, and **as soon as the boot prompt appears**, type the following line at the `boot` prompt and press **Enter**:

```
# boot: linux rescue
```

NOTE

You must type something (at least one character) on the boot prompt line before its timeout period expires. Otherwise, the boot will proceed with incorrect parameters and the keyboard will be disabled. If this happens, you can recover by power-cycling the system and booting again.

5. After several minutes, the Language prompt appears. Follow the prompts.
6. Issue the following command:

```
chroot /mnt/sysimage
```

You may need to remove and re-apply system power before you can reboot from the disk.

If the disks are not detected and automounted, contact the CAC or your service representative for assistance.

Post-Installation Tasks and Considerations

After installing the operating system and ftSSS consider the following topics.

- Default Configuration Notes”
- Configuring the Network”
- Adding Fault-Tolerant Utilities to PATH”

Default Configuration Notes

After installation, the default installed system should appear as described in ‘‘Default System Setup’’ on page 2-3. The following notes apply to the default system configuration.

NOTES

1. After you upgrade or restore the Linux operating system and ftSSS distributions as described in this chapter, you may also need to separately install optional packages using the `rpm` command.

2. The system disk pair was created as a RAID-1 mirrored drive set on drives sda and sdd. Each drive is bootable and configured identically to the other.
3. After installation, the kudzu, haldaemon, and microcode_ctl services are disabled. Do **not** enable these services.

Configuring the Network

Perform the following tasks to configure the system for operating on a network.

1. Issue the following command to launch the Linux graphical network-configuration program:

```
# system-config-network
```

2. In the graphical network configuration tool, specify the network hostname for your system.

See the section on adding hosts in the Red Hat Enterprise Linux 4 System Administration Guide for detailed information about configuring networks.

Adding Fault-Tolerant Utilities to PATH

Stratus fault-tolerant utilities, like `ftsmaint` and `ASNConfig`, reside in the `/opt/ft/bin` and `/opt/ft/sbin` directories. Consider setting your PATH to include these directories.

Performing an Installation Without a Kickstart File

The ftServer System Software for the Linux Operating System CD contains three kickstart files:

- `ks.cfg`—Use this file to perform a full installation. A *full installation* installs the complete set of Red Hat files as well as X Windows.
- `ks_min.cfg`—Use this file to perform a minimal installation. A *minimal installation* installs only those Red Hat files that are required for ftServer systems running a supported Linux distribution together with ftSSS. It does not install X Windows.
- `mfg.cfg`—Stratus uses this file to perform a full installation at the factory. Do not use this file.

You should use either the `ks.cfg` or `ks_min.cfg` kickstart file to perform an installation. However, if you choose not to do so, be aware of the caveats in the following text, and contact the CAC or your customer service representative to ensure that your actions do not void your support agreement.

To install the system software without using an ftServer kickstart file

1. After installation, while the system is booting, the GRUB menu must supply `reboot=warm` at the boot prompt. If you do not type this line, your system will not be fault-tolerant.

   ```
   reboot=warm nmi_watchdog=0 clock=tsc
   ```

   You must type something (at least one character) on the `boot` prompt line before its timeout period expires. Otherwise, the boot will proceed with incorrect parameters and the keyboard will be disabled. If this happens, you can recover by power-cycling the system and booting again.

   This command is necessary for your ftServer system to be fault-tolerant.

   You can find additional information about the kickstart files in the Red Hat Linux OS system administration documentation.

2. Manually make the second disk a bootable disk. At the command prompt, type the following lines to make both system disks bootable:

   ```
   # /sbin/grub
   device (hd0) /dev/sda
   root (hd0, 0)
   setup (hd0)
   device (hd0) /dev/sdb
   root (hd0, 0)
   setup (hd0)
   quit
   ```

- Make sure that the system is running the SMP kernel.
- You must install all required software packages. If you do not use the software package selections from the supplied kickstart files, you may have to manually resolve package dependency failures when installing ftSSS.
- Make sure that the kernel command line does not contain `rhgb`.
- Make sure that services like kudzu, haldaemon, and microcode_ctl are disabled.
- Make sure that partitions are set up correctly.

Additional Documentation and Resources

- The *GNU GRUB* Web page, Free Software Foundation:

  http://www.gnu.org/software/grub/grub.en.html

- *GRUB: GRand Unified Bootloader version 0.5*, original documentation Web site:

  http://www.uruk.org/orig-grub/

- *GRUB file system syntax and semantics* document, by Eric Bolyn, primary author of GRUB:

  http://www.uruk.org/orig-grub/filesystem.txt

Chapter 3
Updating ftServer System Firmware

This chapter discusses the following topics:

- "Updating the System BIOS"
- "Updating BMC Firmware"

Consult the *Release Notes: Stratus ftServer System Software for the Linux Operating System* (R005L) for the ftServer System Software for the Linux Operating System (ftSSS) version you have (or will upgrade to) to determine what firmware version numbers are required.

CAUTION

Update your system firmware only if that version is compatible with your current ftSSS installation, or if you will immediately update to operating system or ftSSS releases that are supported by the updated firmware.

Updating the System BIOS

Make sure that the BIOS you intend to install is compatible with the ftSSS release level that you have (or that you will install immediately after performing the BIOS upgrade). Verify in the applicable version of the *Release Notes: Stratus ftServer System Software for the Linux Operating System* (R005L) that the BIOS is compatible, and also that the Release Notes do not specify a different procedure or sequence for performing the BIOS update. To check the BIOS version, type the following command:

```
# /opt/ft/bin/ftsmaint ls 0
.
.
.
Firmware Rev   : 20.0
```

The preceding example displays a BIOS version number of 20.0 for the top CPU-I/O enclosure (see Table 7-1 for a list of system device IDs).

Stratus ftServer BIOS updates are image files that you must transfer from removable media or download from a network-accessible archive.

Take care when updating firmware. It is a necessary failover characteristic for the CPU-I/O enclosures to be paired in duplexed operation. On system boot and before duplexing, if the ftServer system detects differing firmware between the system CPU elements, the firmware from the CPU-I/O enclosure that is booting is replicated to the other enclosure automatically, allowing the enclosures to synchronize.

NOTE

If an ftServer system boots from the system enclosure that has not been upgraded, the upgraded system firmware will be overwritten with older firmware in order to synchronize to duplex. This defeats the intended upgrade.

CAUTION

Do not interrupt a BIOS firmware update while a *burn* (that is, a write to EEPROM) is in progress. Interrupting power to a CPU-I/O enclosure during a burn can result in EEPROM damage that prevents it from storing the firmware successfully. Correcting this problem may require that you obtain a replacement CPU-I/O enclosure.

Read this entire procedure to review your options before beginning.

To update the system BIOS

1. Insert the ftServer System Software for the Linux Operating System CD in the CD-ROM drive in the CPU-I/O enclosure whose power button is lit.
2. Log on to the system as `root`.
3. Mount the CD-ROM drive by typing the following command:

   ```
   # mount /media/cdrom
   ```

4. Use the `ftsmaint` command to verify that you are starting from a known, good state. At this point, both CPU-I/O enclosures should be operating duplexed.

   ```
   # /opt/ft/bin/ftsmaint ls 0

   H/W Path                 : 0
   Description              : Combined CPU/IO
   State                    : ONLINE
   Op State                 : DUPLEX
   Reason                   : SECONDARY
   LED (Green)              : ON
   LED (Yellow)             : OFF
   LED (White)              : ON
   Modelx                   : AA-G94340
   Firmware Rev             : 1.3:90
   ECO Level                : 37
   Min Partner ECO Level    : 0
   Serial #                 : 401318
   MTBF Fault Count         : 0
   MTBF Last Timestamp      : None
   MTBF Threshold           : 1200
   MTBF Value               : 0
   MTBF Type                : useThreshold
   Logic Revision           : 18023
   ```

   Then run the following command:

   ```
   # /opt/ft/bin/ftsmaint ls 1
   ```

   NOTE

   To determine what model your system is, type the `ftsmaint lsSystem` command. The model is identified in the `Description` field.

5. Obtain the latest BIOS image from the CD. Use the file named `/firmware/bios/aria/g94300bios`*n.n.n*`.rom`. The *n.n.n* represents the revision level of the BIOS.

6. Perform the BIOS burn by issuing the following commands to one of the CPU-I/O enclosures.

```
# /opt/ft/bin/ftsmaint bringDown 0
Completed bringDown on the device at path 0.
# /opt/ft/bin/ftsmaint burnProm
 /media/cdrom/firmware/bios/aria/g94300biosn.n.n.rom 0
Updated firmware on the device at path 0.
# /opt/ft/bin/ftsmaint jumpSwitch 0
Transferred processing to the device at path 0.
# /opt/ft/bin/ftsmaint bringUp 1
Completed bringUp on the device at path 1.
```

   The preceding commands burn the BIOS to both CPUs as follows:

   - `bringDown`: Takes the CPU 0 element, in the top enclosure, out of service.
   - `burnProm`: Burns the new BIOS to CPU 0.
   - `jumpSwitch`: Brings CPU 0 up, runs diagnostics on it, synchronizes it with CPU 1, in bottom enclosure, and then takes CPU 1 down.
   - `bringUp`: Brings CPU 1 up, runs diagnostics on it, automatically burns the new BIOS from CPU 0 onto it, reruns diagnostics, and resynchronizes CPU 1 with CPU 0.

   NOTES

   1. At this point, reboot only if you need to change BIOS configuration settings.
   2. BIOS files are also available on the system in the `/opt/ft/firmware/bios/aria` directory.

7. Repeat step 4 to verify that the CPU-I/O enclosures are again duplexed.
8. If the new BIOS did not perform as you expected, first verify that you do not also need to perform an ftSSS upgrade to use the new BIOS. If so, proceed to step 9 (returning here if the upgrade procedure does not initiate or follow the documented and expected steps). If your operating system is fully up-to-date, it is likely that the BIOS image file was not the correct firmware file for your system, or the EEPROM that holds the BIOS did not properly capture the BIOS. This occurs rarely, but it can happen. In that case, repeating the burn procedure usually works.
9. Double-check that you have a good BIOS image file before attempting the BIOS upgrade again. It should not matter whether you retry the burn on the top CPU-I/O enclosure or the bottom CPU-I/O enclosure, but you should note which option you choose in case troubleshooting is required. Remember that image files are easily corrupted during file transfer if copied from one format to another, as when written out as a regular file rather than stored as an image, or by transfer as a character

file rather than a binary file. You can detect such corruption by computing a checksum with the `md5sum` command before and after copying. A repeated BIOS burn failure is likely to be caused by a command syntax error or by using a damaged or inappropriate BIOS image file.

10. If it is necessary to update the BMC firmware, follow the procedure described in "Updating BMC Firmware" on page 3-5.
11. If you now need to update the Linux operating system and ftSSS, place the ftSSS CD in the top CD-ROM drive and reboot. See the *Release Notes: Stratus ftServer System Software for the Linux Operating System* (R005L) and Chapter 2 for operating system update procedures.

Updating BMC Firmware

Each I/O element contains a socketed Baseboard Management Controller (BMC) chip. Firmware updates are provided on a bootable CD-ROM.

To check the BMC firmware version, type the following command:

```
# /opt/ft/bin/ftsmaint ls 10/120
.
.
.
Firmware Rev     : 7.0.0
```

The preceding example displays a BMC firmware version number of 7.0.0 for CPU element 0.

CAUTION

Do not interrupt a BMC firmware update with a burn in progress. Interrupting power can result in EEPROM damage or corrupted BMC firmware that requires field service or replacement I/O elements to recover.

To update the BMC firmware from a CD

1. Insert the ftServer System Software for the Linux Operating System CD in the CD-ROM drive in the CPU-I/O enclosure whose power button is lit.
2. Log on to the system as `root`.
3. Mount the CD-ROM drive by typing the following command:

```
# mount /media/cdrom
```

4. The latest BMC image, `e90100srabmc.n.n.n.BIN` (where *n.n.n* is the revision level of the BMC firmware) is in the directory `firmware/bmc/aria/g94300srabmc.`*n.n.n*`.BIN` on the CD.

   NOTE

   All ftServers running a supported Linux distribution and ftServer System Software for the Linux Operating System (ftSSS) use the same BMC firmware.

5. Type the following commands to update the BMC firmware on each I/O element:

   ```
   # /opt/ft/bin/ftsmaint burnProm
   /media/cdrom/firmware/bmc/aria/g94300srabmc.n.n.n.BIN
   10/120

   Updated firmware on the device at path 10/120.

   # /opt/ft/bin/ftsmaint burnPROM
   /media/cdrom/firmware/bmc/aria/g94300srabmc.n.n.n.BIN
   11/120

   Updated firmware on the device at path 11/120.
   ```

6. Type the following commands to verify that your BMC firmware is duplexed (be sure that the opstate is `DUPLEX`):

   ```
   # /opt/ft/bin/ftsmaint ls 10/120
   # /opt/ft/bin/ftsmaint ls 11/120
   ```

7. Type the following commands to verify that the I/O elements are duplexed (be sure that the opstate is `DUPLEX`):

   ```
   # /opt/ft/bin/ftsmaint ls 10
   # /opt/ft/bin/ftsmaint ls 11
   ```

Chapter 4
Updating the Operating System and ftServer System Software

This chapter documents how to upgrade the Linux operating system and the ftServer System Software for the Linux Operating System (ftSSS). It discusses the following topics:

- "General Upgrade Considerations"
- "Upgrading or Restoring the Linux Operating System"
- "Upgrading or Restoring Stratus ftSSS"
- "Creating a Backup System Disk"
- "Recovering from a Failed Software Upgrade"
- "Related Information and Resources"

NOTES

1. If you want to update or reinstall an individual software package that is provided in a Red Hat Package Manager (RPM) file, use the `rpm` command (see *rpm*(8)).
2. The Linux operating system upgrade script has been tested only with firmware, hardware, and devices meeting design specifications of the Stratus ftServer system and its system options. See the *Release Notes: Stratus ftServer System Software for the Linux Operating System* (R005L) for issues that may relate to the upgrade process.
3. If you need to perform a complete reinstallation rather than a release-level upgrade or restoration, see Chapter 2 for preparation and for the procedure.

General Upgrade Considerations

When upgrading the Linux operating system or ftSSS, be aware of the following requirements and related considerations.

Upgrade Requirements

First, ensure that the system's BIOS and BMC firmware levels support the new ftSSS version. You can obtain required versions of firmware from the ftServer System Software CD-ROM. If necessary, upgrade the firmware (see Chapter 3).

Optionally, upgrade the Linux operating system, as described in ''Upgrading or Restoring the Linux Operating System'' on page 4-3. Before upgrading the operating system software, check with the CAC or your authorized Stratus service representative to ensure that your ftServer system supports the new version.

Upgrade the ftSSS, as described in ''Upgrading or Restoring Stratus ftSSS'' on page 4-7.

Related Considerations

The upgrade and reinstallation processes do not overwrite the following files, if you have modified them:

- `/etc/modprobe.d/ft-network.conf`
- `/etc/sysconfig/network-scripts/ifcfg-bond0`
- `/etc/sysconfig/network-scripts/ifcfg-bond1`
- `/etc/sysconfig/network-scripts/ifcfg-eth000010`
- `/etc/sysconfig/network-scripts/ifcfg-eth000011`
- `/etc/sysconfig/network-scripts/ifcfg-eth080010`
- `/etc/sysconfig/network-scripts/ifcfg-eth080011`

If the RPM file on an upgrade CD contains updated versions of these files and if you have modified the original file, the upgrade process saves the updated file to the `/etc/opt/ft/network-scripts/ARCHIVE` directory, giving each file a `.rpmnew` extension.

To complete an upgrade to your system, do one of the following:

- If you do not want to preserve your changes, copy the `.rpmnew` file to the appropriate directory, but remove the `.rpmnew` extension.
- To preserve your changes, incorporate the updates into the files you have modified. Compare files in the `/etc/opt/ft/network-scripts/ARCHIVE` directory that have a `.rpmnew` extension to your modified files, and copy the updates from the `.rpmnew` file to your modified file.

For example, if you added a port to the bond defined in the `/etc/sysconfig/network-scripts/ifcfg-bond0` file and if the RPM file for the upgrade contains an update to the `ifcfg-bond0` file from the earlier RPM file, the upgrade process copies the updated `ifcfg-bond0` file from the new RPM as `/etc/opt/ft/network-scripts/ARCHIVE/ifcfg-bond0.rpmnew`. After the upgrade, modify your `/etc/sysconfig/network-scripts/ifcfg-bond0` file with any differences you note in the `/etc/opt/ft/network-scripts/ARCHIVE/ifcfg-bond0.rpmnew` file.

NOTES

1. Resolve the differences in these files immediately, or back up the `rpmsave` or `.rpmnew` files to another location. If the same `.rpmsave` or `.rpmnew` files are generated by uninstalling ftSSS, or by an upgrade or reinstallation, the previous archived versions could be overwritten.
2. If you uninstall ftServer System Software, any of these network files that **you** modified are saved to the `/etc/OPT/ft/network-scripts/ARCHIVE` directory and given a `.rpmsave` extension.

Also whenever you **reinstall** ftSSS, the installation procedure backs up any Ethernet and bonding network configuration files that have been added or edited since the initial installation of ftSSS. The reinstallation procedure backs up the files to the `/etc/opt/ft/network-scripts/ARCHIVE` directory, appending a `.original` extension.

If you **remove** ftSSS or the `lsb-ft-network` package, the removal procedure backs up any Ethernet and bonding network configuration files that have been added or edited since the initial installation of ftSSS. The removal procedure backs up the files to the `/etc/opt/ft/network-scripts/ARCHIVE` directory, appending a `.canonical` extension. Then, the removal procedure restores all backed up network interface files that have the `.original` extension to their original location and file name.

Upgrading or Restoring the Linux Operating System

Use this checklist to prepare for the Linux operating system upgrade or reinstallation:

- ❏ Before performing an upgrade to a system in use, perform a complete file system backup.
- ❏ Carefully read the current *Release Notes: Stratus ftServer System Software for the Linux Operating System* (R005L) document at http://stratadoc.stratus.com for this release. Distribution Release Notes may have been updated with information

about restrictions or problems and workarounds, software updates, and document corrections not found in the Release Notes from earlier distribution CDs. **Verify that the new Linux operating system version you are about to install is intended for installation on *your* ftServer system**.

- ❏ The system that you wish to upgrade **must** be configured to load from a bootable system disk that is in the ftServer system's boot path.

The following topics apply when upgrading or restoring the Linux operating system.

- "Stratus Kernel Modules"
- "Upgrading or Restoring the Linux Operating System"

Stratus Kernel Modules

When the Linux operating system is upgraded, a new Linux kernel is installed.When ftSSS is installed or upgraded, the fault-tolerant ftServer kernel modules are automatically rebuilt at the next boot time. These ftServer modules must be present for fault-tolerant operation. The following requirements must be met, to ensure that the ftServer modules will be rebuilt.

- Because ftServer kernel modules are only built for SMP kernels, make sure that you only use SMP kernels.
- The kernel-smp-devel package must be installed. The version of this package must match the version of the kernel-smp in use.
- The `/boot` file system must be mounted when ftServer kernel modules are rebuilt so the `initrd` can be re-created with these modules (`/boot` is mounted by `/etc/fstab`).
- The directory `/opt/ft` must reside within the root file system.

Immediately after rebuilding kernel modules at boot time, the system automatically reboots to place these rebuilt modules into use. If the system cannot successfully build all the required ftServer kernel modules, a policy of whether a non-fault-tolerant boot is allowed or prohibited takes effect. This policy is set by the `/etc/opt/ft/modules.policy` file and affected by the `/etc/opt/ft/non_ft_boot.sh` script. Execute the `install.sh` ftServer installation script to review or change the policy.

Upgrading the Linux Operating System

Use the Red Hat Update Agent (the `up2date` command) to install a new Linux operating system U update.

Before you install a new Linux operating system U update (for example U5), make sure to do the following:

- See the *Release Notes: Stratus ftServer System Software for the Linux Operating System* (R005L) or check with the CAC to make sure there is an ftSSS version available to support the Linux operating system update.
- Make sure the system is registered with Red Hat and has entitlement on the RHN server.

By default, the Update Agent on a system running ftSSS is configured to access the following servers:

- The Red Hat RHN server for Linux OS patches
- The Stratus YUM server for ftSSS updates. However, at this time, the Stratus YUM server does not supply the latest version of ftSSS.

To upgrade the Linux operating system

1. Start the Red Hat **Update Agent** from the graphical desktop, or by running the `up2date` command on the command line.

   The Update Agent does the following:

   - Queries the RHN server for new versions of RPMs that are already installed on your system.
   - Queries the YUM server for new versions of ftSSS components that are already installed on your system.
   - Lists all available updates.

2. Select the updates you want to install.
   - By default all available OS updates that do **not** affect the kernel are selected for installation.
   - By default all available OS updates that **do** require kernel modification are not selected.

     To install these updates, explicitly select their check boxes.

   - By default, all available ftSSS updates are selected for installation.
3. The Update Agent downloads the updates and prompts you to install them.
4. To install a new version of ftSSS, see “Upgrading or Restoring Stratus ftSSS.”

5. Reboot the system only if you manually selected an update that modifies the kernel.

By default, the Update Agent only updates RPMs that are already installed on the system. To obtain Stratus RPMs whether or not a version of it is currently installed, run the following command:

```
# up2date –installall –channel=Stratus_Technologies_ft_Linux_4.0
```

If the operating system upgrade failed or you want to return the system to the previous operating system version, see ‘‘Recovering from a Failed Software Upgrade’’ on page 4-9.

Restoring the Linux Operating System

1. Make sure that you have a backup system disk, in case the reinstallation fails or you want to return to the previous version of the Linux operating system. See ‘‘Creating a Backup System Disk’’ on page 4-9 for details.
2. Shut down the ftServer system with the command:

   ```
   # halt -p
   ```

3. Ensure that your system is prepared for the reinstallation by checking the items in ‘‘Pre-Installation Checklist’’ on page 2-7.
4. Follow the instructions in ‘‘Booting the Operating System’’ on page 2-10.
5. Follow the instructions in ‘‘Installing the Operating System’’ on page 2-12.

If the operating system reinstallation failed or you want to return the system to the previous operating system version, see ‘‘Recovering from a Failed Software Upgrade’’ on page 4-9.

Your system should now have the same version of operating system software installed as it had previously. But since it has no ftSSS software, its fault-tolerant features are not operational, so you must upgrade or restore ftSSS on your system. See ‘‘Upgrading or Restoring Stratus ftSSS’’ on page 4-7.

Upgrading or Restoring Stratus ftSSS

Prepare for the ftSSS upgrade using this checklist:

- ❏ Before performing an upgrade to a system in use, perform a complete file system backup.
- ❏ If you upgrade the firmware to support the new ftSSS version, backup the BIOS and BMC firmware files on your system. If the upgrade fails or if you choose to return to the previous ftSSS version after the upgrade, you will need these files to reburn the firmware to its earlier versions.
- ❏ If you must restore an installation whose system files are corrupt, this upgrade may fail. To ensure that ftSSS is performing correctly, you should either perform a complete ftSSS reinstallation or restore the system from backups.
- ❏ Carefully read the current *Release Notes: Stratus ftServer System Software for the Linux Operating System* (R005L) document on http://stratadoc.stratus.com. Distribution Release Notes may have been updated with information about restrictions or problems and workarounds, software updates, and document corrections that Release Notes from earlier distribution CDs do not have.

To update or reinstall ftServer System Software for the Linux operating system

1. Make sure that you have a backup system disk, in case the installation fails or you want to return to the previous version of ftSSS. See ‘‘Creating a Backup System Disk’’ on page 4-9 for details.
2. If you want to reinstall the same ftSSS version that is currently installed (for example, to ensure that it is not corrupted), remove the currently installed ftSSS packages by typing the following commands:
   a. Remove all ftSSS packages except `eula_display` by typing the following command:

      ```
      # rpm -e --nodeps --allmatches `rpm -qa | grep lsb-ft | grep -v eula_display`
      ```

   b. Remove `eula_display` by typing the following command:

      ```
      # rpm -e --nodeps lsb-ft-eula_display
      ```

3. If you are upgrading a minimal installation, install the following packages from your Red Hat Enterprise Linux distribution:
   - `libidn.i386`
   - `libidn-devel`

   You can obtain the packages from the Red Hat Network or from a Red Hat Enterprise Linux distribution on CD. The `libidn.i386` package is on Disk 2 and

`libidn-devel` is on Disk 3 of the Red Hat Enterprise Linux Update CD distribution.

Insert Disk 2 in a CD drive and type the following commands to mount the CD, install the `libidn.i386` package, and unmount the CD:

```
# mount -o ro /dev/cdrom /media/cdrom
# cd /media/cdrom/RedHat/RPMS
# rpm -Uvh libidn-*.i386-*.rpm
# cd
# umount /media/cdrom
```

Insert Disk 3 in the CD drive and use the following command to mount the CD, install the `libidn-devel` package, and unmount the CD:

```
# mount -o ro /dev/cdrom /media/cdrom
# cd /media/cdrom/RedHat/RPMS
# libidn-devel-*.rpm --aid
# cd
# umount /media/cdrom
```

Be sure to type two hyphens (`--`) before the `aid` argument in the `rpm` commands.

4. If the value of the if `DefaultDepth` parameter in the `/etc/X11/xorg.conf` file is set to `24` (the default), edit the file to add the last two lines shown in the following:

```
Section "Screen"
     Identifier "Screen0"
     Device     "Videocard0"
     Monitor    "Monitor0"
     DefaultDepth     24
# Comment out next line if DefaultDepth is not 24
     DefaultFbBpp     24
```

5. With the system booted and the top CPU-I/O enclosure the active enclosure, insert the ftServer System Software for the Linux Operating System CD in the top CPU-I/O enclosure.
6. Right click the desktop and click **Open Terminal**.
7. If the CD-ROM drive is not already mounted, mount it by typing the following command:

```
# mount /media/cdrom
```

8. Run the ftSSS installation script by typing:

```
# /media/cdrom/install.sh
```

9. The installation program displays prompts that ask you to answer questions, accept EULAs, run system checks, and advise you of your system's fault-tolerance policy. Press **Enter** at **all of the prompts** the installation program presents **except** the final prompt to reboot your system.
10. Manually reboot the ftServer system by typing the following command.

    ```
    # shutdown -r now
    ```

11. After the system boots, remove the CD from the CD-ROM drive. The installation of the operating system and ftSSS is now complete.

Your system now has a new version of ftSSS installed and the upgrade is complete. If the upgrade failed or you want to return the system to the previous ftSSS version, see ''Recovering from a Failed Software Upgrade'' on page 4-9.

Creating a Backup System Disk

1. Shut down the system with the command:

   ```
   # halt -p
   ```

2. Remove the system disk from the bottom slot of the CPU-I/O enclosure whose power switch is **not** lit and set it aside as a backup disk.
3. To boot the system, lift the switch cover of the CPU-I/O enclosure whose power switch is lit green, and press the switch momentarily.
4. After the system has booted, insert a spare disk in the bottom slot of the CPU-I/O enclosure whose power switch is **not** lit.
5. Perform the procedures described in ''Manually Creating Partitions on Blank Disks and Adding to RAID-1 Arrays'' on page 5-19.

NOTE

Backup disks can be new, factory-fresh disks or disks recycled from other systems. However, care must be taken with recycled disks. The partition table and RAID superblocks that exist on a recycled disk can confuse the system.

Recovering from a Failed Software Upgrade

Use this procedure if an upgrade procedure failed or if you want to go back to the software versions installed before an upgrade procedure was performed.

To recover from a failed software upgrade procedure

1. Shut down the system with the command:

   ```
   # halt -p
   ```

2. Remove the system disks from the bottom slots of both CPU-I/O enclosures and set them aside.
3. Insert the backup system disk you prepared before performing the upgrade in the bottom slot of the CPU-I/O enclosure whose power switch **is** lit (green).
4. Boot the system by lifting the switch cover of the core CPU-I/O enclosure whose power switch **is** lit (green), and by pressing the switch momentarily.
5. After the system has booted, insert the one of the disks you set aside in step 2 in the bottom slot of the CPU-I/O enclosure whose power switch **is not** lit.

   The drive spins up, the system adds it to the RAID array and resynchronizes it so that it mirrors the backup system disk. You can monitor the mirroring process of drives sda and sdd by using the tools described in Chapter 5 or by monitoring the `/proc/mdstat` file.

6. If you upgraded the BIOS and BMC firmware during the upgrade process, you must restore the firmware to its previous version. See Chapter 3 for details.

The system is now restored to the pre-upgrade state.

Related Information and Resources

Release Notes: Stratus ftServer System Software for the Linux Operating System (R005L)

Chapter 5
Setting Up the ftServer System

This chapter discusses the following topics:

- "Setting Up Internal Disk Storage"
- "Setting Up RAID Arrays on Internal Disks"
- "Removing and Replacing Internal Disks"
- "Administering RAID Arrays on Internal Disks"
- "System Backup and Disaster Recovery"
- "Ethernet Devices"
- "Other System Configuration Information"
- "Additional Documentation and Resources"

At system startup, the operating system autoprobes hardware for legacy devices and attached devices that are not already configured for use on the system. Often, the device is recognized and automatically supported, requiring no direct configuration.

While the operating system may recognize legacy devices, Stratus ftServer System Software for the Linux Operating System (ftSSS) does not support them as fault-tolerant devices. To be supported as a fault-tolerant device, a device must have a special hardened driver that supports surprise removal and fault management by ftSSS.

Some system components, such as data storage, may require additional configuration. The following sections discuss them, as well as some of the automated features of the ftServer system that support fault-tolerant operation and ease system administration.

Setting Up Internal Disk Storage

This section discusses the following topics:

- "Internal Disk Storage Overview"
- "The Console Log and the /var/log/messages File"
- "Configuring Internal Disks"
- "Managing Partitions"
- "Default Internal Disk Configuration for a Newly Installed System"
- "Checking the Current State of the Internal Disk Subsystem"
- "Storage Device Definition"

Internal Disk Storage Overview

ftServer systems support up to six internal Serial Advanced Technology Attachment (SATA) disks, three in each CPU-I/O enclosure.

You can use RAID 1 to mirror the disks in one enclosure with the corresponding disks in the other enclosure for fault tolerance. (You should not mirror disks in the same internal storage enclosure.) RAID 1 directs I/O flow to the appropriate disks in the two CPU-I/O enclosures. When a CPU-I/O enclosure (or a disk in it) is pulled, the RAID-1 mirrors are broken, the other CPU-I/O enclosure becomes simplex, and its disk status LED becomes amber. (Pulling the CRU would cause all disk LEDs in the remaining CRU to go amber. Pulling one disk would cause only its partner's LED to go amber.) An amber LED indicates that the device is no longer safe to pull. The Opstate Manager (OSM) administratively removes the missing mirrors from their RAID 1 arrays.

When the CPU-I/O enclosure is re-inserted, the SATA driver spins up the disk(s), and the OSM administratively adds the mirrors back into their RAID arrays. The other CPU-I/O enclosure remains simplex until all mirror synchronization completes. (The RAID-1 fast resync feature can greatly reduce resync time.)

SATA, RAID, DM-multipath, and OSM plug-ins are independent modules that interact with each other.

The Console Log and the /var/log/messages File

The system console displays messages from the internal storage subsystem. This includes messages when disks are inserted and removed and when disk errors occur. The console messages are also in the system log (`/var/log/messages`). Tailing the messages file while configuring disks is very helpful.

```
# tail -f /var/log/messages
```

Since some disk-configuration operations produce considerable console output, it can be helpful to log on to another session.

Configuring Internal Disks

The six internal storage disks are persistently named based on the slot that they occupy. As shown in Figure 5-1, in the top CPU-I/O enclosure, the disks are /dev/sda, /dev/sdb, and /dev/sdc, from bottom to top. In the bottom CPU-I/O enclosure, they are /dev/sdd, /dev/sde, and /dev/sdf, from bottom to top. The name is associated with the slot, not the disk.

Figure 5-1. CPU-I/O Enclosures: Front Panel with Drive Slots Fully Populated

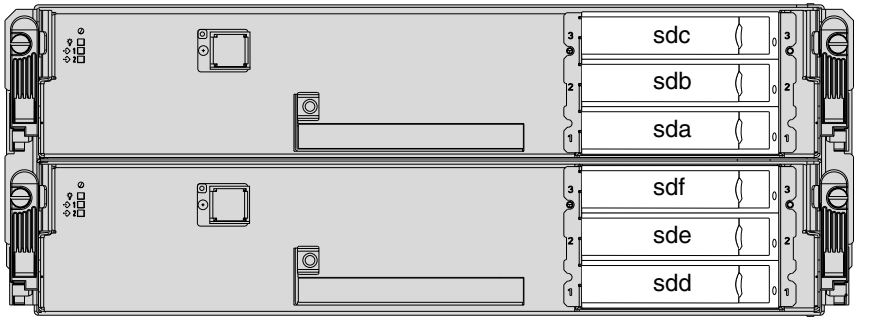

asys076

The Linux operating system allows many possible configurations of these six disks. To simplify administration and reduce confusion, this section presents recommended configurations.

For convenience, the disks are used in pairs based on vertical grouping of the disks in the CPU-I/O enclosures; sda is paired with sdd, sdb with sde, and sdc with sdf. RAID-1 arrays are created by placing one mirror on each disk of the pair. For example, RAID array `/dev/md0` occupies sda1 and sdd1.

Each SATA disk in a pair should be the same size.

In the CPU-I/O enclosure internal disk slots, insert only hard drives provided by Stratus. Inserting any other type of device may cause data loss or system failure.

Managing Partitions

You can use the `fdisk` or `sfdisk` utilities to display and change a disk's partition table and geometry (see *fdisk*(8) and *sfdisk*(8) for details). During the Linux operating system installation, all of the mirrored boot partitions are created as type 0xfd (Linux RAID

autodetect). After the installation, you use the `fdisk` utility to add data disks with the type 0x83 (Linux). Both disks of a RAID pair must have the same geometry, partition table, and type.

You can use the `fdisk` command to manage disk partitions.The following example uses the internal storage enclosure disk sdb.

To display the partition table

1. Enter the `fdisk` command.

```
# fdisk /dev/sdb
The number of cylinders for this disk is set to 17849.
There is nothing wrong with that, but this is larger than
1024,and could in certain setups cause problems with:
1) software that runs at boot time (e.g., old versions of
LILO)
2) booting and partitioning software from other OSs
   (e.g., DOS FDISK, OS/2 FDISK)
Command (m for help):
```

2. Enter the **p** argument of the `fdisk` command.

```
Command (m for help): p
Disk /dev/sdb: 146.8 GB, 146815737856 bytes
255 heads, 63 sectors/track, 17849 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Device    Boot Start     End      Blocks   Id   System
/dev/sdb1          1   17849  143372061   83  Linux
Command (m for help):
```

3. In this example, there is one partition, sdb1, that is 143,372,061 1K blocks long. Note the geometry, `255 heads, 63 sectors/track`. This is the information required when you need to adjust the geometry of a replacement partner disk to match (see ‘‘Manually Creating Partitions on Blank Disks and Adding to RAID-1 Arrays’’ on page 5-19). Enter the **q** command to quit, or continue with other commands, as required.

The following example creates a new partition table and adds a primary partition, sdb1, of type 0xfd on sdb.

To create a new partition table and add a partition

1. If `fdisk` is not already running, enter the `fdisk` command.

```
# fdisk /dev/sdb
The number of cylinders for this disk is set to 17849.
There is nothing wrong with that, but this is larger than
1024, and could in certain setups cause problems with:
1) software that runs at boot time (e.g., old versions of
LILO)
2) booting and partitioning software from other OSs
   (e.g., DOS FDISK, OS/2 FDISK)
Command (m for help):
```

2. Enter the **o** command to create a new empty DOS partition table. **Note the caution displayed by the command.**

```
Command (m for help): o
Building a new DOS disklabel. Changes will remain in memory
only, until you decide to write them. After that, of
course, the previous content won't be recoverable.

The number of cylinders for this disk is set to 17849.
There is nothing wrong with that, but this is larger than
1024, and could in certain setups cause problems with:
1) software that runs at boot time (e.g., old versions of
LILO)
2) booting and partitioning software from other OSs
   (e.g., DOS FDISK, OS/2 FDISK)
Warning: invalid flag 0x0000 of partition table 4 will be
corrected by w(rite)
Command (m for help):
```

3. Enter the **n** command to add a new partition.

```
Command (m for help): n
Command action
   e   extended
   p   primary partition (1-4)
```

4. Enter **e** or **p** to specify the desired type.

```
p
Partition number (1-4):
```

5. Enter the partition number you wish to assign (the choices depend on the type specified).

```
Partition number (1-4): 1
First cylinder (1-8924, default 1):
```

6. Enter the desired starting cylinder number for the partition, or press **Enter** to accept the default (this example accepts the default).

```
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-8924, default
8924):
```

7. Enter the desired last cylinder number for the partition, the size in megabytes or kilobytes, or press **Enter** to accept the default (this example accepts the default).

```
Using default value 8924
Command (m for help):
```

8. Enter the **t** command.

```
Command (m for help): t
Selected partition 1
Hex code (type L to list codes):
```

9. Press **Enter** to accept the default ID (83). Enter **fd** if you want the partition system ID to be Linux RAID autodetect.

```
Hex code (type L to list codes): fd
Changed system type of partition 1 to fd (Linux RAID
autodetect)
Command (m for help):
```

10. Enter the **w** command to write the partition table to the disk and exit `fdisk`.

```
Command (m for help): w
```

Default Internal Disk Configuration for a Newly Installed System

The Linux operating system is installed on the sda/sdd pair of SATA disks. See Table 5-1. The administrator must partition the remaining disks.

Table 5-1. Default Internal Storage Allocation *(Page 1 of 2)*

Restrictions	Directory / File system	Size	RAID Array	Mirrored Partitions
Required	/boot	256MB	/dev/md0	/dev/sda1 /dev/sdd1
	(swap)	2GB	/dev/md1	/dev/sda2 /dev/sdd2
	/ (root)	32GB	/dev/md2	/dev/sda3 /dev/sdd3

Table 5-1. Default Internal Storage Allocation *(Page 2 of 2)*

Restrictions	Directory / File system	Size	RAID Array	Mirrored Partitions
User configurable†	Not applicable (NA)	Extended partition to end of disk	NA	/dev/sda4 /dev/sdd4
		16 or more GB unused space		

† See Redhat Linux system administration documentation for information relating to configuration of extended disk-partition space.

NOTE

The RAID array is not deterministic.

Checking the Current State of the Internal Disk Subsystem

The `/proc/scsi/scsi` file displays the current state of the internal disk subsystem. It shows all of the disks and their names and states, as well as additional information. This example shows a system with two SATA disks.

Example 5-1. Checking the Current State of the Internal Storage Subsystem

```
# cat /proc/scsi/scsi
Attached devices:
Host: scsi0 Channel: 00 Id: 00 Lun: 00
  Vendor: ATA      Model: ST380013AS       Rev: 3.00
  Type:   Direct-Access                 ANSI SCSI revision: 05
Host: scsi4 Channel: 00 Id: 00 Lun: 00
  Vendor: ATA      Model: ST380013AS       Rev: 3.00
  Type:   Direct-Access                 ANSI SCSI revision: 05
```

Storage Device Definition

The Linux operating system automatically creates device nodes for all devices in a system.

SCSI tape drives are addressable as /dev/st* devices. Miscellaneous SCSI devices such as scanners are generally mapped as /dev/sg* devices. Note that the Linux operating system also allows some non-SCSI devices to be addressable as SCSI pseudo-devices. This can be useful to allow certain SCSI software packages to work with non-SCSI devices.

Setting Up RAID Arrays on Internal Disks

This section discusses the following topics related to the internal disk drives in the CPU-I/O enclosures:

- "RAID Array Overview"
- "Creating a RAID-1 Array"
- "Creating a RAID-0 Array"
- "Creating and Mounting a File System"
- "Checking the Current State of RAID"

RAID Array Overview

RAID is the basis for fault-tolerant file system availability. As disks come in and go out of service, the only way to keep the file system available is to mirror it on multiple disks with a disk in each CPU-I/O enclosure.

All of the file systems are created on RAID devices.

The system supports RAID-1 (mirrored) and RAID-0 (striped) on RAID-1. RAID-0 is configured using RAID-1 devices, since the underlying devices must be fault tolerant.

Each RAID array has a number (for example, `/dev/md23`) that must be unique among the running RAID arrays. The RAID array numbers are in the range 0 through 128. Device files are created for the first 128 RAID arrays. Use the `mknod` command (see *mknod*(1)) to create additional device files as needed. The number is the minor device number, and it is also used in the name.

The smaller numbers are used by the installer, so it is convenient to add new RAID arrays above 10. When RAID arrays are intended to be moved between systems, try to pick numbers that are unique among all of the systems.

The `/etc/mdadm.conf` file, which describes all of the RAID arrays for the system, is created during installation. This file contains one line for each RAID array. When

creating a new RAID array, it is convenient to copy an existing line and modify it to suit your needs. Here is a typical line:

```
ARRAY /dev/md2 super-minor=2
```

You can create, start, and stop the RAID arrays. You can manage RAID-1 arrays by adding and removing mirrors.

The following sections discuss working with RAID arrays. This includes configuring RAID arrays, administration, and tools.

NOTE

To ensure fault-tolerant operation of your system, only use the `mdadm` command to create RAID arrays.

Creating a RAID-1 Array

The example in this discussion uses the disk pair sdb and sde. The RAID array is called /dev/md20 and consists of sdb1 and sde1.

To create a RAID-1 array

1. Select a pair of same-sized disks and insert them into two corresponding slots in different CPU-I/O enclosures. This example uses the middle slot (sdb) of the top CPU-I/O enclosure, and the middle slot (sde) of the bottom enclosure.
2. When the disks have spun up, partition them for the desired RAID array (see ‘‘Managing Partitions’’ on page 5-3). You can mark the partitions with code 83.
3. Edit the `/etc/mdadm.conf` file so that the new RAID array will start each time the system boots.
   a. Use an existing ARRAY line as a model. Copy it to the bottom of the file.
   b. Edit the device and the two disks. In this example, the device is changed to md20 and the two disks are changed to /dev/sdb1 and /dev/sde1. The result should look like the following:

      ```
      ARRAY /dev/md20 level=1 num-devices=2
         devices=/dev/sdb1,/dev/sde1
      ```

      Always define RAID-1 arrays on disks in the embedded slots to have two disks and no spares.

   c. Save the file and exit the editor.

4. Make the RAID array with the following command:

   ```
   # mdadm --create /dev/md20 --level=1 --raid-devices=2
   /dev/sdb1 /dev/sde1
   ```

   This command creates the RAID array and starts it.

   NOTE

   After an array is created, `/proc/mdstat` shows that disks are in the process of resynchronizing, but the LEDs on those disks do not light. Although `/proc/mdstat` reports a resynchronization in progress, none is occurring and no disk I/O is involved. In this case, ignore the resynchronization information in `/proc/mdstat`.

You can use the `mdadm` command to see the status of the new RAID array:

```
# mdadm -Q --detail /dev/md20
/dev/md20:
        Version : 00.90.01
  Creation Time : Wed Sep 28 15:20:08 2005
     Raid Level : raid1
     Array Size : 143371968 (136.73 GiB 146.81 GB)
    Device Size : 143371968 (136.73 GiB 146.81 GB)
   Raid Devices : 2
  Total Devices : 2
Preferred Minor : 20
    Persistence : Superblock is persistent

    Update Time : Wed Sep 28 15:20:58 2005
          State : clean
 Active Devices : 2
Working Devices : 2
 Failed Devices : 0
  Spare Devices : 0

    Number   Major   Minor   RaidDevice State
       0       8       17        0      active sync   /dev/sdb1
       1       8       49        1      active sync   /dev/sde1
           UUID : 866e4ecd:12657190:79293b72:6b774c0d
         Events : 0.3
```

Creating a RAID-0 Array

When the desired file system is larger than a single disk, use RAID-0 to combine multiple RAID-1 arrays into a single RAID array.

This example assumes that two RAID-1 arrays have been created: md20, consisting of sdb1 and sde1; and md21, consisting of sdc1 and sdf1 on the sdc/sdf pair of disks.

To create a RAID-0 array

1. Select a pair of RAID-1 arrays.
2. Edit the `/etc/mdadm.conf` file so that the new RAID-0 array starts each time the system boots.
   a. Use an existing ARRAY line as a model. Copy it to the bottom of the file.
   b. Edit the device and the two disks. In this example, the device is changed to md30 and the two devices take the RAID-1 array names. The result should look like the following:

      ```
      ARRAY /dev/md30 level=0 DEVICES=/dev/md20,/dev/md21
      ```

   c. Save the file and exit the editor.
3. Make the RAID array with the following command:

   ```
   # mdadm --create /dev/md30 --level=0 --raid-devices=2 /dev/md20
   /dev/md21
   ```

   This command creates the RAID array and starts it.

If you type the `mdadm` command again, it shows the active RAID array:

```
# mdadm -Q --detail /dev/md30
/dev/md30:
        Version : 00.90.01
  Creation Time : Wed Sep 28 15:20:08 2005
     Raid Level : raid0
     Array Size : 286743808 (273.46 GiB 293.63 GB)
   Raid Devices : 2
  Total Devices : 2
Preferred Minor : 30
    Persistence : Superblock is persistent

    Update Time : Wed Sep 28 15:20:58 2005
          State : clean
 Active Devices : 2
Working Devices : 2
 Failed Devices : 0
  Spare Devices : 0
```

```
     Chunk Size : 64K

    Number   Major   Minor   RaidDevice State
      0      9       20        0      active sync   /dev/md20
      1      9       21        1      active sync   /dev/md21
           UUID : f5762850:f7dd0c84:5c720b6b:b328ca20
         Events : 0.1
```

To stop the device, use the `mdadm` command with the `-S` argument, as follows:

```
mdadm -S /dev/md30
```

Creating and Mounting a File System

The RAID arrays created in the preceding examples are raw disk block devices. You can mount a file system on the RAID array.

The following command creates an ext-3 journaled file system in the RAID-0 array created above:

```
# mkfs.ext3 /dev/md30
```

You can mount the file system on a convenient mount point, `/big_data`, as follows:

```
# mkdir /big_data

# mount /dev/md30 /big_data
```

To verify the work, use the command `df` to show the size of the file system:

```
# df /big_data
```

Use the command `ls` to show the lost+found directory in the file system:

```
# ls -l /big_data
```

At this point, be sure to add the mount to the `/etc/fstab` file, so the file system is mounted during boot. Reboot the system to make sure it works.

NOTE

A single disk that is not part of a RAID array must have a file system mounted on it to ensure that its operation state is reported correctly.

Checking the Current State of RAID

The `mdstat` file displays the current state of RAID. It shows all running RAID arrays and their current status, including which mirrors are present, whether they are synchronized, and more. See Example 5-2.

NOTE

The device names displayed in `/proc/mdstat` are the kernel names for each device. These are different from the user device names displayed by the `mdadm` command.

Example 5-2. Checking the Current State of RAID

```
# cat /proc/mdstat
Personalities : [raid0] [raid1]
md30 : active raid0 md21[1] md20[0]
   286743808 blocks 64k chunks

md21 : active raid1 sdf1[1] sde1[0]
   143371968 blocks [2/1] [_U]

md20 : active raid1 sdd1[1] sdb1[0]
   143371968 blocks [2/2] [UU]

md1 : active raid1 sdc2[1] sda2[0]
   2096384 blocks [2/2] [UU]

md2 : active raid1 sdc3[2](F) sda3[0]
   31647936 blocks [2/1] [U_]

md0 : active raid1 sdc1[1] sda1[0]
   2096384 blocks [2/2] [UU]

unused devices: <none>
```

In this example, md30 is a RAID-0 array made up of md20 and md21. The remaining stanzas are for RAID-1 arrays. Note that md21 is operating in degraded mode. It is missing a disk: `[2/1] [_U]`.

Removing and Replacing Internal Disks

For disk fault tolerance, disk mirrors must be maintained when disks or CPU-I/O enclosures are removed and replaced. For information about the recommended partnering confirmation for internal disks, see ‘‘Configuring Internal Disks’’ on page 5-3.

This section discusses the following topics:

- ‘‘Disk Insertion”
- ‘‘Administering RAID Arrays on Internal Disks”

Disk Insertion

When you reinsert a pulled disk, the OSM storage plugin attempts to match it with an existing disk. If it finds a match, it hot-adds the mirror partitions on the inserted disk back into the existing RAID arrays and resynchronizes them (see ‘‘Resynchronization’’ on page 5-16).

Similarly, if you replace a failed disk, the OSM plugin automatically adds the replacement disk to a running RAID array.

Administering RAID Arrays on Internal Disks

This section discusses the following topics:

- ‘‘To Stop a RAID Array and Move It to Another System”
- ‘‘Errors and Faulty Mirrors”
- ‘‘Removing a Faulty Mirror”
- ‘‘Resynchronization”
- ‘‘Replacing a Failed Disk”
- ‘‘Manually Creating Partitions on Blank Disks and Adding to RAID-1 Arrays”

You can use the `mdadm` command to administer RAID arrays. The following sections provide examples of how to perform some common administrative procedures using `mdadm`.

NOTE

Never remove both member disks of a RAID-1 array. The Linux operating system does not support that operation.

To Stop a RAID Array and Move It to Another System

You can stop RAID-0 and RAID-1 arrays, if they are not in use. Unmount the file system (if one is mounted) and stop the array as shown in the following example:

```
# umount /dev/md30
# mdadm -S /dev/md30
```

Before physically removing the disks from the system, check that the RAID array no longer appears in `/proc/mdstat`. Edit the `/etc/fstab` and `/etc/mdadm.conf` files on the current system to delete it. Edit the information into the files in the new system. If the new system already has the device in use, you cannot start the RAID array.

You can start a RAID array when it is stopped. Use the following command to start a RAID array that was already configured in `/etc/mdadm.conf`:

```
# mdadm -A /dev/md30
```

Errors and Faulty Mirrors

When an error is reported on a mirror, the mirror is marked faulty and it is no longer used. The last active mirror is never marked faulty even if errors are reported against it.

The SCSI subsystem (which comprises the SCSI midlayer and the low-level drivers) returns errors to RAID software when its error management code determines that the I/O request cannot succeed. When this happens, RAID software marks the mirror faulty and retries reads on another mirror but otherwise ignores write errors.

Every time an active disk is pulled, all outstanding I/O is returned as errors. When a disk is pulled, all RAID members or mirrors that have active I/O on the missing disk are marked faulty. This may not be all RAID arrays that use the disk.

When a mirror becomes faulty, the disk and CPU-I/O enclosure of the remaining active mirror becomes simplex and that enclosure is no longer safe to pull.

If you want to remove a good disk from a RAID array, you must mark it faulty. You can use the `mdadm` command to simulate an error and mark the mirror faulty.

```
# mdadm /dev/md20 -f /dev/sdb1
```

The `/proc/mdstat` file shows an `F` after a faulty mirror in the display:

```
md20 : active raid1F sdc1[1] sdf1[0]F
   4095872 blocks 0 active chunks [2/1] [_U]
```

Removing a Faulty Mirror

Before removing a mirror, check the `/proc/mdstat` file to make sure that it is marked faulty. You can use the `mdadm` command to remove a faulty mirror from a RAID array, as shown in the following example:

```
# mdadm /dev/md20 -r /dev/sdb1
```

In the preceding example, the mirror (/dev/sdb1) is removed from the RAID array (/dev/md20). After running this command, the `/proc/mdstat` file shows the RAID array without the mirror.

You cannot remove a mirror that is not faulty. Before a disk is pulled, all of the mirrors on it must be marked faulty (either by the operating system or with the `mdadm -f` command) and removed with the `mdadm -r` command before the system can completely remove the disk (the OSM storage plugin automates these tasks). This means that until all mirrors are removed, a replacement disk inserted in the same slot will not spin up.

You can use the `mdadm` command to add a mirror into a running RAID array. The following example shows how to do this.

```
# mdadm /dev/md20 -a /dev/sdb1
```

In the preceding example, `/dev/md20` is the RAID array and `/dev/sdb1` is the *mirror*. After running this command, the `/proc/mdstat` file shows the RAID array with the new mirror as a spare. The command starts resynchronization of the mirror. Resynchronization can take a while, depending on how much data must be written to the new mirror. Resynchronization must finish before the add operation is complete.

Resynchronization

There are two classes of resynchronization: fast and full. A *fast resync* is done when the mirror was recently part of the array and RAID software knows what has been written while the mirror was missing. If the fast resync is not possible (for example, when a clean new disk is added), a full resync is performed. The *full resync* synchronizes all of the mirror that is in use. If nothing was written while the mirror was missing, no resync is needed.

Synchronization begins automatically when the mirror is added to a RAID array. RAID software limits one resync per disk at a time so that, for example, on the system disks, there may be one active resync and several delayed resyncs. Example 5-3 provides an example of resynchronization.

Example 5-3. Resynchronization

```
# cat /proc/mdstat
Personalities : [raid1]
md21 : active raid1 sde1[0] sdf1[1]
       104320 blocks [2/1] [U_]
       resync=DELAYED

md20 : active raid1 sdc1[2] sdd1[0]
       33640000 blocks [2/1] [U_]
       [===>.................]  recovery = 19.0%
       (6400000/33640000) finish=2.1min
       speed=206451K/sec

md1 : active raid1 sdb2[1] sda2[0]
      2096384 blocks [2/2] [UU]

md2 : active raid1 sdb3[1] sda3[0]
      31647936 blocks [2/2] [UU]
md0 : active raid1 sdb1[1] sda1[0]
      104320 blocks [2/2] [UU]

unused devices: <none>
```

NOTE

The device names displayed in `/proc/mdstat` are the kernel names for each device. These are different from the user device names displayed by the `mdadm` command.

As long as there is a missing mirror or a resynchronization in process, RAID and the CPU-I/O enclosure are simplex for the active mirror.

Replacing a Failed Disk

When you need to replace a failed disk, the OSM plugin can automatically add the replacement disk to a running RAID array, provided the following conditions exist:

- The replacement disk must be blank, as defined by the current *safe mode* setting. If safe mode is active, zero the disk's partition table and RAID superblocks. Then remove and reinsert the disk to start the automatic disk replacement. For more information about safe mode, see “Configuring Safe Mode.”
- Do not reboot the system or stop and restart OSM after you remove the failed disk until you have inserted the replacement disk and it has synchronized with its partner. The information necessary to perform automatic disk replacement is not

persistent, so if OSM is restarted, the replacement disk must be paired using a different method.

- The failed disk must have been paired with one (and only one) partner disk. For example, if /dev/md4 consisted of partitions sda1 and sdb1, and /dev/md5 consisted of sdb2 and sdc2, automatic disk replacement would not work for disk sdb. In addition, partition numbers on the failed disk and its partner, for any partitions belonging to RAID1 arrays, must be the same.
- The failed disk must belong to a RAID 1 on top of a disk, partition, or multipath. If the failed disk belongs to a RAID0 (even if that RAID0 is part of a RAID1), the blank disk will not be added to the RAID array.

NOTE

If the running member of the RAID array was a system disk, the bootloader (grub) is added to the newly-inserted disk.

To replace a failed disk

1. While the system and RAID array are running, remove the failed disk.
2. Insert a blank disk. The blank disk is automatically added to the array.

Configuring Safe Mode

By default, the OSM configuration file, `/opt/ft/osm/config.xml`, configures this automatic pairing of disks in *safe-mode*. In safe mode, a newly inserted disk is considered blank if it has no valid partition table or RAID superblocks. Otherwise, a newly inserted disk is considered blank even if it has a valid partition table or RAID superblocks, as long as it **does not** belong to a running RAID array.

The following entry in the OSM configuration file configures safe mode:

```
<entry key="blankDiskSafeMode" value="true"/>
```

To specify that automatic pairing of disks **not** operate in safe-mode, replace the word `true` with `false` so that the file contains the following entry:

```
<entry key="blankDiskSafeMode" value="false"/>
```

Because OSM preferences are only read at start-up time, you must reboot the system for a change to take effect.

Manually Creating Partitions on Blank Disks and Adding to RAID-1 Arrays

When you create a backup system disk (which is described in ‘‘Creating a Backup System Disk’’ on page 4-9), you must create its partitions, use the `mdadm` command to add it to a RAID-1 array, and run the `grub` command.)

NOTE

Do not perform these procedures if you are replacing a failed disk. See, instead, “Replacing a Failed Disk” on page 5-17.

- “Replacing Defective Disks Interactively”
- “Replacing Defective Disks Manually”

Replacing Defective Disks Interactively

To replace a defective disk, perform the following procedures:

- Remove the defective disk and insert a spare disk.
- Run the duplex_blank_disk command (see “The `duplex_blank_disk` Command” on page 5-23).

NOTE

Replacement disks can be new, factory-fresh disks or disks recycled from other systems. Care must be taken with recycled disks. The partition table and RAID superblocks that exist on the disk can confuse the system.

Replacing Defective Disks Manually

To replace a defective disk by manually issuing commands for each step of the process, perform the following procedures:

- Remove the defective disk and insert a spare disk.
- Verify that the spare disk is not in use.
- Zero the spare disk.
- Partition the spare disk to match the running disk.
- Add partitions on the spare disk to RAID-1 arrays.
- Run the GRUB boot loader (only if the running disk is the system disk).

NOTE

Replacement disks can be new, factory-fresh disks or disks recycled from other systems. Care must be taken

with recycled disks. The partition table and RAID superblocks that exist on the disk can confuse the system.

To remove a defective disk and insert a spare disk

1. Remove the defective disk from any RAID arrays that it belongs to, using the instructions presented in “Removing a Faulty Mirror” on page 5-16.
2. Physically remove the disk.
3. Insert a spare disk in the slot previously occupied by the defective disk. The drive spins up automatically.

To verify that the spare disk is not in use

Type the following commands and check the resulting output:

```
# mdadm --detail --scan

# swapon -s

# cat /etc/mtab
```

To zero the spare disk

Perform **one** of the following procedures:

- Zero the spare disk’s RAID superblocks by typing a command such as the following for each partition on the spare disk (substitute the device node of the partition you wish to zero for sdb1 in this example):

  ```
  # mdadm --zero-superblock /dev/sdb1
  ```

  NOTE

  Zeroing the disk’s RAID superblocks takes very little time but may not remove everything from the disk. If you are concerned about this, zero the entire disk as described in the following step.

- Zero the entire spare disk by typing a command such as the following (substitute the device node of the disk you wish to zero for `sdf` in this example).

  ```
  # dd if=/dev/zero of=/dev/sdf bs=1024k
  ```

  NOTE

  Zeroing the entire disk takes a long time but removes everything from the disk, eliminating many problems.

To partition the spare disk to match the running disk

1. Save the partition table of the running disk to a file with a command like the following:

   ```
   # sfdisk -d /dev/sda > sda_partition_table
   ```

2. Write the saved partition table to the spare disk with a command like the following:

   ```
   # sfdisk /dev/sdd < sda_partition_table
   ```

   Occasionally, `sfdisk` returns the following error while writing the saved partition table to the spare disk:

   ```
   Checking that no-one is using this disk right now ...
   BLKRRPART: Input/output error
   ```

   This error indicates that the disk is currently in use, so you should not repartition it. Perform these steps to correct this error:

   a. Unmount all file systems.
   b. Swap off all swap partitions on this disk.
   c. Use the `--no-reread` flag to suppress this check.
   d. Use the `--force` flag to overrule all checks.

   NOTE

   The preceding error does not occur if the spare disk already contained a valid partition table.

   If you are sure that the spare disk is not in use, force `sfdisk` to write the partition table by using the `--no-reread` flag.

To add partitions on the spare disk to RAID-1 arrays

1. Type the following command to determine which RAID-1 arrays the running disk belongs to:

   # **mdadm --detail --scan**

   ```
   ARRAY /dev/md2 level=raid1 num-devices=2
      UUID=5ddb14c7:d5e0b2d6:ad80086d:8db2a245
   devices=/dev/sda2

   ARRAY /dev/md1 level=raid1 num-devices=2
      UUID=3838df6e:60caf7e6:695d0f62:de94e821
   devices=/dev/sda3

   ARRAY /dev/md0 level=raid1 num-devices=2
      UUID=3e4ad330:c8ee5dfc:f48bd88a:401ada25
   devices=/dev/sda1
   ```

2. Add each partition on the spare disk to the RAID-1 array containing the corresponding partition on the running disk with commands like the following:

   ```
   # mdadm -a /dev/md0 /dev/sdd1
   mdadm: hot added /dev/sdd1

   # mdadm -a /dev/md1 /dev/sdd3
   mdadm: hot added /dev/sdd3

   # mdadm -a /dev/md2 /dev/sdd2
   mdadm: hot added /dev/sdd2
   ```

Perform the following procedure only if the running disk is the system disk.

To run GRUB

If the running disk is the system disk, run the GRUB boot loader on the boot partition of the spare disk, after resynchronization is complete on that partition. Example 5-4 shows a typical use of GRUB:

Example 5-4. Running GRUB

```
# /sbin/grub
       GNU GRUB  version 0.95  (640K lower / 3072K upper memory)
[ Minimal BASH-like line editing is supported. For the first word, TAB
   lists possible command completions. Anywhere else TAB lists the possible
   completions of a device/filename.]
grub> device (hd0) /dev/sdd
grub> root (hd0,0)
 Filesystem type is ext2fs, partition type 0xfd
grub> setup (hd0)
 Checking if "/boot/grub/stage1" exists... no
 Checking if "/grub/stage1" exists... yes
 Checking if "/grub/stage2" exists... yes
 Checking if "/grub/e2fs_stage1_5" exists... yes
 Running "embed /grub/e2fs_stage1_5 (hd0)"...  16 sectors are embedded.
 succeeded
 Running "install /grub/stage1 (hd0) (hd0)1+16 p (hd0,0)/grub/stage2
/grub/grub.conf"..
. succeeded
Done.
grub> quit
```

On your own system, replace the `/dev/sdd` shown in Example 5-4 with the device node for your spare disk. In the `root (hd0,0)` command, the second zero is the number of the partition to GRUB. GRUB partitions are zero-based rather than one-based, so these commands actually indicate that partition 1 on `/dev/sdd` has been GRUBbed.

The `duplex_blank_disk` Command

The `duplex_blank_disk` command prompts you for all of the information required to pair a spare disk with a running disk. You can run it by typing:

```
# /opt/ft/bin/duplex_blank_disk
```

In Example 5-5, the command prompts you for information that is needed to pair a spare internal disk with the running system disk.

Example 5-5. Pairing a Spare Internal Disk with the Running System Disk

```
# /opt/ft/bin/duplex_blank_disk

Device Path ID of blank disk (e.g. 10/40/1 or 70/1): 11/40/1

Device node(s) for 11/40/1: /dev/sdd

Is this the correct blank disk device? (yes/no) y

Device Path ID of source disk (e.g. 10/40/1 or 70/1): 10/40/1

Device node(s) for 10/40/1: /dev/sda

Is this the correct source disk device? (yes/no) y

Source disk is partitioned: partitioning blank disk to match.

Source disk partition 1 belongs to RAID 1 /dev/md0.

Adding blank disk partition 1 to RAID 1 /dev/md0.
mdadm: hot added /dev/sdd1

Source disk partition 2 belongs to RAID 1 /dev/md2.

Adding blank disk partition 2 to RAID 1 /dev/md2.
mdadm: hot added /dev/sdd2

Source disk partition 3 belongs to RAID 1 /dev/md1.

Adding blank disk partition 3 to RAID 1 /dev/md1.
mdadm: hot added /dev/sdd3

Waiting for resync to complete before grubbing /dev/sdd1.

Grubbing /dev/sdd1
```

Setting Up External ftScalable Storage

ftServer systems also support external ftScalable Storage arrays. See the following manuals for information about installing, configuring, and administering ftScalable Storage arrays:

- *ftScalable Storage: Getting Started Guide* (R601)
- *ftScalable Storage: Operation and Maintenance Guide* (R600)
- *ftScalable Storage: Commands Reference Manual* (R599)

System Backup and Disaster Recovery

Your ftServer system provides many safeguards against losing data due to hardware failures. However, it cannot cover all contingencies, so **it is still important to perform regular backups and enact a good disaster-recovery program**.

Ethernet Devices

Network interface naming on ftServer systems running a supported Linux distribution together with ftSSS is different from that on other Linux systems. On ftServer systems, physical devices are given names corresponding to their hardware location. After installing ftSSS, the interfaces associated with the Ethernet adapters are operational. Multiple interfaces can be bonded together in a channel-bonding interface.

This section discusses the following topics:

- "Physical Device Naming"
- "Monitoring and Configuring Channel-Bonding Interfaces"
- "MAC Addresses"

Physical Device Naming

On many Linux systems, Ethernet devices are normally assigned names based on the order of discovery at system startup. The names begin with the letters "eth," followed by a number starting with 0 and counting up. This is convenient because the first (and often only) device on a host is predictably named eth0 and can be configured without detailed knowledge of the device type.

On an ftServer system, configuration may change dynamically when hardware failures occur, repairs are made, or when an administrator adds or removes components. Creating new Ethernet device names when new hardware is installed, tracking the name of an device while it is removed and replaced and matching it up again, or deleting the name would be difficult and the results confusing.

Instead, ftSSS assigns to network devices names that are derived from their physical location in the system. Table 5-2 shows the names of the embedded Ethernet devices in ftServer CPU-I/O enclosures.

Table 5-2. Ethernet Devices in ftServer CPU-I/O Enclosures

Device	Location	Ethernet Interface Device Name
Embedded 10/100/1000-Mbps Ethernet PCI adapter	CPU-0, I/O-10, slot 5, port 0	eth000010
Embedded 10/100/1000-Mbps Ethernet PCI adapter	CPU-0, I/O-10, slot 5, port 1	eth000011
Embedded 10/100/1000-Mbps Ethernet PCI adapter	CPU-1, I/O-11, slot 5, port 0	eth080010
Embedded 10/100/1000-Mbps Ethernet PCI adapter	CPU-1, I/O-11, slot 5, port 1	eth080011
Ethernet PCI adapter	CPU-0, I/O-10, PCI slot 9, port 0	eth000008
Ethernet PCI adapter	CPU-0, I/O-10, PCI slot 9, port 1	eth000009
Ethernet PCI adapter	CPU-0, I/O-10, PCI slot 10, port 0	eth000218
Ethernet PCI adapter	CPU-0, I/O-10, PCI slot 10, port 1	eth000219
Ethernet PCI adapter	CPU-0, I/O-10, PCI slot 11, port 0	eth000220
Ethernet PCI adapter	CPU-0, I/O-10, PCI slot 11, port 1	eth000221
Ethernet PCI adapter	CPU-1, I/O-11, PCI slot 9, port 0	eth080008
Ethernet PCI adapter	CPU-1, I/O-11, PCI slot 9, port 1	eth080009
Ethernet PCI adapter	CPU-1, I/O-11, PCI slot 10, port 0	eth080218
Ethernet PCI adapter	CPU-1, I/O-11, PCI slot 10, port 1	eth080219
Ethernet PCI adapter	CPU-1, I/O-11, PCI slot 11, port 0	eth080220
Ethernet PCI adapter	CPU-1, I/O-11, PCI slot 11, port 1	eth080221

Monitoring and Configuring Channel-Bonding Interfaces

By default, the physical Ethernet interfaces listed in Table 5-2 are bound together into two channel-bonding interfaces, called bond0 and bond1. The two channel-bonding interfaces are set to operate in active-backup mode (mode 1) with Dynamic Host Configuration Protocol (DHCP) enabled.

In many cases, no additional configuration is necessary. However, you may want to change the default configuration to better meet your particular networking requirements.

You configure and administer the Ethernet interfaces on your ftServer system just as you would on any standard Linux system. Additionally, you can use the `ftsmaint` command to obtain information about the fault-tolerant status of the interfaces.

This section discusses the following topics:

- "Monitoring Channel-Bonding Interfaces"
- "Configuring Channel-Bonding Interfaces"
- "Determining Interface Device Names"

Monitoring Channel-Bonding Interfaces

You can monitor the fault-tolerant status of channel-bonding interfaces by using the `ftsmaint` command. Example 5-6 shows the default configuration of the embedded Ethernet devices:

Example 5-6. Default Configuration of Embedded Ethernet Devices

```
# /opt/ft/bin/ftsmaint lsVnd
Virtual Network Device (VND) Groups
===================================

Group Name   Status  Inet Address    RX Errors  TX Errors  Collisions
-------------------------------------------------------------------
bond0        ONLINE  134.111.78.103          0          0           0
bond1        ONLINE  192.168.4.10            0          0           0
bond2        OFFLINE -                       0          0           0
bond3        OFFLINE -                       0          0           0
bond4        OFFLINE -                       0          0           0

VND Group Members
=================

Member           Group Name  Status     Interface  Link State  Link Speed
-------------------------------------------------------------------------
eth000010        bond0       DUPLEX     UP         -           -
eth000011        bond1       DUPLEX     UP         -           -
eth080010        bond0       DUPLEX     UP         -           -
eth080011        bond1       DUPLEX     UP         -           -
```

In Example 5-6, there are two online channel-bonding interfaces (masters), bond0 and bond1, each composed of two physical interfaces (slaves). The output shows the four physical slave interfaces in the system and also shows their status and the name of the bond to which they belong. Note that three other channel-bonding interfaces are defined by default, but they are not configured and are therefore offline.

You can monitor additional information about the currently installed channel-bonding and physical interfaces by running the Linux `ifconfig` tool.

Configuring Channel-Bonding Interfaces

You configure and administer channel-bonding interfaces using standard Linux procedures. Configuration of channel-bonding interfaces is controlled by the `ifcfg-*` files in the `/etc/sysconfig/network-scripts` directory. You modify existing channel bond configurations by editing the bond's `ifcfg-bond`*n* file or the bond's slaves' interface `ifcfg-eth*` files. Additionally, you use standard Linux network utilities like `ifdown`, `ifup`, `service`, `ifconfig`, `ip`, and `route`.

By default, the system supports five channel-bonding interfaces. In the unlikely event that you must configure more than five bonding master interfaces (up to a maximum of 10), modify the options line in the `/etc/modprobe.d/ft-network.conf` file so that `max_bonds=`*x*, where *x* is the desired number of channel-bonding interfaces. Perform the following procedure to activate configuration-file modifications.

To activate configuration-file modifications

1. Type the `ifdown` command to stop all network interfaces.
2. Type the `rmmod` command to unload the bonding kernel module. This disables network access.
3. Type the `modprobe` command to reload the bonding module. This enables network access with the newly designated number of channel-bonding interfaces.
4. If necessary, type an `ifup bond`*N* command for each bonded interface you wish to restart.

All channel-bonding interfaces must operate in the same mode. If you want to change the mode from the default setting (mode 1, active-backup mode), modify the options line in the `/etc/modprob.d/ft-network.conf` file so that `mode=`*x*, where *x* is the desired mode of operation. To activate configuration-file modifications, perform the preceding procedure.

NOTES

1. There must be at least one alias for an active bond in the `/etc/modprobe.d/ft-network.conf` file, or bonding cannot occur.
2. The `/etc/modprobe.d` directory should contain no more than one `ft-network.conf` file.

Determining Interface Device Names

When you add a PCI Ethernet adapter to a system, you must determine the device names of the physical interfaces on the adapter before you can configure it. See Figure 7-1 and Figure 7-2 for the hardware device path of PCI Ethernet adapters, and then refer to Table 5-2 to determine the device names.

When installing a PCI adapter, you can use the output of the `lspci` command to confirm its Ethernet interface device name. By comparing the output of `lspci` before and after installing the adapter, you can identify the newly displayed output that corresponds to the newly installed adapter. The output might look like the following:

```
04:02.0 Ethernet controller: Intel Corp. 82546EB Gigabit
Ethernet Controller (Copper) (rev 01)
```

The device names for the two physical interfaces on the newly added, dual-Ethernet adapter shown above are `eth000008` and `eth000009`.

CAUTION

Do not issue the `lspci` command with the `-xxx` option. Doing so may temporarily interfere with the fault-tolerant operation of the system.

To add two physical interfaces and configure a new channel-bonding interface

1. Install the first Ethernet PCI adapter in a supported slot in one CPU-I/O enclosure.

   NOTE

   When adding a pair of Ethernet PCI adapters to the system, be sure to install one adapter in the top CPU-I/O enclosure and the other in the bottom CPU-I/O enclosure, in same-numbered slots. This is necessary to maintain fault tolerance. If you install both devices in the same enclosure, and that enclosure fails, you will lose connectivity.

2. Determine the interface device name for each physical Ethernet interface on this newly installed adapter. See ‘‘Determining Interface Device Names’’ on page 5-29 for details.

   You must add this device name to the physical interface’s configuration file (see step 5).

3. Repeat steps 1 and 2 for the second adapter in the corresponding slot in the second paired CPU-I/O enclosure.
4. Create a new bond file (for example, `ifcfg-bond2`) in the `etc/sysconfig/network-scripts` directory. Use the contents of an existing `bond`*n* file as a guide.
5. Create two new physical interface configuration files for the two new physical interfaces. Use the contents of an existing `ifcfg-eth*` file as a guide. Be sure to use the device names of the newly installed adapters (see step 2).
6. Type the following command to bring up the new interface:

   ```
   # ifup bond2
   ```

MAC Addresses

You can use the `ifconfig` command to determine the current MAC address of an Ethernet interface. Alternatively, you can examine the interface's address file in the `sys/class/net/`*`interfacename`* directory.

For the embedded Ethernet adapter interfaces, Ethernet MAC addresses are algorithmically generated from a *base* MAC address assigned to the machine as a whole. Each physical device receives a different MAC address. Because of this, a channel-bonding interface (and all of its physical devices) may get a different MAC address from one reboot to the next, depending on which physical device is available first (based on which CPU-I/O enclosure is used during the boot).

Other System Configuration Information

In addition to setting up storage and network devices, you may want to perform the following tasks to set up your system:

- Install and configure an ActiveService Network (ASN) modem and software to allow monitoring of system alarms and remote troubleshooting. See the *Stratus ActiveService Network Configuration Guide* (R072) for instructions.
- Configure ftISNMP to allow remote management of your ftServer system. See Chapter 8 for instructions.
- Disable hyperthreading.

Information about video display settings is also included.

You also need to perform the following configuration tasks, using standard Linux procedures:

- Configuring the IP address for the bond0 and bond1 interfaces (static or DHCP, and gateway in `/etc/sysconfig/network-scripts/ifcfg-bond0` and `/etc/sysconfig/network-scripts/ifcfg-bond1`)
- Configuring DNS resolution for the system (`/etc/nsswitch.conf` and `/etc/resolv.conf`)
- Configuring static routes for the system (`/etc/sysconfig/static-routes`)
- Configuring the system hostname (`/etc/hosts` and `/etc/sysconfig/network`)
- Configuring the system time zone (`/etc/sysconfig/clock`)

Disabling Hyperthreading

Some system installations may disable hyperthreading to facilitate application execution.

CAUTION

Disable hyperthreading only if your system satisfies the minimum BIOS version requirements. Contact the CAC or your authorized Stratus service representative to confirm that you have the correct BIOS version.

To disable hyperthreading, you enter the ftServer Setup utility.

To disable hyperthreading

1. Turn on or restart your system. When the Stratus ftServer logo screen appears, press F2 to enter your system's BIOS setup utility.

   The BIOS setup utility's **Main** menu appears after the system completes more of the POST (power-on self-test) process.

2. On the Main menu, use the right-arrow key to select the **Advanced** tab.
3. Use the down-arrow key to select **Advanced Processor Options** and press Enter.
4. Select **Hyper Threading Technology** and press Enter.
5. Change the value from Enabled to **Disabled** and press Enter.
6. Press **Esc**, then select **Exit Saving Changes** and press Enter.

   The system resumes booting.

Configuring the System Video Display

Your ftServer system's video comes configured by default. There is normally no need to change the video displays settings, and the system is strictly limited in some of its parameters. For instance, the screen resolution is limited to 1024x768 pixels. However, it is possible, though not advisable, to change the video configuration.

CAUTION

Using other means to configure the video —including any of those available from a Red Hat icon—may result in loss of system fault tolerance, or may cause the system to boot only in text mode, or may return an error message.

If you do alter video settings, change them only by using a text editor to change the entries in the `/etc/X11/xorg.conf` file. Use one of the configured video modes that the ftSSS installation program installed in the `xorg.conf` file.

To do this, select one of the available modes for the pixel depth you are using by putting the desired mode first in the list for that depth. Changes made to `xorg.conf` file are preserved during an upgrade.

Managing the System Clock

You may see the following message after the system boots or after you attempt to use the `system-config-time` utility:

```
Cannot access the Hardware Clock via any known method.
Use the --debug option to see the details of our search for
an access method.
```

The message occurs when you are trying to use Coordinated Universal Time (UTC). Instead, use the Network Time Protocol (NTP).

The following message does not indicate a problem with the system. The clock will be properly reset and you can safely ignore the message.

```
Losing some ticks... checking if CPU frequency changed
```

Additional Documentation and Resources

Linux System Administrator's Guide v0.8, Linux Documentation Project:

http://www.ibiblio.org/pub/Linux/docs/linux-doc-project/system-admin-guide/

http://unthought.net/Software-RAID.HOWTO/

Managing RAID on Linux, Derek Vadala, O'Reilly & Associates, 2003:
http://www.oreilly.com/catalog/mraidlinux/

Chapter 6
Managing Data Storage Devices

This chapter discusses the following topics:

- “CD-ROM Drives”
- “SCSI Tape Drives”
- “USB Storage Devices”
- “Additional Resources”

Chapter 5 explains basic storage device definition and the configuration and management of the internal disk drives embedded in CPU-I/O enclosures. This chapter briefly discusses other data storage devices that are included with or can be optionally attached to the system.

CD-ROM Drives

An ftServer system supports two CD-ROM or DVD+RW drives, which may appear as the following names on the system, depending on whether a CD-ROM or DVD writer is in the drive:

- In CPU-0, I/O-10: /dev/cdrom, /dev/dvdwriter, /dev/hde
- In CPU-1, I/O-11: /dev/cdrom1, /dev/dvdwriter1, /dev/hdi

It is good practice to remove media from the drive and unmount the drive as soon as the CD or DVD device is no longer being used.

Typically, if a CPU-I/O enclosure fails or is taken out of service, its CD or DVD drive will again be accessible when the enclosure is brought back into service. However, the CD or DVD device will **not** be accessible after the enclosure is brought back into service if either of the following conditions are true:

- The CD or DVD device contains media which is open by an application when the CPU-I/O enclosure fails or is removed from service
- The CD or DVD device is mounted when the CPU-I/O enclosure fails or is removed from service

If the CD or DVD drive is inaccessible after its CPU-I/O enclosure is brought back into service, you must reboot the system to regain access to the drive.

SCSI Tape Drives

ftServer systems running a supported Linux distribution together with ftServer System Software for the Linux Operating System (ftSSS) support several optional tape drives. See the installation guide for your system for information on connecting tape drives to ftServer systems. Also, see the *Stratus ftServer Systems Peripherals Site Planning Guide* (R582) for information about supported tape drives and enclosures.

Autoloader tape drives may require configuration into an operational mode that is fully addressable by applications through switch settings on the drive. The operation manual for the drive should provide you with needed configuration information.

USB Storage Devices

USB storage devices, including floppy-disk and solid-state storage, are supported through the SCSI driver. These devices appear as SCSI devices. You can get information about these devices by examining the file `/proc/scsi/scsi`, by running the command `lsusb`, or by examining the system log (`/var/log/messages`).

When you connect a USB device to the USB bus, the SCSI driver scans it once and assigns a name (for example, sde). This is the device's internal name that is displayed by commands such as `/proc/mdstat` and in the system log. When mounting and unmounting a USB device, **do not use this name**. Instead, use the name assigned by the `udev` command. This name has the format `sd*usb`. For example, `sd1usb` is the name of the device attached to port 1 of the root USB hub.

If a device is plugged into a USB hub, the name has two numbers. For example, `sd1.3usb` is the name of the device attached to port 3 of a hub connected to port 1 of the root USB hub. If you add another hub to the chain, the device name would contain a third number.

The `udevinfo` command translates the internal name into the name assigned by the `udev` command. For example, the following command and output show that the `udev` command has assigned the name sd1usb to the device with the internal name sde:

```
# udevinfo -q name -p /sys/block/sde
sd1usb
```

The `ls` command displays the device node for this device:

```
# ls -l /dev/sd1usb
brw-rw---- 1 root disk 8, 64 Oct 24 16:18 /dev/sd1usb
```

You can also translate the name assigned by the `udev` command into the internal name. For example:

```
# ls -l /dev/sd*usb
brw-rw---- 1 root disk 8, 64 Oct 24 16:18 /dev/sd1usb
```

Once you have the name assigned by the `udev` command (in this case, sd1usb), you can use `udevinfo` to find the internal name (in this case, sde):

```
# udevinfo -q path -n /dev/sd1usb
/block/sde
```

For more information about the `udevinfo` command, see *udevinfo*(8).

CAUTION

Before unplugging the device, make sure that it is not being used (the usage count is 0). If a file system is mounted, unmount it (and make sure the `umount` command completes) before unplugging the device. The `umount` command flushes any buffered pages back to the device, so failing to wait for `umount` to complete can cause data corruption.

NOTES

1. An important consequence of the fact that the SCSI subsystem scans USB devices only on connection is that simply removing a floppy disk from the floppy drive or inserting a disk does not cause a rescan. You must unplug the floppy drive and plug it back in to cause a rescan.
2. During failovers, access to USB storage devices is not robust.
3. If an AC switch occurs, some USB devices are not properly reset and disappear from `ftsmaint` output. To reset these devices, you must unplug them and then plug them back in.
4. When you unplug a device and later plug it back in, the internal name may change if the SCSI subsystem has added other devices while it was removed. The name assigned by the `udev` command, however, does not change.

If the active CPU-I/O enclosure fails over to the other enclosure, a mounted USB disk-drive device may become unusable. If this happens, remove the device from the system and then insert it back in the system.

USB storage devices are not bootable devices.

Most floppy disks and solid-state devices come with a virtual file allocation table (VFAT) file system. You can create ext-2 or other file systems on the device as well. You can mount them on a convenient mount point, for example:

```
# mkdir /mnt/floppy
# mount /dev/sdg1 /mnt/floppy
```

USB Floppy Drives

The USB floppy drive appears as follows in `/proc/scsi/scsi`.

```
Host: scsi0 Channel: 00 Id: 00 Lun: 00
  Vendor: ATA      Model: ST380013AS     Rev: 3.00
  Type:   Direct-Access                  ANSI SCSI revision: 05
Host: scsi4 Channel: 00 Id: 00 Lun: 00
  Vendor: ATA      Model: ST380013AS     Rev: 3.00
  Type:   Direct-Access                  ANSI SCSI revision: 05
```

The system log shows details about the device:

```
scsi4 : SCSI emulation for USB Mass Storage devices
  Vendor: NEC       Model: USB UF000x        Rev: 1.50
  Type:  Direct-Access                 ANSI SCSI revision: 02
SCSI device sdaz: 2880 512-byte hdwr sectors (1 MB)
sdaz: Write Protect is on
```

`Write Protect is` on indicates that the disk is read-only.

USB Solid-State Devices

The following is an example of the `/proc/scsi/scsi` display for a solid-state device.

```
Host: scsi5 Channel: 00 Id: 00 Lun: 00
  Vendor: LEXAR    Model: JUMPDRIVE SECURE Rev: 3000
  Type:  Direct-Access                 ANSI SCSI revision: 02
```

The system log provides details about the device, including its size:

```
scsi5 : SCSI emulation for USB Mass Storage devices
  Vendor: LEXAR     Model: JUMPDRIVE SECURE  Rev: 3000
  Type:   Direct-Access                    ANSI SCSI revision: 02
SCSI device sdaz: 506880 512-byte hdwr sectors (260 MB)
```

Additional Resources

Linux Allocated Devices, LANANA:
http://www.lanana.org/docs/device-list/devices.txt

Chapter 7
Using ftServer Fault-Tolerant Utilities and Software

This chapter discusses the following topics:

- "The `ftsmaint` Command"
- "ActiveService Network Support"
- "Kernel Memory Dump File Management"

The ftServer System Software for the Linux Operating System (ftSSS) provides a special command interface, `ftsmaint`, for managing the fault-tolerant components of your ftServer system. It also includes a monitoring and diagnostic package, the ASN, that enables your ftServer system to interact with the Stratus ActiveService Network (ASN). When you configure the ASN, the Stratus Customer Assistance Center (CAC) or your authorized Stratus service representative can receive alarm notifications when faults or other significant events occur on your system, and can remotely diagnose problems. The following sections explain how to use the `ftsmaint` command and how to configure the ASN, and also information about managing dump files and about system load.

The `ftsmaint` Command

The `ftsmaint` command provides a control interface for managing your ftServer system's fault-tolerant functions. To see basic command options, type the `ftsmaint` command.

The `ftsmaint` command arguments support both device query and management tasks. Some of the command arguments only apply to certain devices or systems, which you must specify following the command argument.

Most of the `ftsmaint` command task arguments require an enumerated hardware specification argument following them, indicated by *path* in the command descriptions that follow. Any hardware *`path`* value shown in `ftsmaint ls` output can be used as a *`path`* argument.

The task arguments are as follows (See also `ftsmaint`(8)):

- `ftsmaint ls` *`path`*

  This command displays the status of the hardware specified by the enumerated path. Specifying a path displays a detailed status of the hardware at that path. Omitting the *`path`* argument displays a less-detailed table of all fault-tolerant devices on the system. See “Device Path Enumeration” on page 7-5 for more information.

  Output from `ftsmaint ls` *`path`* reflects what the OSM reports about the state of a given component. Because of system latency, this may not reflect the immediate state of the device. However, you cannot, as a result of this discrepancy, issue a command that would take the system offline. (See `ftsmaint bringDown`.)

  To verify the actual state of the device, check the opstate of its LED.

- `ftsmaint lsLong`

  This command displays the status of all fault-tolerant devices on the ftServer system. This command also returns the status of “empty” devices such as unpopulated PCI slots. This command is useful to study the addressable fault-tolerant devices that can be queried or controlled with ftSSS software.

- `ftsmaint lsPeriph`

  This command displays information about peripheral devices, such as CD-ROMs, DVDs, and modems.

- `ftsmaint lsVND`

  This command displays the status of the Ethernet channel-bonding interfaces in the ftServer system.

- `ftsmaint acSwitch` [10 | 11]

  If you do not provide an enumerated hardware specification, this command toggles the active compatibility of the I/O elements between I/O element 10 and I/O element 11. If you do specify an I/O element, it forces the enumerated enclosure to active status.

- `ftsmaint bringDown` *`path`*

  This command removes from service the CPU element, I/O element, or CPU-I/O enclosure slot specified by *`path`*. No other devices are supported. When you bring down a device, the effect on the system is the same as physically removing CPU-0, I/O-10.

  NOTE

  The `ftsmaint bringDown` command will not permit you to bring down a simplex device, because this would disable the system.

- `ftsmaint bringUp` *`path`*

  This command brings into service the CPU element, I/O element, or CPU-I/O enclosure slot specified by *`path`*. No other devices are supported.

- `ftsmaint burnProm fw_file` *`path`*

  This command updates the firmware contained in the file *`fw_path`* into the EPROM devices on the ftServer device specified by *`path`*. This command can only be used to update BMC and BIOS firmware.

- `ftsmaint clearMtbf` *`path`*

  This command clears the MTBF value of the CPU-I/O enclosure, CPU-I/O enclosure, or CPU-I/O enclosure slot specified by *`path`*.

- `ftsmaint identify [start|stop]` *`path`*

  This command starts or stops the LEDs on the device specified by *`path`*. The device can be a CPU board, an I/O board, or an I/O slot.

- `ftsmaint dump` *`path`*

  This command generates a dump of the BMC or CPU element specified by *`path`*. BMC dumps are stored in `/var/crash/date/bmcx.dmp`, where *`x`* is 10 or 11. CPU dumps are stored in `/var/crash/`*`YYYY-MM-DD-hh-mm`*`/vmcore`. See Table 7-1 for *`path`* values for BMC and CPU.

- `ftsmaint powerOn modem`

  This command supplies electrical power to the modem.

- `ftsmaint powerOff modem`

  This command removes electrical power from the modem.

- `ftsmaint reset modem`

  This command restores modem settings to their factory defaults.

- `ftsmaint resetMtbf` *`path`*

  This command resets the MTBF *`value`* of the CPU element, I/O element, or CPU-I/O enclosure slot specified by *`path`*.

  NOTE

  Do not use this feature to retain a faulty or degraded device in service. It may be useful if the MTBF for a device has been degraded by testing or configuration error.

- `ftsmaint runDiag` *`path`*

  This command starts diagnostics on the CPU element or I/O element specified by *`path`*.

- `ftsmaint setPriority` *`level path`*

  This command sets the priority level of the CPU element specified by *`path`* to the value in the *`level`* argument.

- `ftsmaint setMtbfThresh` *`value path`*

  This command sets the MTBF threshold to *value* of the CPU element, I/O element, or CPU-I/O enclosure slot specified by *`path`*.

- `ftsmaint setMtbfType` *`policy path`*

  This command sets the MTBF type to *policy* on the CPU element, I/O element, or CPU-I/O enclosure slot. The *`policy`* argument can take one of the following values:

  - `useThreshold`
  - `neverRestart`
  - `alwaysRestart`

- `ftsmaint setSensorThresh` *`th_name value path`*

  This command sets the threshold specified by *`th_name`* on the sensor device specified by *`path`* to *`value`*. The *`th_name`* argument can take one of the following values:

  - `uf` (upper fatal)
  - `uc` (upper critical)
  - `unc` (upper noncritical)
  - `lf` (lower fatal)
  - `lc` (lower critical)
  - `lnc` (lower noncritical)

  The opstates for the sensors are as follows:

  - FATAL: above `uf` or below `lf`
  - CRITICAL: above `uc` or below `lc`
  - WARNING: above `unc` or below `unc`
  - NORMAL: default

- `ftsmaint -version`

  This command returns the build number of the `ftsmaint` command on your system. This number coincides with the build number of ftSSS installed on the system.

Device Path Enumeration

Some subsystems and components of the ftServer system are addressable by device path IDs. Device path enumerators uniquely identify the various devices in an ftServer system.

ftServer System Device Path Enumeration

Table 7-1 lists the device paths for devices in an ftServer system.

Table 7-1. Device Paths of ftServer Devices *(Page 1 of 3)*

Location	Device	Path
Top CPU element	Top CPU element	0
	DIMMs (addressed by slot)	0/0—0/7
	Processors	0/20, 0/23
	CPU internal temperature sensor	0/20/130, 0/23/130
	CPU 12v sensors	0/20/150, 0/23/150
	Ambient air temperature sensor	0/130
	Fan sensors	0/140, 0/141
	Voltage sensors	0/150—0/152
Bottom CPU element	Bottom CPU element	1
	DIMMs (addressed by slot)	1/0—1/7
	Processors	1/20, 1/23
	CPU internal temperature sensor	1/20/130, 1/23/130
	CPU 12v sensors	1/20/150, 1/23/150
	Ambient air temperature sensor	1/130
	Fan sensors	1/140, 1/141
	Voltage sensors	1/150—1/152

Table 7-1. Device Paths of ftServer Devices *(Page 2 of 3)*

Location	Device	Path
Top I/O element	Top I/O element	10
	Mass storage controller —EIDE controller	10/0 —05:00.0[†]
	SATA controller —SATA controller	10/1 —05:01.0
	USB controllers —USB host controller	10/2 —05:02.0–05:02.2
	VGA controller —Graphics controller	10/3 —05:03.0
	Ethernet controller —Ethernet card	10/5 —04:02.0, 04:02.1
	PCI device —Core logic	10/8 —03:01.0
	PCI slots 9, 10, 11	10/9—10/11
	Storage enclosure in top I/O element	10/40
	Internal disk slot 1 (maps to sda)	10/40/1
	Internal disk slot 2 (sdb)	10/40/2
	Internal disk slot 3 (sdc)	10/40/3
	BMC	10/120
	Fan speed sensor	10/140
	Voltage sensors	10/150—10/162

Table 7-1. Device Paths of ftServer Devices *(Page 3 of 3)*

Location	Device	Path
Bottom I/O element	Bottom I/O enclosure	11
	Mass storage controller —EIDE controller	11/0 —7c:00.0
	SATA controller —SATA controller	11/1 —7c:01.0
	USB controllers —USB host controller	11/2 —7c:02.0–7c:02.2
	VGA controller —Graphics controller	11/3 —7c:03.0
	Ethernet controller —Ethernet card	11/5 —7b:02.0, 7b:02.1
	PCI device —Core logic	11/8 —7a:01.0
	PCI Slots 9, 10, 11	11/9—11/11
	Storage enclosure in top I/O enclosure	11/40
	Internal disk slot 1 (maps to sdd)	11/40/1
	Internal disk slot 2 (sde)	11/40/2
	Internal disk slot 3 (sdf)	11/40/3
	BMC	11/120
	Fan speed sensor	11/140
	Voltage sensors	11/150—11/162
Optional ftScalable™ Storage Array‡	RAID controller tray	70
	Expansion tray	71

† IDs in the format nn:nn.n indicate PCI bus, slot, and function. These numbers may change as a result of certain system events and are provided here as representative sample data only.
‡ See the *ftScalable Storage: Operation and Maintenance Guide* (R600) for more information about ftScalable Storage device IDs and sample `ftsmaint` output.

Figure 7-1 and Figure 7-2 show the locations of the major enumerated devices.

Figure 7-1. ftServer Enclosures: Locations of Major Enumerated Devices (Front View)

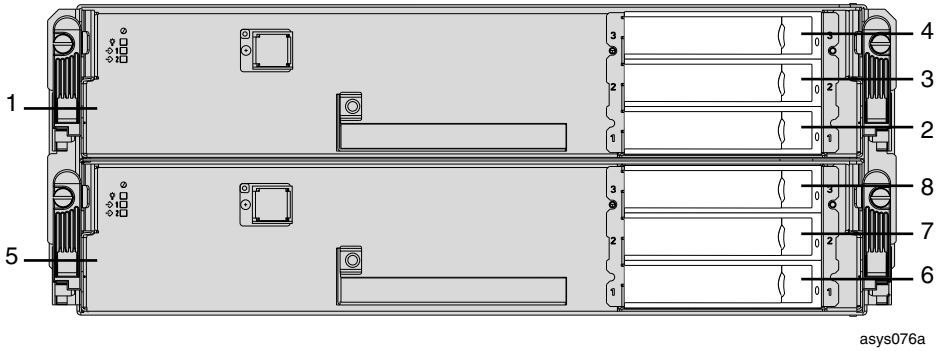

Callout	Device ID	Component
1	0	CPU-0, I/O-10
2	10/40/1	Internal disk drive 1, sda
3	10/40/2	Internal disk drive 2, sdb
4	10/40/3	Internal disk drive 3, sdc
5	1	CPU-1, I/O-11
6	11/40/1	Internal disk drive 1, sdd
7	11/40/2	Internal disk drive 2, sde
8	11/40/3	Internal disk drive 3, sdf

Figure 7-2. ftServer Enclosures: Locations of Major Enumerated Devices (Rear View)

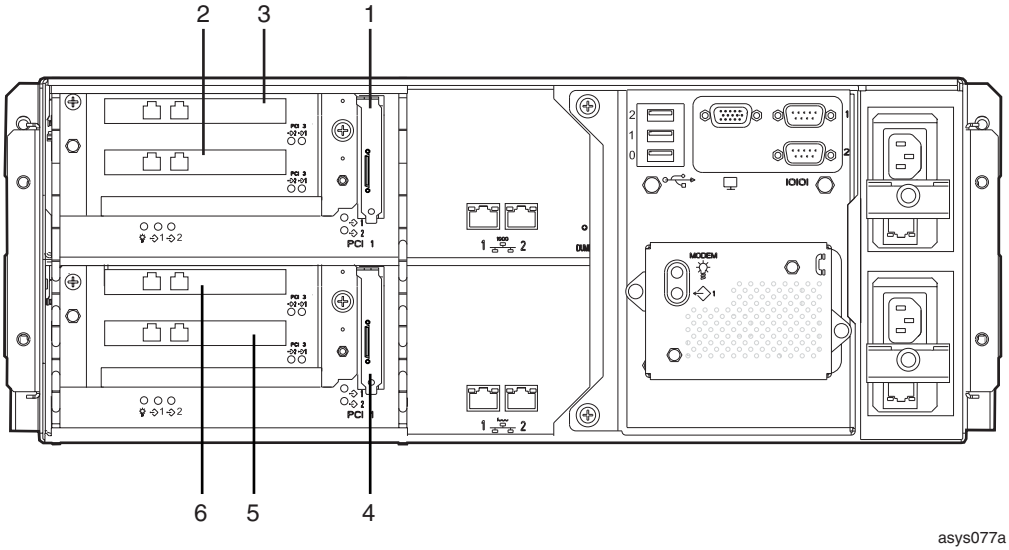

asys077a

Callout	Device ID	Component
1	10/9	Slot 1, low profile (PCI Slot - 9)
2	10/10	Slot 2, full height (PCI Slot - 10)
3	10/11	Slot 3, full height (PCI Slot - 11)
4	11/9	Slot 1, low profile (PCI Slot - 9)
5	11/10	Slot 2, full height (PCI Slot - 10)
6	11/11	Slot 3, full height (PCI Slot - 11)

`ftsmaint` Examples

The following sections provide examples of how to use the `ftsmaint` command:

- "Displaying System Status"
- "Bringing System Components Down and Up"
- "Removing a PCI Adapter From Service and Bringing It Into Service"

Displaying System Status

To display the status of the fault-tolerant devices and subsystems in your ftServer system, issue the following command:

```
# ftsmaint ls
```

Example 7-1. Displaying System Status with the `ftsmaint` Command

```
root@lstlinux14 14:37:27 ~> /opt/ft/bin/ftsmaint ls

Modelx       H/W Path    Description                 State   OPState FRev   Fct
================================================================================
AA-G90730    0           Combined CPU/IO             ONLINE  DUPLEX  *      0
AA-M23100    0/0         1GB DDR-2 DIMM              ONLINE  ONLINE  -      -
AA-M23100    0/1         1GB DDR-2 DIMM              ONLINE  ONLINE  -      -
-            0/2         -                           MISSING EMPTY   -      -
-            0/3         -                           MISSING EMPTY   -      -
-            0/4         -                           MISSING EMPTY   -      -
-            0/5         -                           MISSING EMPTY   -      -
-            0/6         -                           MISSING EMPTY   -      -
-            0/7         -                           MISSING EMPTY   -      -
-            0/20        15 4 1                      ONLINE  ONLINE  -      -
-            0/20/130    Internal Temp               -       NORMAL  -      -
-            0/20/150    +12V                        -       NORMAL  -      -
-            0/23        0 0 0                       ONLINE  ONLINE  -      -
-            0/23/130    Internal Temp               -       NORMAL  -      -
-            0/23/150    +12V                        -       NORMAL  -      -
-            0/130       Ambient Air Temp            -       NORMAL  -      -
-            0/140       Fan1                        -       NORMAL  -      -
-            0/141       Fan2                        -       NORMAL  -      -
-            0/150       +1.2V VTT                   -       NORMAL  -      -
-            0/151       +1.8V VDD                   -       NORMAL  -      -
-            0/152       +12V                        -       NORMAL  -      -
AA-G90730    1           Combined CPU/IO             ONLINE  DUPLEX  *      0
AA-M23100    1/0         1GB DDR-2 DIMM              ONLINE  ONLINE  -      -
AA-M23100    1/1         1GB DDR-2 DIMM              ONLINE  ONLINE  -      -
-            1/2         -                           MISSING EMPTY   -      -
-            1/3         -                           MISSING EMPTY   -      -
-            1/4         -                           MISSING EMPTY   -      -
-            1/5         -                           MISSING EMPTY   -      -
-            1/6         -                           MISSING EMPTY   -      -
-            1/7         -                           MISSING EMPTY   -      -
-            1/20        15 4 1                      ONLINE  ONLINE  -      -
-            1/20/130    Internal Temp               -       NORMAL  -      -
-            1/20/150    +12V                        -       NORMAL  -      -
```

```
-           1/23        0 0 0                       ONLINE  ONLINE  -      -
-           1/23/130    Internal Temp               -       NORMAL  -      -
-           1/23/150    +12V                        -       NORMAL  -      -
-           1/130       Ambient Air Temp            -       NORMAL  -      -
-           1/140       Fan1                        -       NORMAL  -      -
-           1/141       Fan2                        -       NORMAL  -      -
-           1/150       +1.2V VTT                   -       NORMAL  -      -
-           1/151       +1.8V VDD                   -       NORMAL  -      -
-           1/152       +12V                        -       NORMAL  -      -
AA-G90730   10          Combined CPU/IO             ONLINE  DUPLEX  -      0
-           10/0        Mass Storage Ctlr           ONLINE  ONLINE  -      0
-           05:00.0     Fast Track TX EIDE Ctlr     ONLINE  ONLINE  -      -
-           10/1        Mass Storage Ctlr           ONLINE  DUPLEX  -      0
-           05:01.0     PCI/PCI-X SATA Ctlr         ONLINE  DUPLEX  -      -
-           10/2        Serial Bus Ctlrs            ONLINE  ONLINE  -      0
-           05:02.0     USB 1.0 Host Ctlr           ONLINE  ONLINE  -      -
-           05:02.1     USB 1.0 Host Ctlr           ONLINE  ONLINE  -      -
-           05:02.2     USB 2.0 Host Ctlr           ONLINE  ONLINE  -      -
-           10/3        Display Ctlr                ONLINE  DUPLEX  -      0
-           05:03.0     ATI Rage Mobility           ONLINE  DUPLEX  -      -
-           10/4        -                           MISSING EMPTY   -      -
-           10/5        Network Ctlr                ONLINE  DUPLEX  -      0
-           04:02.0     2-port 1GB Enet NIC         ONLINE  DUPLEX  -      -
-           eth000010   Network Interface           ONLINE  DUPLEX  -      -
-           04:02.1     2-port 1GB Enet NIC         ONLINE  DUPLEX  -      -
-           eth000011   Network Interface           ONLINE  DUPLEX  -      -
-           10/6        Bridge                      ONLINE  ONLINE  -      0
-           10/7        Bridge                      ONLINE  ONLINE  -      0
-           10/8        Misc                        ONLINE  ONLINE  -      0
-           03:01.0     ftSwitch Core Logic         ONLINE  ONLINE  -      -
-           10/9        Network Ctlr                ONLINE  ONLINE  -      0
AA-U57500   04:01.0     2-port 1GB Copper Enet NIC  ONLINE  ONLINE  -      -
-           eth000008   Network Interface           BROKEN  BROKEN  -      -
AA-U57500   04:01.1     2-port 1GB Copper Enet NIC  ONLINE  ONLINE  -      -
-           eth000009   Network Interface           BROKEN  BROKEN  -      -
-           10/10       -                           MISSING EMPTY   -      -
-           10/11       -                           MISSING EMPTY   -      -
-           10/40       SATA Enclosure              -       -       -      -
AA-D64200   10/40/1     160GB SATA Disk Drive       ONLINE  DUPLEX  3.00   -
AA-D64300   10/40/2     74GB SATA Disk Drive        ONLINE  ONLINE  33.0   -
-           10/120      Baseboard Management Ctlr   ONLINE  DUPLEX  4.0.0  -
-           10/140      Fan                         -       NORMAL  -      -
-           10/150      -12V                        -       NORMAL  -      -
-           10/151      +1.3V                       -       NORMAL  -      -
-           10/152      +1.5V GB                    -       NORMAL  -      -
-           10/153      +2.5V GB                    -       NORMAL  -      -
-           10/154      +2.5V SATA                  -       NORMAL  -      -
-           10/155      +2.5V VGA                   -       NORMAL  -      -
-           10/156      +3V CLK                     -       NORMAL  -      -
-           10/157      +3.3V                       -       NORMAL  -      -
-           10/158      +3.3Vs                      -       NORMAL  -      -
-           10/159      +3.3V GBE                   -       NORMAL  -      -
-           10/160      +5V                         -       NORMAL  -      -
-           10/161      +5Vs                        -       NORMAL  -      -
-           10/162      +12V                        -       NORMAL  -      -
AA-G90730   11          Combined CPU/IO             ONLINE  DUPLEX  -      0
-           11/0        Mass Storage Ctlr           ONLINE  ONLINE  -      0
-           7c:00.0     Fast Track TX EIDE Ctlr     ONLINE  ONLINE  -      -
```

```
-            11/1        Mass Storage Ctlr           ONLINE  DUPLEX  -      0
-            7c:01.0     PCI/PCI-X SATA Ctlr         ONLINE  DUPLEX  -      -
-            11/2        Serial Bus Ctlrs            ONLINE  ONLINE  -      0
-            7c:02.0     USB 1.0 Host Ctlr           ONLINE  ONLINE  -      -
-            7c:02.1     USB 1.0 Host Ctlr           ONLINE  ONLINE  -      -
-            7c:02.2     USB 2.0 Host Ctlr           ONLINE  ONLINE  -      -
-            11/3        Display Ctlr                ONLINE  DUPLEX  -      0
-            7c:03.0     ATI Rage Mobility           ONLINE  DUPLEX  -      -
-            11/4        -                           MISSING EMPTY   -      -
-            11/5        Network Ctlr                ONLINE  DUPLEX  -      0
-            7b:02.0     2-port 1GB Enet NIC         ONLINE  DUPLEX  -      -
-            eth080010   Network Interface           ONLINE  DUPLEX  -      -
-            7b:02.1     2-port 1GB Enet NIC         ONLINE  DUPLEX  -      -
-            eth080011   Network Interface           ONLINE  DUPLEX  -      -
-            11/6        Bridge                      ONLINE  ONLINE  -      0
-            11/7        Bridge                      ONLINE  ONLINE  -      0
-            11/8        Misc                        ONLINE  ONLINE  -      0
-            7a:01.0     ftSwitch Core Logic         ONLINE  ONLINE  -      -
-            11/9        Network Ctlr                ONLINE  ONLINE  -      0
AA-U57500   7b:01.0     2-port 1GB Copper Enet NIC  ONLINE  ONLINE  -      -
-           eth080008   Network Interface           BROKEN  BROKEN  -      -
AA-U57500   7b:01.1     2-port 1GB Copper Enet NIC  ONLINE  ONLINE  -      -
-           eth080009   Network Interface           BROKEN  BROKEN  -      -
-           11/10       -                           MISSING EMPTY   -      -
-           11/11       -                           MISSING EMPTY   -      -
-           11/40       SATA Enclosure              -       -       -      -
AA-D64200   11/40/1     160GB SATA Disk Drive       ONLINE  DUPLEX  3.00   -
AA-D64300   11/40/2     74GB SATA Disk Drive        ONLINE  ONLINE  33.0   -
-           11/120      Baseboard Management Ctlr   ONLINE  DUPLEX  4.0.0  -
-           11/140      Fan                         -       NORMAL  -      -
-           11/150      -12V                        -       NORMAL  -      -
-           11/151      +1.3V                       -       NORMAL  -      -
-           11/152      +1.5V GB                    -       NORMAL  -      -
-           11/153      +2.5V GB                    -       NORMAL  -      -
-           11/154      +2.5V SATA                  -       NORMAL  -      -
-           11/155      +2.5V VGA                   -       NORMAL  -      -
-           11/156      +3V CLK                     -       NORMAL  -      -
-           11/157      +3.3V                       -       NORMAL  -      -
-           11/158      +3.3Vs                      -       NORMAL  -      -
-           11/159      +3.3V GBE                   -       NORMAL  -      -
-           11/160      +5V                         -       NORMAL  -      -
-           11/161      +5Vs                        -       NORMAL  -      -
-           11/162      +12V                        -       NORMAL  -      -

IO Enclosure 11 is the Active Compatibility Node.

This is an ftServer 2400, 1-way DMR, 3.2 GHz system, P-Package P3403R-1D,
...
* Use lsLong to see this value.
```

Bringing System Components Down and Up

You can use the `ftsmaint` command to bring down and restart a fault-tolerant subsystem. After bringing up a system, it attempts to synchronize and duplex the corresponding components automatically.

For example, the first command below brings down the bottom I/O element; the second command brings it back up:

```
# /opt/ft/bin/ftsmaint bringDown 11
# /opt/ft/bin/ftsmaint bringUp 11
```

NOTE

Before removing an essential component, like an I/O element, from service, first verify that its partner is running.

When you issue the `bringUp` command, the system should automatically synchronize, the RAID array drives should update and become mirrored, and the system should resume duplex operation.

Removing a PCI Adapter From Service and Bringing It Into Service

You can also use the `ftsmaint` command to remove a PCI adapter from service. For example, use the following command to remove the PCI adapter in slot 9 of I/O element 10 from service:

```
# /opt/ft/bin/ftsmaint bringDown 10/9
```

You can bring that PCI adapter back into service by typing the following command:

```
# /opt/ft/bin/ftsmaint bringUp 10/9
```

ActiveService Network Support

To activate ASN support, refer to the *Stratus ActiveService Network Configuration Guide* (R072).

You can prepare a user account for CAC support in advance, in case remote access to your system is needed for support or analysis. You need not activate this account until access is needed; you can limit privileges to those necessary for anticipated tasks; you can monitor ASN service as it is performed, and then deactivate the account once the service access is completed. Access should be supported by modem dial-in serial communications using Internet point-to-point protocol (PPP).

The sra_pppd daemon is installed on the system to support ASN communications that have been pre-configured to use different port assignments than default PPP installations with a serial communications setup that will not conflict with standard software package installations. Usual asynchronous communications, TCP/IP communications, or other PPP communications on the Linux system are not affected.

NOTES

1. Do not use the ASN modem for other purposes. The ASN provides continuous system monitoring. It requires dedicated assignment to a system serial port, as well as a dedicated modem and telephone line for switched line communications.
2. A dedicated phone line provides the most reliable service. ASN calls routed through a PBX may be slow due to load on the PBX, or may not complete successfully due to disconnections. If you must use a PBX, do not route the telephone extension through a switchboard; instead, provide a direct-dial analog number.
3. The `sra_alarm` and `sra_ras` services will not run until the ASN is configured.

You can prepare a user ID for ASN service support with sufficient privileges for anticipated tasks, but you should not unnecessarily leave active an account where administrative privileges can be exercised. Some ASN-monitored events may trigger the need for the CAC to log on to your system remotely to perform some system maintenance or analysis tasks. You should generally activate the account for this purpose and deactivate it on completion. If you use techniques like dial-up and call-back authentication, good password security, and secure shell communications, you may prefer to allow access without getting prior notification. However, you should carefully monitor the use of active accounts on a continuing basis to guard system and network security.

Kernel Memory Dump File Management

CAUTION

To ensure that the ftServer dumping mechanism works successfully, do not enable Diskdump or Netdump. Diskdump or Netdump can interfere with the completion of an ftServer dump.

By default, the supported Linux operating system is installed with kernel memory dump enabled. If the system CPU-I/O enclosures are duplexed at the time of a crash, the system creates a memory dump and stores the dump in the memory of one CPU-I/O enclosure. The system restarts with the remaining CPU-I/O enclosure. After the system restarts, the dump is automatically written in a compressed format to a disk file in `var/crash/`*`YYYY-MM-DD-HH:mm/`*`vmcore`, and an alert is logged to the ASN. When the ASN starts, it reads its notification log, and when the system reaches run level 3 or 5, it files an alert over the ASN server to which it is configured to report. If authorized, the CAC can log into the system and obtain the dump file for analysis if required. Or, you can send the dump file to them.

It is important that you monitor and maintain the size of the `/var/crash` directory. Back up old crash dump data before deleting it.

Chapter 8

Simple Network Management Using Net-SNMP and ftlSNMP

This chapter discusses the following topics:

- "Installing and Configuring ftlSNMP"
- "SNMP Foundations and Concepts"
- "Installing Remote Network Management Services"
- "Managing SNMP"
- "SNMP and MIBS"
- "SNMP Network Management Station Considerations"
- "Initial SNMP Testing"
- "Trap Filtering"

If you are reading this chapter for the first time, be sure you first read *Release Notes: Stratus ftServer System Software for the Linux Operating System* (R005L) as well.

Because Simple Network Management Protocol (SNMP) is new to some administrators, this chapter provides a conceptual introduction following the immediate discussions of installing and configuring ftlSNMP.

Installing and Configuring ftlSNMP

When installing the ftSSS, you have the option of installing ftlSNMP. If you chose not to install ftlSNMP at that time, you can install it later by rerunning the ftServer System Software for the Linux Operating System (ftSSS) installation program and selecting to install ftlSNMP. Installing ftlSNMP this way will not reinstall packages already installed on the system, and therefore not affect unrelated configurations.

On ftServer systems running a supported Linux distribution together with, SNMP is composed of two separate packages:

- Net-SNMP (net-snmp), which is installed during the Linux operating system installation. SNMP is a widely used protocol for monitoring network equipment (for example, routers), computer equipment, and devices such as uninterruptible

power supplies (UPS). *Net-SNMP* is a suite of applications used to implement SNMP v1, SNMP v2c, and SNMP v3 using both IPv4 and IPv6. This suite includes:

  - Various command-line applications for retrieving, manipulating, converting, and displaying information
  - A daemon application for receiving SNMP notifications
  - An extensible agent for responding to SNMP queries for management information
  - A library for developing new SNMP applications

  See www.net-snmp.org for more information about Net-SNMP.

- ftlSNMP (lsb-ft-snmp), which is preinstalled with ftSSS. The ftlSNMP package consists of the following components:

  - SRA-ftLinux-MIB—This MIB supports Stratus's fault-tolerant hardware.
  - ftlsubagent—This subagent supports SRA-ftLinux-MIB SNMP GET and SET operations.
  - ftltrapsubagent—This subagent supports SRA-ftLinux-MIB traps.
  - Various startup, restart, and shutdown scripts
  - Man pages

The ftlSNMP package is currently built against a particular Net-SNMP version. Review the *Release Notes: Stratus ftServer System Software for the Linux Operating System* (R005L) for compatibility information.

During most ftSSS upgrades, ftlSNMP is also upgraded. Install ftlSNMP upgrades with the provided installer or upgrade script. If a script is not provided, use `rpm` to upgrade with the Stratus binary ftlSNMP distribution package. This keeps the RPM database correct and provides a simple package management approach to maintaining ftlSNMP.

ftlSNMP Inventory

This ftlSNMP distribution is provided as a single binary RPM package on the ftSSS distribution called `lsb-ft-snmp-4.0-`*`nnnn`*`.x86_64.rpm`.

The number *nnnn* is the Stratus build number.

When binary RPMs are installed, the `rpm` database will track the revision level of your installed distribution. ftlSNMP binaries are built with performance optimizations for use with the Linux operating system as installed on ftServer systems. See also the *Release Notes: Stratus ftServer System Software for the Linux Operating System* (R005L).

The ftlSNMP package is preinstalled with the default ftSSS installation, and only requires configuration and deployment. Files in the ftlSNMP package are located in the following directories:

- `/etc/opt/ft/snmp`—Contains the fault-tolerant subagent configuration templates and the Net-SNMP master agent configuration template.
- `/etc/opt/ft/snmp/scripts`—Contains the start, stop, and restart scripts.
- `/opt/ft/doc/lsb-ft-snmp-4.0`—Contains the Stratus README file.
- `/opt/ft/mibs`—Contains the Stratus SRA-ftLinux-MIB file.
- `/opt/ft/sbin`—Contains the Stratus fault-tolerant subagents (ftlsubagent and ftltrapsubagent).
- `/opt/ft/share/man/en/man8`—Contains the fault-tolerant subagent man pages.

Manually Installing and Upgrading the ftlSNMP RPM

Whenever possible, avoid manual installation of ftlSNMP. Use the standard ftSSS installation and upgrade procedures to manage ftlSNMP installations and upgrades. Remove earlier versions of ftlSNMP before you install a later ftlSNMP release. Whenever the ftSSS is upgraded, the ftlSNMP distribution may require updating as well. Among other reasons, changes in the `/proc` virtual file system, in supported hardware configurations, and in device support may need to be reflected in the SRA-ftLinux-MIB file and possibly in supporting scripts. You can expect ftlSNMP to be upgraded with most significant ftSSS upgrades.

Search for any prior ftlSNMP with the following command:

```
# rpm -qa | grep lsb-ft-snmp
```

Run the following command to uninstall any conflicting SNMP package installation that is found:

```
# rpm -e lsb-ft-snmp
```

Before you upgrade, make sure the prerequisite Net-SNMP packages are installed. To install or upgrade ftlSNMP from the distribution RPM, enter the following command:

```
# rpm -Uvh path/lsb-ft-snmp-4.0-nnnn.x86_64.rpm
```

The string *path* is the directory containing the RPM, and *nnnn* is the Stratus build number. The executables will be placed in `/opt/ft/sbin/ftltrapsubagent` and `/opt/ft/sbin/ftlsubagent`. A plain text file representing SRA-ftLinux-MIB will be placed in the `/opt/ft/mibs` directory.

Enter the following commands, and optionally, add them to the login user's profile (for example `/etc/.bash_profile`).

```
# export MIBDIRS=/usr/share/snmp/mibs:/opt/ft/mibs

# export MIBS=ALL
```

This installs or upgrades the MIBs and subagents.

ftlSNMP Prerequisites

ftlSNMP requires ftSSS and the following Net-SNMP packages to be installed. Note that *n.n.n.n* represents the current supported Net-SNMP release number.

- `net-snmp-libs-`*`n.n.n.n`*
- `net-snmp-`*`n.n.n.n`*
- `net-snmp-utils-`*`n.n.n.n`*
- `net-snmp-perl-`*`n.n.n.n`*

SNMP Configuration File Updates

To configure ftlSNMP, edit the following files:

- `/etc/opt/ft/snmp/snmpd.conf` (see "The `snmpd.conf` File" on page 8-5 and *snmpd.conf*(5))

  NOTE

  The ftSSS installation automatically creates `snmpd.conf` in the `/etc/opt/ft/snmp` directory, while Net-SNMP creates it in the `/etc/snmp` directory. You should use the `/etc/opt/ft/snmp` version because it is the directory that the ftlSNMP scripts reference.

- `/etc/opt/ft/snmp/ftlsubagent.conf` (see "The `ftlsubagent.conf` and `ftltrapsubagent.conf` Files" on page 8-5)
- `/etc/opt/ft/snmp/ftltrapsubagent.conf` (see "The `ftlsubagent.conf` and `ftltrapsubagent.conf` Files" on page 8-5)

These files are created when you install or upgrade ftlSNMP, provided that they do not already exist. If they already exist, installing or upgrading ftLSNMP does not overwrite them. Use the information in the template files (for example, `/etc/opt/ft/snmp/snmpd.conf.template`) as a guide when editing the files. However, values of parameters to be set in these files are system-, network-, and SNMP-manager specific.

The `snmpd.conf` File

CAUTION

Use SNMPv3 when the manager and master agent are separated on a public network.

The following is an example only. Failure to use SNMPv3 when communicating over a public network is a server and network security risk.

SNMP V3 includes true authentication and encryption. The three authentication models are NoAuthnoPriv, authNoPriv, and authPriv. Note that you must have auth status for encryption.

An SNMP engine identifier takes the first IP address as the default that identifies the agent in the device. Each device must have a user login account for the device.

SNMPv3 also has concepts of groups, views, and privileges for access control. These are referred to as the view-based access control model (VACM) and user-based security model (USM).

You **must** keep the following two lines in the `snmpd.conf` file for the master agent to function properly:

```
master agentx
agentxTimeout 60
```

ftlSNMP requires agentX services.

To avoid timeouts when the subagents are running under abnormal system stress (for example, 80% CPU usage and disks heavily stressed), raise the value of `agentxTimeout`. If a timeout occurs, there will be a short delay while the subagents reinitialize their communications.

The `ftlsubagent.conf` and `ftltrapsubagent.conf` Files

These files require no editing for default operation, but you may want to adjust logging. You can change the trace level from `off` to `brief` or `verbose` as desired, or as suggested by Stratus to aid in diagnosing any problems. Debugging information will be logged.

These files contain the following configuration lines:

```
sraTraceLevel off
sraTraceLog /var/opt/ft/log/ftlsubagents.log
```

With `sraTraceLevel` set to `brief`, data flows to and from ftlSNMP external items are traced. With `sraTraceLevel` set to `verbose`, internal items are also traced. You can also change the location of the log file.

Agent and subagent startup and shutdown events are separately logged in `syslog`.

With trace levels other than `off`, logs may grow rapidly (depending on the number of managed objects and their activity). In order to limit the size of the logs, you can use `logrotate` to manage log size and archiving. See *logrotate*(8).

Configuring SNMP to Start at System Initialization

After installation of the lsb-ft-snmp package, the `ft-snmp` initialization script is installed in the `/etc/init.d` directory, and the `ft-snmp` service is added. During system initialization, the `ft-snmp` service is automatically started at run level 3, 4, or 5. This initialization script provides start, stop, restart, and status functionality.

Configuring SNMP for Service Management

There are several ways of adding SNMPv3 user accounts, which allow local and remote access to SNMP services.

SNMP provides an `/etc/opt/ft/snmp/snmpd.conf` file that contains instructions and sample entries. Uncomment the right lines, then run the commands.

Uncomment these lines in `/etc/opt/ft/snmp/snmpd.conf`:

```
createUser admin MD5 your_passwd DES
rwuser v3user
group  v3usergroup     usm     admin
group  v3usergroup     usm     v3user
view   v3view       included   .1.3.6.1
access v3usergroup ""  usm     authNoPriv    exact  v3view
   v3view   v3view
```

Run the following commands each time SNMP is restarted (or write a script to manage this task):

```
# snmpusm -v3 -u admin -n "" -l authNoPriv -a MD5 -A
your_passwd localhost    create v3user admin

# snmpusm -v3 -u v3user -n "" -l authNoPriv -a MD5 -A
your_passwd localhost    passwd old_passwd new_passwd
```

These commands clone an initial (template) SNMPv3 user, `admin`, as v3user, and then change the password of `v3user`. The string *`old_passwd`* is the password previously assigned to the user admin.

Another way to add SNMPv3 users is to run the following three commands:

```
# stop_snmp
# net-snmp-config --create-snmpv3-user
# start_snmp
```

This series of commands adds lines in the proper configuration files to add a user to the system. Be sure to run `stop_snmp` before running the `net-snmp-config` command, which prompts you for a user name and password. The command adds a `create user snmpv3-user MD5 password DES` line entry into `/etc/opt/ft/snmp/snmpd.conf` file. This line is automatically replaced with a key on SNMP restart because it contains the correct password.

The `net-snmp-config` command also puts an `rwuser` entry into the `/etc/opt/ft/snmp/snmpd.conf` file. You will be prompted for a user name and password when you enter the `net-snmp-config` command.

You will need to configure `snmpd.conf` with VACM entries for a new user.

1. To create a new user configuration, enter the following command:

   # `net-snmp-config --create-snmpv3-user`

2. When prompted, enter the user name and password you want to assign.
3. Edit `/etc/opt/ft/snmp/snmpd.conf` and add the user to the VACM using a current group and view or creating new ones. The following example lines add a new user `paul` to the current view and group in `snmd.conf` by inserting the (highlighted) line for `paul`:

   ```
   group   v3group     usm      admin
   group   v3group     usm      v3user
   group   v3group     usm      paul
   view    v3view      included .1.3.6.1
   access  v3group     ""  usm  authNoPriv exact v3view
      v3view v3view
   ```

   Note that this example shows that you previously ran `snmpusm` commands to create the admin user.

4. Start SNMP and run the following two commands to clone a new user and change the password:

   # snmpusm -v3 -u admin -n "" -l authNoPriv -a MD5 -A *your_passwd* localhost create paul admin

   # snmpusm -v3 -u paul -n "" -l authNoPriv -a MD5 -A *your_passwd* localhost passwd *old_passwd new_passwd*

When you run `start_snmp`, you will be able to use this user ID and password in SNMPv3 `snmpwalk` and `snmpget` commands, for example:

```
# snmpwalk -v 3 -t 40 -l authNoPriv -u paul -A new_passwd
localhost    1.3.6.1.4.1.458
```

Unless you plan to use SNMP to monitor the local server, you should configure SNMP to shut down when entering runlevel 1. At any rate, SNMP daemons need to exit before shutting the system down. It is particularly important that you shut down SNMP before performing an ftSSS upgrade.

SNMP Foundations and Concepts

The Net-SNMP and ftlSNMP packages support the SNMP protocol and many of the capabilities of SNMP for managing network objects using protocols and interface features described in numerous Internet Engineering Task Force (IETF) documents. Net-SNMP and ftlSNMP are packages for network administration that compatibly support Stratus ftServer fault-tolerant operations using standard network communications. There are few network administration tools available that readily support fault-tolerant capabilities of networked systems and devices. ftlSNMP allows Stratus ftServer systems to be monitored and managed by any remote-networked system running SNMP-based network management software.

Net-SNMP provides a functional network administration package for use on ftServer systems to meet identified customer needs. ftlSNMP is a unique extension of Net-SNMP that provides the SRA-ftLinux-MIB to define manageable systems and components of ftServer Linux-based systems. ftServer subagents and MIB provide SNMP support and services for fault-tolerant operations.

ftlSNMP Management Commands

The following commands are provided to start and stop the SNMP master agent daemon and subagents on the ftServer host system:

- `start_snmp`, `stop_snmp`, `restart_snmp`—These commands, start, stop, and restart (stop, then restart) the two Stratus subagents, `ftlsubagent` and `ftltrapsubagent`, and the master agent.

Use the preceding commands for most situations.

CAUTION

The following commands are also available to deal with special situations. However, you should **not** normally use them except under the guidance of the CAC or your authorized Stratus service representative, as starting and stopping agents in untested order can have unforeseen consequences.

- `start_all_subagents`, `stop_all_subagents`, `restart_all_subagents`—These commands start, stop, and restart only the subagents.
- `start_ftlsubagent`, `stop_ftlsubagent`, `restart_ftlsubagent`—These commands start, stop, and restart only `ftlsubagent`.
- `start_ftltrapsubagent`, `stop_ftltrapsubagent`, `restart_ftltrapsubagent`—These commands start, stop, and restart only `ftltrapsubagent`.
- `start_snmp_daemon`, `stop_snmp_daemon`, `restart_snmp_daemon`—These commands start, stop, and restart only `snmp_daemon` (the master agent).

The Basic Net-SNMP Commands

These tools provide a basic set of features for exercising and managing objects using a standard command syntax and core functionality:

NOTE

Although these commands are documented as user commands (*man* (1)), you should treat SNMP utilities as the administrative tools they are, and closely limit privileges to execute these commands.

- `snmpwalk`—This command uses SNMP `GETNEXT` requests to query a network entity for a tree of information that maps the managed objects by object ID hierarchically. See *snmpwalk*(1). While this can return much information, take care not to use this command on a heavily loaded net, since it can add significantly to traffic.
- `snmpget`— This command queries a single SNMP object using an SNMP `GET` request. See *snmpget*(1).
- `snmpgetnext`— This command uses `GETNEXT` requests to query network entities for information.
- `snmpgetbulk`—This command uses the SNMP `GETBULK` request to query a network entity for quantities of information. See *snmpgetbulk*(1).

- `snmpdf`—This command replicates `df` command functionality on a network-accessible drive. The `snmpdf` command checks disk space on the remote machine by examining the system's HOST-RESOURCES-MIB `hrStorageTable`, or a UCD-SNMP-MIB's `dskTable` value. See *snmpdf*(1).
- `snmpstatus`—This command queries a network entity to retrieve significant information about a communicating object. See *snmpstatus*(1). The following information is retrieved:
  - The IP address of the entity.
  - A textual description of the entity (`sysDescr.0`)
  - The uptime of the entity's SNMP agent (`sysUpTime.0`)
  - The sum of received packets on interfaces (`ifInUCastPkts.* + ifInNUCastPkts.*`)
  - The sum of transmitted interface packets (`ifOutUCastPkts.* + ifOutNUCastPkts.*`)
  - The number of IP input packets (`ipInReceives.0`)
  - The number of IP output packets (`ipOutRequests.0`)
- `snmptranslate`—This command converts object ID values into more easily understood forms. See *snmptranslate*(1).
- `snmptable`—This command repeatedly uses SNMP `GETNEXT` or `GETBULK` requests to get information on a network entity, which is specified as, and must be mapped by, a table. See *snmptable*(1).
- `snmpset`—This command uses the SNMP `SET` request to control, or set information on, a network entity. See *snmpset*(1).
- `snmptrap`—This command uses the SNMP `TRAP` operation to send information to a network manager when a trigger condition is met. See *snmptrap*(8).
- `snmpinform`—This command essentially works like `snmptrap`, but uses a different form of signal, and can require a response in order to suppress resending. See *snmptrap*(1).
- `snmptest`—This command is a flexible test utility that can send a variety of signals and retrieve a variety of information. It is best used within shell scripts that can hide its complexity and focus on particular test queries. See *snmptest*(1).
- `snmpnetstat`—This command is a powerful data retrieval tool to query a remote system and retrieve a variety of information about communications objects. See *snmpnetstat*(1).
- `snmpdelta`—This command is a tool used to monitor values of a network object over time, and respond if the values deviate from established parameters. See *snmpdelta*(1).

MIBs

A *management information base* (MIB) uses ISO Abstract Syntax Notation 1 to assign a unique object identifier to any object to be managed by SNMP. This syntax is a hierarchical model that is intended to provide unique object identification. Under this notation, there is conceptually only one true MIB; everything fits within it. To the extent that developers observe syntactic standards, various MIB definitions will not conflict, because any MIB used should use only unique identifiers. Thus Net-SNMP and other SNMP implementations allow a large number of MIBs to be loaded from various paths at initialization, under the presumption that all identifications are unique. In current standard SNMP implementations, at least the IETF MIB-II definitions supporting RFC1213 must be used. RFC1213, *Management Information Base for Network Management of TCP/IP-based internets: MIB-II* can be downloaded from the IETF Web site.

SNMPv1 MIBs supported only strings of data, but later MIBs typically use a columnar layout of information that can be easily manipulated by a scripting language that handles textual data, such as Perl. The supplied MIBs are stored in the `/opt/ft/mibs` and `/usr/share/snmp/mibs` directories in the default search path.

NOTE

The ftSSS installation automatically creates the SRA-ftLinux-MIB file in the `/opt/ft/mibs` directory, while Net-SNMP creates its MIBs in the `/usr/share/snmp/mibs` directory.

You can store MIBs in a variety of locations, but, if the MIB is added after an SNMP agent is running, you must direct SNMP agents to the location of a MIB the first time the agent is used. If the MIB is not in the path when the SNMP services are started (and SRA-ftLinux-MIB exported), use a command similar to the following example to identify the SRA-ftLinux-MIB file to the SNMP tools:

```
# snmpget -m SRA-ftLinux-MIB -v 2c -c public myhost.com
   ftcBdState.1

SRA-ftLinux-MIB::ftcBdState.1 = INTEGER: duplex(21)
```

As you begin to develop your own MIB (or MIBs) for your management requirements, you can save a lot of work by adopting defined variables from other MIBs. A large number of MIBs are defined by the IETF and are available as plain text files. You should use standardized MIBs where they define objects to avoid non-standard implementation of networked objects.

Do not alter standard MIBs. If you need additional object definitions, you can add another MIB or create your own. The MIB-defined objects can be queried and data

recovered that provides a basis for SNMP agent operations. You can create scripts that the SNMP agent or a subagent executes according to MIB definitions.

Some Objects Defined by Standard MIBs

For a practical implementation of SNMP, a number of objects simply must be defined. Some of these are introduced here.

For UDP or TCP/IP communications and collection of statistical data about communications and communications channels, MIB-II defines some necessary objects. MIB-II defines these objects for querying:

```
system
interfaces
at
ip
icmp
tcp
udp
snmp
```

The Net-SNMP implementation requires basic support of the Host Resources MIB. The objects defined in RFC1514 Host Resources MIB include:

```
hrSystem
hrStorage
hrDevice
hrSWRun
hrSWRunPerf
hrSWinstalled
hrSWRunID
```

NOTE

A Host Resources MIB should support these objects, which are defined in RFC1514. A Newer Host Resources MIB may comply with RFC2790, which extends and replaces RFC1514. The Net-SNMP Host Resources MIB implementation has been tested for the RFC1514 features.

The Net-SNMP package also implements the Net-SNMP version of the UCD Extensions MIB, which defines these objects:

- `prTable`
- `memory`
- `laTable`
- `systemStats`
- `fileTable`

SNMPv3 Support

SNMPv3 support includes implementation of IETF RFCs 3410 through 3418. The third version of the Simple Network Management Protocol, presented by the IETF as the Internet Standard Management Framework RFC3410, SNMPv3 incorporates elements of SNMPv1 and SNMPv2, and shares the same basic modular architecture. This framework consists of four structures: a data definition language (SMIv1), a management information base (MIB) defining management information, a separately defined communication protocol, and security and administration applications and engines.

Features of SNMP version 1, SNMP version 2, and SNMP version 3 are not mutually exclusive. IETF Best Current Practices 74 (BCP74) describes how to implement these protocols compatibly on networks and on internetworked environments so that objects can be managed using the least sophisticated protocol required. In this way, networked and internetworked objects may be managed using SNMPv1, for example, without becoming obsoleted if the network is commingled into a larger network where objects are managed using SNMPv2 or SNMPv3.

This is necessary because the SNMP schema treats all networks as potentially a single network, providing for addressing every object uniquely with a single MIB. Accordingly, implementing conformant extensions to SNMP should not cause interoperability conflicts with existing standards-conforming SNMP implementations. In the SNMP network universe, any number of SNMP servers can exist, and they can manage the objects they know about using SNMPv1, SNMPv2, SNMPv3, and with confidence that any SNMPv4 or subsequent protocol that may be defined will not obsolete existing SNMP servers and their managed objects.

SNMP's View of a Network

SNMPv1 defines a simple and robust internet protocol-based communications method for tracking the status of and managing almost any network-interactive item that is sufficiently defined as an object in a MIB.

SNMP normally uses UDP protocol implemented on socket-based IP communications, but may also be implemented using TCP/IP and another IP-based protocol, and also on non-IP protocols such as RS-232 serial communications by spoofing an IP-based communication or by piggybacking it on another transmission or transfer protocol. SNMP can also take advantage of common security enhancements implemented over IP, such as the Secure Socket Layer and other encryption, authentication, and remote access technologies provided by, for example, `ssh`, the Openssh package.

SNMPv2 expands management capabilities of SNMPv1 by providing a mechanism for more easily defining the managed objects that SNMP communicates with. SNMP refines SNMPv2 definitions and adds important security features. Net-SNMP supports SNMPv1, SNMPv2, and SNMPv3 protocols. Because of the simple basic structure of SNMP, applications developed for any SNMP implementation tend to be easily adaptable and useful with other SNMP implementations.

Conceptually, every managed object on a network is uniquely identifiable. SNMP uses ISO Abstract Syntax Notation Standard 1 (ASN.1) to place every SNMP object within the internet hierarchy of managed objects. All these unique managed objects can be managed by their defined characteristics in the MIB. While in the theoretical schema there is only one MIB, it is usual to refer to any file that provides SNMP MIB definitions as a MIB. MIBs can be formally registered and entered into defined namespace or used locally as experimental MIBs.

An SNMP server only knows of objects for which it has definitions. This allows distributed SNMP services to co-exist on networks without interfering with each other. SNMP agents can, however, interact. SNMP agents can act as subagents of a master agent. A managed object can be a host computer or subsystem, an arbitrary interactive device, or a software application (including an operating system), basically anything whose interactivity over the network can be defined in a MIB so that it can be interfaced via SNMP.

Extensions and Fault-Tolerant SNMP Operation

While Net-SNMP supports the security features of SNMPv3, it can also interact compatibly with distributed SNMP services that use SNMPv1 and SNMPv2. Net-SNMP is the most widely-adopted open source SNMP utility package. This facilitates interfacing the Net-SNMP and ftlSNMP implementations with other servers deploying Net-SNMP-based distributed SNMP services and service management utilities in a heterogeneous network environment. Net-SNMP has been ported to Linux, UNIX, and other operating systems, such as Windows NT and Stratus VOS. Note, however, that the Net-SNMP and ftlSNMP combined packages provide support only for ftServer systems running a supported Linux distribution together with ftSSS.

The Net-SNMP and ftlSNMP packages encourage deployment of distributed SNMP services on heterogeneous networks featuring both ftServer systems running a supported Linux distribution together with ftSSS, and Stratus ftServer hosts running other supported operating systems. The Net-SNMP and ftlSNMP packages also can be used to complement Stratus ActiveService Network alarm and notification functions to strengthen server management support for Linux administrators. See ‘‘ActiveService Network Support’’ on page 7-14.

The Net-SNMP and ftlSNMP packages interact with and manage networked objects defined in MIB files. The ftlSNMP package includes the SRA-ftLinux-MIB file (SRA-ftLinux-MIB.txt) to support fault-tolerant Stratus ftServer systems. Net-SNMP also supports MIB-II and Host Resource MIB features. The ftlSNMP package follows the SNMP master/agent daemon management model, and extends the basic model using AgentX subagents. This allows the subagents associated with different MIBs to be kept separate so that failure of one does not bring down the others. Also, the `ftltrapsubagent` was kept separate from the `ftlsubagent` to avoid blocking on serious traps. AgentX extensions are defined in RFC2741, *Agent Extensibility (AgentX) Protocol version 1*. RFC2741 defines a standardized framework for extensible SNMP agents, and then defines master agents and subagents as processing daemons. An AgentX protocol is defined for communication between an AgentX-capable master agent and subagents. RFC2741 also defines elements of procedure for an AgentX daemon to process SNMP protocol messages.

Traditional CMU SNMP management utilities are modestly refined and enhanced in Net-SNMP. Most of the ftlSNMP extensions to Net-SNMP come through added MIB definitions and correspondingly augmented configuration files for managing SNMP agent and subagent daemons.

Installing Remote Network Management Services

The remote host system must have an SNMP system installed that supports SNMPv3 (an example is the Red Hat Net-SNMP distribution, which the following sample procedure assumes).

Transfer the SRA-ftLinux-MIB.txt file to the host system and install it in `/opt/ft/mibs/SRA-ftLinux-MIB.txt` (or wherever the host system stores its MIBs). You can use `ftp` to transfer the MIB file from the ftServer system to the host system as long as the two systems can communicate with each other over a network.

Run these commands to set up MIB path environment variables and reinitialize Net-SNMP:

```
# export MIBDIRS=/usr/share/snmp/mibs:/opt/ft/mibs
# export MIBS=ALL
# service snmpd stop
# service snmptrapd stop
# service snmpd start
# service snmptrapd start
```

The host SNMP manager now can execute commands managing a network-accessible Stratus ftServer system running the Linux operating system and the Net-SNMP and ftlSNMP packages.

To receive traps, configure the `/etc/opt/ft/snmp/snmpd.conf` file on the ftServer system, adding trapsink entries pointing to the host server running `snmptrapd`. The configuration lines should look something like this:

```
trapcommunity   public
trapsink        192.168.33.75   public
trap2sink       192.168.33.75   public
```

In this sample, the community is `public`. This would not be the usual case on an internet-accessible system.

Configuring SNMP for Remote Service Management

The procedure for configuring Net-SNMP is very similar to “Configuring SNMP for Service Management” on page 8-6, which describes enabling remote services by adding SNMP users and groups. If you are using a network management station, you may have some other procedure provided with your software.

Deploying SNMP Agents and Subagents

The basic SNMP model has a master agent on the SNMP server system, with behavior configured by MIBs. The master agent manages one or more subagents. Typically, a single subagent is used per system (including the SNMP server system). However, with AgentX extensions, multiple agents can be deployed on a system, performing different tasks under control of the master agent. Even a MIB-II subagent can be extended to provide new functionality using AgentX. Although logging is flexible, on some SNMP systems logging is simply merged with syslog output. In the default ftlSNMP configuration, logging is configured as shown in Figure 8-1.

Figure 8-1. AgentX-Enabled Extensions and Subagents

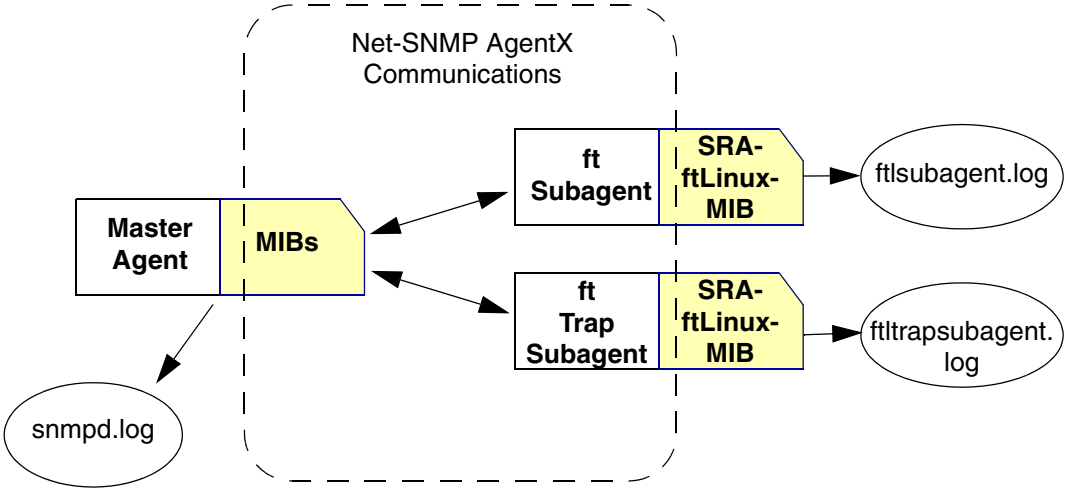

By default, `snmpd.log`, `ftlsubagent.log`, and `fttrapsubagent.log` are located in `/var/opt/ft/log`. You can relocate the subagent logs by modifying `/etc/opt/ft/snmp/ftlsubagent.conf` and `/etc/opt/ft/snmp/ftltrapsubagent.conf` (see "The `ftlsubagent.conf` and `ftltrapsubagent.conf` Files" on page 8-5). If you have `sraTraceLevel` set to `brief` or `verbose` in these files, you may want to relocate the logs to a file system with ample space.

Verifying Traps

You can easily verify traps using `snmptrapd` on a remote Linux system with Net-SNMP installed.

1. On the remote Linux system, set up Net-SNMP to autostart, and verify it using the `chkconfig` command, or manually start Net-SNMP.
2. On the ftServer system with the Linux operating system, ftSSS, Net-SNMP, and ftlSNMP installed, configure `/etc/opt/ft/snmp/snmpd.conf` with trapsink entries that point to the IP address of a remote management system running `snmptrapd`, by adding lines as follows:

   ```
   trapcommunity   public
   trapsink        ip_address  public
   trap2sink       ip_address  public
   ```

   In this example, *`ip_address`* is the IP address of the remote management system.

3. Start (or restart) SNMP using the `/etc/opt/ft/snmp/scripts/start_snmp` (or `/etc/opt/ft/snmp/scripts/restart_snmp`) command for changes to take effect.
4. From the remote management system, you can view the traps as they are generated, by tailing /`var/log/messages` (or wherever the remote management system is configured to log `snmptrapd` messages):

   ```
   tail -f /var/log/messages
   ```

Pulling CPU-I/O enclosures and/or pulling Ethernet cables on the ftServer system will generate traps, as will exercising the system using the `ftsmaint` command.

Managing SNMP

The following sections discuss how to manage SNMP on your system:

- “Testing Your SNMP Configuration”
- “Managing ftServer Hardware Components”
- “Testing Ethernet Ports”

NOTE

The sample command lines in the following sections are for general guidance only. Some of the command-line details and command output shown—for instance, PCI adapter device names—may differ from what is applicable to your system.

Testing Your SNMP Configuration

The following are some Net-SNMP commands that you can use to test or exercise MIBs. If you run these remotely, the target name and IP address will differ.

To walk the SRA-ftLinux-MIB file:

```
# snmpwalk -v 1 -c public -t 120 localhost 1.3.6.1.4.1.458
# snmpwalk -v 2c -c public localhost 1.3.6.1.4.1.458
```

To walk the Stratus `ftcPcidevcnf` table:

```
# snmpwalk -v 2c -c public localhost
1.3.6.1.4.1.458.107.1.2.5.2.1
```

To walk the Stratus `EtherState` by symbolic *object identifier* (OID) name:

```
# snmpwalk -v 2c -c public localhost ftcEtherState
```

To walk the UCDAVIS MIB:

```
# snmpwalk -v 2c -c public localhost .1.3.6.1.4.1.2021.2
# snmpwalk –v 2c –c public localhost
 .iso.org.dod.internet.private.enterprises.ucdavis.memory
# snmpwalk -v 2c -c public localhost ucdavis
```

To use SNMPv3 with `snmpwalk`:

```
# snmpwalk -v 3 -l authNoPriv -u v3user -A new_passwd localhost
ucdavis

# snmpwalk -v 3 -l authNoPriv -u v3user -A new_passwd localhost
system

# snmpwalk -v 3 -l authNoPriv -u v3user -A new_passwd localhost
1.3.6.1.4.1.458

# snmpwalk -v 3 -t 40 -l authNoPriv -u v3user -A new_passwd
localhost    1.3.6.1.4.1.458
```

In these command examples, `v3user` and *new_passwd* are the user name and password set up in “Configuring SNMP for Service Management” on page 8-6.

Managing ftServer Hardware Components

In the following command examples, `v3user` and *new_passwd* are the user name and password set up in “Configuring SNMP for Service Management” on page 8-6.

To bring down a CPU element

Note that a final octet 1 identifies CPU element 0 and a final octet 2 identifies CPU element 1.

```
# snmpset -v 3 -t 40 -l authNoPriv -u v3user -A new_passwd
localhost    1.3.6.1.4.1.458.107.1.2.1.2.3.1.13.1 s test
```

To bring up a CPU element

```
# snmpset -v 3 -t 40 -l authNoPriv -u v3user -A new_passwd
localhost    1.3.6.1.4.1.458.107.1.2.1.2.3.1.11.1 s test
```

To bring down an I/O element

Note that a final octet 1 identifies I/O element 10 and a final octet 2 identifies I/O element 11.

```
# snmpset -v 3 -t 40 -l authNoPriv -u v3user -A new_passwd
localhost    1.3.6.1.4.1.458.107.1.2.1.3.3.1.14.1 s test
```

To bring up an I/O element

```
# snmpset -v 3 -t 40 -l authNoPriv -u v3user -A new_passwd
localhost    1.3.6.1.4.1.458.107.1.2.1.3.3.1.12.1 s test
```

To control firmware burn (FWBURN)

```
# snmpset -v 3 -t 40 -l authNoPriv -u v3user -A new_passwd
localhost    1.3.6.1.4.1.458.107.1.2.1.2.3.1.15.1 s FWBURN
```

Example: Managing Hardware

In this example, only relevant portions of the `ftsmaint` command output are shown. The following example illustrates bringing a CPU-I/O enclosure down and then back up.

Issue the following command to check the status of CPU element 0, I/O element 10:

```
# /opt/ft/bin/ftsmaint ls 0

H/W Path                : 0
Description             : Combined CPU/IO
State                   : ONLINE
Op State                : DUPLEX
Reason                  : SECONDARY
LED (Green)             : ON
LED (Yellow)            : OFF
LED (White)             : ON
...
```

The command output shows that CPU element 0 is online and duplexed, so it is safe to remove if from service. To bring down CPU element 0 by invoking the `ftcCpubdInitiateBringDown` command, use the whole numeric OID for that command (`1.3.6.1.4.1.458.107.1.2.1.2.3.1.13`) plus the CPU element 0 index (1) as the final octet. Thus, the complete OID is:

```
1.3.6.1.4.1.458.107.1.2.1.2.3.1.13.1
```

The following example assumes that the community string “private” has been defined in `/etc/opt/ft/snmp/snmpd.conf`.

```
# snmpset -v 1 -c private localhost
1.3.6.1.4.1.458.107.1.2.1.2.3.1.13.1 s test

SRA-ftLinux-MIB::ftcCpubdInitiateBringDown.3 = STRING:
"test"
#
```

To check that the CPU-I/O enclosure status has changed:

```
# /opt/ft/bin/ftsmaint ls 0

H/W Path                  : 0
Description               : Combined CPU/IO
State                     : OFFLINE
Op State                  : REMOVED_FROM_SERVICE
Reason                    : OK_FOR_BRINGUP
LED State                 : RED
...
```

To bring CPU element 0 back up by invoking the `ftcCpubdInitiateBringUp` command, use the numeric OID (see ''SRA-ftLinux-MIB OID Values and Properties'' on page 8-33) for that command (1.3.6.1.4.1.458.107.1.2.1.2.3.1.11), again followed by CPU element 0's index:

```
# ./snmpset -v 1 -c private localhost
1.3.6.1.4.1.458.107.1.2.1.2.3.1.11.1 s test

SRA-ftLinux-MIB::ftcCpubdInitiateBringUp.3 = STRING: "test"
#
```

If you check CPU element 0's status immediately, you can see that it has started initializing:

```
# /opt/ft/bin/ftsmaint ls 0

H/W Path                  : 0
Description               : Combined CPU/IO
State                     : INRESET
Op State                  : INITIALIZING
Reason                    : NONE
LED State                 : RED
...
```

After a while, it is fully back up again:

```
# /opt/ft/bin/ftsmaint ls 0

H/W Path                  : 0
Description               : Combined CPU/IO
State                     : ONLINE
Op State                  : DUPLEX
Reason                    : PRIMARY
LED State                 : GREEN
...
```

Testing Ethernet Ports

You can test Ethernet ports for proper traps and changes to OIDs. On an ftServer system running a supported Linux distribution together with ftSSS, Ethernet ports are uniquely identified.

When testing cable pulls or bringdowns, the system should generate traps, and the data that Stratus MIB objects returned should reflect these changes.

One approach is to set up an `snmptrapd` on a Linux system to verify the traps as they are generated (see "Verifying Traps" on page 8-18) and to run `snmpwalk` on the SRA-ftLinux-MIB file before and after a fault insertion to verify object data changes. A `diff` of these two walks will reveal changes that Stratus MIB objects return.

Example: Testing Ethernet Ports

The following example demonstrates how to test Ethernet ports. The example assumes dual-port 10/100/1000-Mbps Ethernet PCI adapters installed in slot 1 of both CPU-I/O enclosures. First, determine the instance name of the Ethernet device and which slot it is in:

1. Run an `snmpwalk` on `ftcEtherInstanceName` OID. This gives you a list of `EtherInstance` names mapped to Ethernet device names:

   # **`snmpwalk -v 1 -c private -t 40 localhost ftcEtherInstanceName`**

   ```
   SRA-ftLinux-MIB::ftcEtherInstanceName.1 = STRING: "lo"
   SRA-ftLinux-MIB::ftcEtherInstanceName.2 = STRING: "sit0"
   SRA-ftLinux-MIB::ftcEtherInstanceName.3 = STRING: "eth080010"
   SRA-ftLinux-MIB::ftcEtherInstanceName.4 = STRING: "eth080011"
   SRA-ftLinux-MIB::ftcEtherInstanceName.5 = STRING: "eth000010"
   SRA-ftLinux-MIB::ftcEtherInstanceName.6 = STRING: "eth000011"
   ```

2. Run an `snmpwalk` on `ftcEtherDevPathID` OID. This gives you a list of `EtherDevPath` names mapped to `EtherInstance` names:

   # **`snmpwalk -v 1 -c private -t 40 localhost ftcEtherDevPathID`**

   ```
   SRA-ftLinux-MIB::ftcEtherDevPathID.1 = STRING: "[unknown]"
   SRA-ftLinux-MIB::ftcEtherDevPathID.2 = STRING: "[unknown]"
   SRA-ftLinux-MIB::ftcEtherDevPathID.3 = STRING: "[10/5]"
   SRA-ftLinux-MIB::ftcEtherDevPathID.4 = STRING: "[10/5]"
   SRA-ftLinux-MIB::ftcEtherDevPathID.5 = STRING: "[11/5]"
   SRA-ftLinux-MIB::ftcEtherDevPathID.6 = STRING: "[11/5]"
   ```

The instances of interest are the following:

Instance	Instance Name (I/O element / Slot)	Device
ftcEtherInstanceName.3	10/5	eth080010
ftcEtherInstanceName.4	10/5	eth080011
ftcEtherInstanceName.5	11/5	eth000010
ftcEtherInstanceName.6	11/5	eth000011

These are the instances to check for when pulling cables. State changes include DUPLEX, SIMPLEX, BROKEN, and of course, various counters such as frames and collisions. (See "OpState:State Definitions" on page 8-30 for state change identification.)

3. Run snmpwalk on the ftcEtherState OID before and after pulling each cable:

```
# snmpwalk -v 1 -c private -t 40 localhost ftcEtherState

SRA-ftLinux-MIB::ftcEtherState.1 = INTEGER: triplex(22)
SRA-ftLinux-MIB::ftcEtherState.2 = INTEGER: device-ready(13)
SRA-ftLinux-MIB::ftcEtherState.3 = INTEGER: simplex(20)
SRA-ftLinux-MIB::ftcEtherState.4 = INTEGER: device-ready(13)
SRA-ftLinux-MIB::ftcEtherState.5 = INTEGER: simplex(20)
SRA-ftLinux-MIB::ftcEtherState.6 = INTEGER: device-ready(13)
```

In practice, you will actually redirect your snmpwalk output to files for before and after diff comparison. For example, in your work area, run snmpwalk for the entire SRA-ftLinux-MIB file and dump that data to a file. Pull the cable, then run snmpwalk again and dump it to another file.

Finally, run diff on the two files to see all Stratus objects that have changed because of the fault insertion. You may want to put these commands into a shell script for easier testing.

SNMP and MIBS

The SRA-ftLinux-MIB file maps ftServer device definitions for management by Net-SNMP and ftlSNMP. These device definitions map to addressable devices in the /proc virtual file system. ftlSNMP can retrieve operation state data on these devices. The contents of SRA-ftLinux-MIB provide useful remarks about objects that can be managed.

Device Enumeration

See Table 7-1 for information on the enumeration of hardware components for ftServer systems running a supported Linux distribution together with ftSSS.

ftServer System Operation State Management

Figure 8-2 illustrates the operational states and state changes in an ftServer system.

Figure 8-2. Operational State Management on ftServer Systems

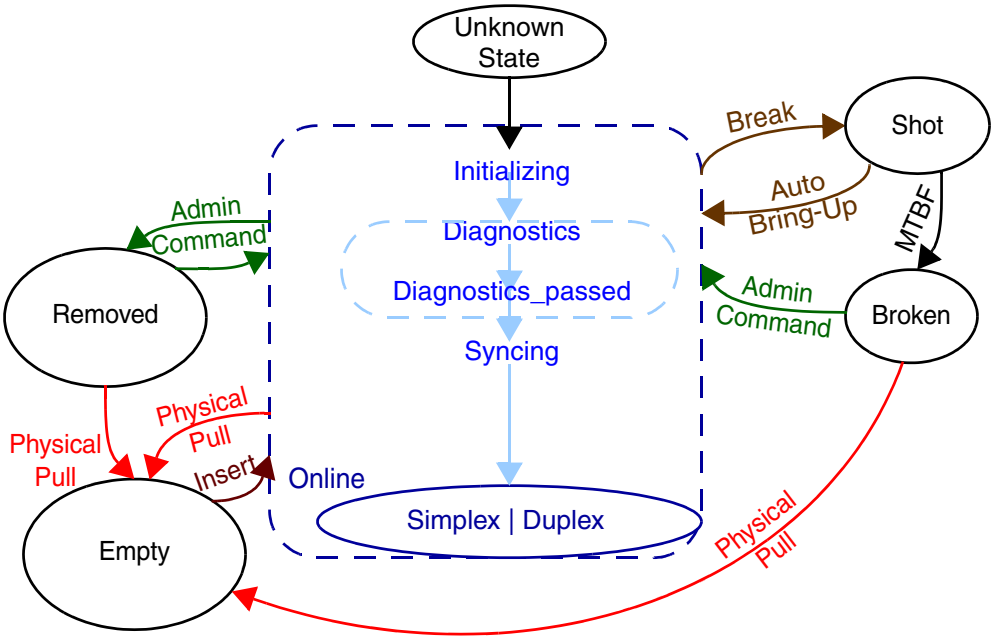

The device physical states can be Empty, or an online state such as Simplex or Duplex. A device can be administratively *Removed*, which the system then treats as an Empty state. However, a Removed device can be restarted by administrative command, while an *Empty* device is restarted by physical removal (if actually present) and insertion of the device into a system slot.

The assigned state *Shot* is an automatically determined or managed state based on a failure threshold or MTBF. A Shot state can be transitioned to an operational state by bringing the device up automatically, or transitioned to Broken through a defined system administrative process. A *Broken* device requires intervention: an administrator can choose to return the device to service or may physically remove the device.

Another state is *Unknown*. The system will try to bring up any device and resolve it to a known state, initializing it and performing diagnostics, and if diagnostics are passed,

bringing the device into an Online state for fault-tolerant operations. A partnered device on an ftServer system typically reaches a Simplex state (if its partner is missing or not functioning) or a Duplex state. The interpretation of Duplex depends on the individual device type, as shown in Table 8-1.

Table 8-1. Meaning of *Duplex* for ftServer System Components

Component	Meaning of Duplex
CPU element	A partner CPU element is present and online, and the two partners are synchronized and running in lockstep.
I/O element	A partner I/O element is present, online, and able to become primary (to assume *active compatibility*).
I/O Device	A partner I/O device (for example, an Ethernet adapter) is present, online, and available for failover.
Disk Drive	A partner disk drive is online, and the partitions of the two partners are mirrored and synchronized.

You can use ftlSNMP to track and log these states, and to control some operations. See Table 8-2 for a complete list of operational states.

SNMP Network Management Station Considerations

The ftlSNMP package provides SNMP subagents. Net-SNMP and ftlSNMP do not provide an SNMP-capable network management station (NMS). However, you can use a commercial or open source NMS to manage the Net-SNMP and ftlSNMP packages remotely; you can also manage these packages directly from a remote system using Net-SNMP. The SRA-ftLinux-MIB file must be provided on the managing host(s) and must support the ftSSS release installed on the managed system(s). If different systems use different Linux operating system releases, the SRA-ftLinux-MIB file must reconcile differences or your SNMP NMS will not be able to manage mismatched object IDs. The SRA-ftLinux-MIB file for a later Linux operating system release will likely work with the earlier Linux operating system releases but may require some adjustment for different defined objects in the `/proc` file system.

The MIBs in ASN.1 encoded text form are located in `/opt/ft/mibs`, `usr/share/snmp/mibs`, and subordinate directories by default. Note that the SRA-ftLinux-MIB file is present in the `/opt/ft/mibs` directory and is named SRA-ftLinux-MIB.txt.

Load all of the MIB files you require into the SNMP NMS; certainly, SRA-ftLinux-MIB will be necessary to manage ftServer objects. Configure the SNMP NMS to avoid verbose OID (object ID) printouts that may clutter the display. The minimum part of the OID needs to be displayed to provide the object's unique name.

NOTES

1. Net-SNMP and ftlSNMP do not require the SNMP NMS, and the package does not provide one. Choice, installation, and configuration of the SNMP NMS is your responsibility.
2. The SRA-ftLinux-MIB file is only useful for managing ftServer LInux-based systems.

Initial SNMP Testing

On a system with an SNMP-aware NMS, you start the NMS before starting SNMP servers.

Start the master agent and the ftlSNMP subagents by typing:

```
# start_snmp
```

After this command completes, master agent and subagent processes with the names **`snmpd`**, `ftlsubagent`, and `ftltrapsubagent` should be running. Verify this:

```
# ps -aux | grep snmnp
# ps -aux | grep ftl
```

In this command, *`process`* is one of the process names listed above. Any errors or warnings generated during the startup script's execution are posted to `syslog` and `stderr`. See "Deploying SNMP Agents and Subagents" on page 8-17 for the default destinations of messages logged by the master agent and subagent processes and "The `ftlsubagent.conf` and `ftltrapsubagent.conf` Files" on page 8-5 for information on how to change the location of the subagent logs.

Note that there are many other commands available for managing Net-SNMP and ftlSNMP. See "ftlSNMP Management Commands" on page 8-8 and *ftlsnmp_scripts*(8) for descriptions of other commands, and read comments in the script files.

To terminate these processes, run the `stop_snmp` command.

Initial Testing of `ftltrapsubagent`

The `ftltrapsubagent.conf` file allows you to control trap filtering. For detailed information about controlling trap filtering, see "Trap Filtering" on page 8-33.

To perform initial testing of `ftltrapsubagent`, determine which system enclosure can be safely brought down.

NOTES

1. Do not use this procedure on a deployed network host.
2. Before continuing, read *ftsmaint*(8) for information on single-digit device path IDs, and "ftServer System Device Path Enumeration" on page 7-6 if you have not already done so.

Select an enclosure that can be safely brought down. To get a listing of options, type:

```
# /opt/ft/bin/ftsmaint ls
```

To get a report on a single device (see "`ftsmaint` Examples" on page 7-11 for an explanation of *hw_path*), type:

```
# /opt/ft/bin/ftsmaint ls hw_path
```

The status display must indicate `DUPLEXED` and `ONLINE` for a system enclosure that can be safely brought down.

NOTE

The `ftsmaint` command allows you to bring down an enclosure only if it has an operational duplexed partner. When an enclosure is operating duplexed, the simple/duplex LED is lit steady white. The simplex/duplex LED is the bottom LED of the three LEDs on the left side of the front of the enclosure, and is the right-most LED on the rear of the enclosure. You can also use `ftsmaint` to see if the enclosures are duplexed.

Suppose that CPU 1, I/O 11 is duplexed. You can bring it down (leaving its partner, CPU 0, I/O 10, functioning) with the following commands:

```
# /opt/ft/bin/ftsmaint bringDown 11
# /opt/ft/bin/ftsmaint bringDown 1
```

NOTE

The term DIAGNOSTICS may appear when remotely bringing down an CPU-I/O enclosure. This term is typically returned if a diagnostic test is in progress, without regard to whether the test will succeed or fail. DIAGNOSTICS is a transient state.

Use the following command to bring the CPU-I/O enclosure up again:

```
# /opt/ft/bin/ftsmaint bringUp 11
# /opt/ft/bin/ftsmaint bringUp 1
```

Initial Testing of `ftlsubagent`

Use the `snmpwalk` tool to perform a `get next` operation on a system where an SNMP master agent is running. See *snmpwalk*(1). For example, for the `ftcPcidevcnf` table:

```
# ./snmpwalk -Os -c public -v 1 -t 40 localhost
   1.3.6.1.4.1.458.107.1.2.5.2.1

 . . .
ftcPcidevcnfMasterDataParityError
iso.3.6.1.4.1.458.107.1.2.5.2.1.14.0 = INTEGER: 2
iso.3.6.1.4.1.458.107.1.2.5.2.1.14.1 = INTEGER: 2
...
ftcPcidevcnfSignaledSERR
iso.3.6.1.4.1.458.107.1.2.5.2.1.15.0 = INTEGER: 2
iso.3.6.1.4.1.458.107.1.2.5.2.1.15.1 = INTEGER: 2
...
ftcPcidevcnfDetectedParityError
iso.3.6.1.4.1.458.107.1.2.5.2.1.16.0 = INTEGER: 2
iso.3.6.1.4.1.458.107.1.2.5.2.1.16.1 = INTEGER: 2
 . . .
```

Notice that the `snmpwalk` tool can provide symbolic decoding of absolute numbers/OIDs.

Removing ftlSNMP

Whenever possible, avoid manual removal of the Net-SNMP and ftlSNMP packages. Use the standard Linux operating system and ftSSS installation and upgrade procedures to manage Net-SNMP and ftlSNMP installations and upgrades. This section documents the manual removal process.

First, stop all server SNMP processes:

```
# /etc/opt/ft/snmp/scripts/stop_snmp
```

This stops the SNMP subagents, followed by the master agent. Any errors or warnings resulting are written to the `syslog` or `stderr`. To remove the installed binary RPM, enter:

```
# rpm -e lsb-ft-snmp
```

OpState:State Definitions

Table 8-2 lists operation state (OpState) names, SRA-ftLinux-MIB codes, and definitions for ftServer systems running a supported Linux distribution together with ftSSS.

Table 8-2. Operation State Values, Names, and Definitions *(Page 1 of 2)*

Value	Operation State (OpState)	Definition
1	UNKNOWN	The state of a component could not be determined.
2	EMPTY	The component slot does not have a component present, or the component does not have power.
3	REMOVED	A component is present in the slot, but main power is not turned on and the component is out of service.
4	SHOT	A component has an error and was taken out of service by system logic. When in this state, the component is electrically isolated from the rest of the system.

Table 8-2. Operation State Values, Names, and Definitions *(Page 2 of 2)*

Value	Operation State (OpState)	Definition
5	BROKEN	A component has a problem; an associated reason (see OpState: Reason) describes the problem. This is a terminal state; some user action must occur to change this state. User actions that cause a transition out of the BROKEN state include bringing the component up or down, or removing the component.
6	DUMPING	A CPU-I/O enclosure is recovering crash dump information.
7	DIAGNOSTICS	A component is running diagnostics.
8	DIAGNOSTICS_PASSED	A component has passed diagnostics.
9	INITIALIZING	Software is preparing a device to be brought online.
11	FIRMWARE_UPDATE	Board firmware code is being updated.
12	FIRMWARE_UPDATE_COMPLETE	Board firmware code is updated.
14	OFFLINE	The unit has been brought down.
15	STOPPED	The driver has stopped the component; the component is no longer running.
19	ONLINE	The unit can be communicated with.
20	SIMPLEX	A component is online and has no partner; it is not safe to remove this component. Applies to components that can be partnered.
21	DUPLEX	The component is online and has a partner component that is running in lockstep, mirrored, or available for failover (depending on the type of component). This component is safe to remove. Applies to components that can be partnered.

OpState:Reason Definitions

Table 8-3 lists reason names, SRA-ftLinux-MIB codes, and definitions for ftServer systems running a supported Linux distribution together with ftSSS.

Table 8-3. Reason Codes, Names, and Definitions

Code	Reason	Definition
1	UNKNOWN	The cause is not known
2	NONE	No reason is available
3	BELOW_MTBF	The current MTBF is below the MTBF threshold specified for this component.
4	DIAGNOSTICS_FAILED	This component failed diagnostic testing.
5	HARDWARE_INCOMPATIBLE	The component hardware is incompatible with the online system hardware.
6	HOLDING_DUMP	Bring-up failed for dump is in process
9	MEDIA_DISCONNECT	Simplex state was entered because a cable was unplugged.
10	FIRMWARE_BURN_FAIL	Failed to update the enclosure's BIOS or firmware.
11	FIRMWARE_FILE_NOT_FOUND	The entered firmware file path is either incorrect or the file does not exist.
12	FIRMWARE_FILE_ERROR	There was an error in the firmware image on disk.
13	FIRMWARE_PROM_ERROR	Could not write to the firmware PROM.
14	AUTOBURN_DISABLED	Cannot match a new enclosure's BIOS or firmware with that of the existing enclosure.
16	PRIMARY	With duplex devices, this indicates that the specific device is primary in the pair.
17	SECONDARY	With duplex devices, this indicates that the specific device is secondary in the pair.

GET and SET Operations for ftlSNMP MIB Objects

See the SRA-ftLinux-MIB.txt file for information on objects that have GET and SET operations. Table 8-4 lists the operations.

Table 8-4. Set Operations Currently Implemented in ftlSNMP

Operation	
ftcCpubdBurnFirmware	ftcIobdInitiateBringUp
ftcCpubdClearMTBF	ftcIobdSetMTBFThreshold
ftcCpubdInitiateBringDown	ftcIobdSetMtbfType
ftcCpubdInitiateBringUp	ftcEtherClearMTBF
ftcCpubdSetCPUBoardPriority	ftcEtherSetMTBFThreshold
ftcCpubdSetMTBFThreshold	ftcEtherSetMtbfType
ftcCpubdSetMtbfType	ftcScaClearMTBF
ftcIobdClearMTBF	ftcScaSetMTBFThreshold
ftcIobdInitiateBringDown	ftcScaSetMtbfType

SRA-ftLinux-MIB OID Values and Properties

The SRA-ftLinux-MIB file contains all ftlSNMP defined object identifiers, and associated operation classes and properties, with some notes inserted to assist you. You can `grep` through the SRA-ftLinux-MIB.txt file for information about OID values and properties or open it with an editor or browser to view and search for information. The default installation places SRA-ftLinux-MIB at `/opt/ft/mibs/SRA-ftLinux-MIB.txt`.

Trap Filtering

This section discusses the following topics:

- "Trap-Filtering Capability"
- "Activating and Deactivating Trap Filtering"
- "Trap-Filtering Examples"

Trap-Filtering Capability

ftlSNMP provides the ability to filter out transitional traps. *Traps* are messages that inform you about network events. Hardware components that go in and out of service trigger a number of traps that are seen at the management client. Some of these traps are actually transitional state information for devices. For example, when you bring up a CPU element, the CPU board's state changes from DIAGNOSTICS to INITIALIZING, ONLINE, and then DUPLEX. However, if you are interested in only the end-state (for example, ONLINE and DUPLEX), the trap-filtering capability is useful.

Another reason to use the trap-filtering capability is that some SNMP traps are triggered by obvious reason codes. For example, when you bring down an I/O element, the display controller with the device path 10/0 or 11/0 will change state from DUPLEX to OFFLINE with the reason code of PARENT_EMPTY. If you are not interested in this type of trap, use the trap-filtering capability.

Activating and Deactivating Trap Filtering

To activate trap filtering, specify the following configuration line in the `/etc/opt/ft/snmp/ftltrapsubagent.conf` file:

```
sraTrapFiltering on
```

When you activate trap filtering, traps with the following operational states are filtered out:

- DUMPING
- DIAGNOSTICS
- DIAGNOSTICS_PASSED
- INITIALIZING
- SYNCING
- FIRMWARE_UPDATE
- FIRMWARE_UPDATE_COMPLETE
- UNKNOWN

Traps with the following reason codes are also filtered out:

- PARENT_EMPTY
- PARENT_BROKEN

To deactivate the trap-filtering capability, specify the following configuration line:

```
sraTrapFiltering off
```

By default, trap filtering is turned off (that is, `sraTrapFiltering` is set to `off` in the configuration file).

Trap-Filtering Examples

Example 8-1 shows some traps that can occur when I/O element 11 is brought down and trap filtering is off.

NOTE

The following examples show sample data only. Data from your system may be different.

Example 8-1. Traps that Can Occur for I/O Element 11 When Trap Filtering Is Off

```
RFC1213-MIB::sysUpTime.0 = Timeticks: (11829) 0:01:58.29
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "10 0"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "SIMPLEX"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "NONE"
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206134938.609247-300"

RFC1213-MIB::sysUpTime.0 = Timeticks: (11929) 0:01:59.29
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "10"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "SIMPLEX"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "PRIMARY"
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206134939.616628-300"

RFC1213-MIB::sysUpTime.0 = Timeticks: (12030) 0:02:00.30
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "10 40 1"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "SIMPLEX"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "NONE"
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206134940.622184-300"
```

```
RFC1213-MIB::sysUpTime.0 = Timeticks: (12130) 0:02:01.30
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "11 40 1"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "ONLINE"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "NONE"
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206134941.625363-300"

RFC1213-MIB::sysUpTime.0 = Timeticks: (12231) 0:02:02.31
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "10 5"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "SIMPLEX"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "NONE"
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206134942.629238-300"

RFC1213-MIB::sysUpTime.0 = Timeticks: (12431) 0:02:04.31
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "11"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "OFFLINE"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "NONE"
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206134944.633611-300"

RFC1213-MIB::sysUpTime.0 = Timeticks: (12435) 0:02:04.35
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "11 0"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "OFFLINE"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "PARENT_EMPTY"
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:

RFC1213-MIB::sysUpTime.0 = Timeticks: (12535) 0:02:05.35
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "11 2"
```

```
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "OFFLINE"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "PARENT_EMPTY"
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206134945.675906-300"

RFC1213-MIB::sysUpTime.0 = Timeticks: (14244) 0:02:22.44
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "11 3"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "OFFLINE"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "PARENT_EMPTY"
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206134945.675906-300"

RFC1213-MIB::sysUpTime.0 = Timeticks: (14444) 0:02:24.44
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "11 4"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "OFFLINE"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "PARENT_EMPTY"
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206135004.762343-300"

RFC1213-MIB::sysUpTime.0 = Timeticks: (17244) 0:02:52.44
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "11 5"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "OFFLINE"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "PARENT_EMPTY"
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"

SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206135032.759280-300"

RFC1213-MIB::sysUpTime.0 = Timeticks: (18744) 0:03:07.44
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "11 6"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "OFFLINE"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "PARENT_EMPTY"
```

```
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206135047.759920-300"

RFC1213-MIB::sysUpTime.0 = Timeticks: (20244) 0:03:22.44
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "10 6"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "SIMPLEX"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "NONE"
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206135102.760559-300"

RFC1213-MIB::sysUpTime.0 = Timeticks: (21744) 0:03:37.44
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "11 40 1"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "REMOVED"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "NONE"
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206135117.760098-300"

RFC1213-MIB::sysUpTime.0 = Timeticks: (23244) 0:03:52.44
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "10 120"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "SIMPLEX"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "PRIMARY"
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206135132.761581-300"


RFC1213-MIB::sysUpTime.0 = Timeticks: (26444) 0:04:24.44
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "11 120"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "OFFLINE"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "NONE"
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206135204.763945-300"
```

Example 8-2 shows some traps that can occur when CPU-I/O enclosure 11 is brought down and trap filtering is on.

Example 8-2. Traps That Can Occur for I/O Element 11 When Trap Filtering Is On

```
RFC1213-MIB::sysUpTime.0 = Timeticks: (4223) 0:00:42.23
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "10 0"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "SIMPLEX"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "NONE"
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206141302.277883-300"

RFC1213-MIB::sysUpTime.0 = Timeticks: (5445) 0:00:54.45
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "10"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "SIMPLEX"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "PRIMARY"
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206141314.504766-300"

RFC1213-MIB::sysUpTime.0 = Timeticks: (5546) 0:00:55.46
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "11"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "OFFLINE"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "NONE"


SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206141315.508290-300"

RFC1213-MIB::sysUpTime.0 = Timeticks: (5746) 0:00:57.46
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "10 40 1"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "SIMPLEX"
```

```
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "NONE"
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206141317.514636-300"

RFC1213-MIB::sysUpTime.0 = Timeticks: (5847) 0:00:58.47
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "11 40 1"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "ONLINE"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "NONE"
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206141318.517971-300"

RFC1213-MIB::sysUpTime.0 = Timeticks: (5948) 0:00:59.48
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "10 6"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "SIMPLEX"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "NONE"
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206141319.534567-300"

RFC1213-MIB::sysUpTime.0 = Timeticks: (6064) 0:01:00.64
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "10 5"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "SIMPLEX"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "NONE"
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206141320.697382-300"

RFC1213-MIB::sysUpTime.0 = Timeticks: (6845) 0:01:08.45
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "10 120"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "SIMPLEX"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "PRIMARY"
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206141328.502684-300"
```

```
RFC1213-MIB::sysUpTime.0 = Timeticks: (8545) 0:01:25.45
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "11 120"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "OFFLINE"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "NONE"
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206141345.505798-300"

RFC1213-MIB::sysUpTime.0 = Timeticks: (8846) 0:01:28.46
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "11 40 1"
SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE"
SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "REMOVED"
SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "NONE"
SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN"
SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051206141348.510919-300"
```

Example 8-3 shows some traps that can occur when CPU 1, I/O 11 is brought up and trap filtering is off.

Example 8-3. Traps That Can Occur When Trap Filtering Is Off

```
RFC1213-MIB::sysUpTime.0 = Timeticks: (461810) 1:16:58.10
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "2
"SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE
"SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "DIAGNOSTICS
"SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "NONE


"SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN
"SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051207143640.794126-300"

RFC1213-MIB::sysUpTime.0 = Timeticks: (466169) 1:17:41.69
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "2
"SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE
"SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "INITIALIZING
```

```
"SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "NONE
"SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN
"SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051207143724.381849-300"

RFC1213-MIB::sysUpTime.0 = Timeticks: (467013) 1:17:50.13
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "0
"SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE
"SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "DUPLEX
"SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "PRIMARY
"SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN
"SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051207143732.824273-300"

RFC1213-MIB::sysUpTime.0 = Timeticks: (467114) 1:17:51.14
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "2
"SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE
"SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "DUPLEX
"SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "SECONDARY
"SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN
"SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051207143733.828058-300"
```

Example 8-4 shows some traps that can occur when CPU 1, I/O 11 2 brought up and trap filtering is on.

Example 8-4. Traps That Can Occur When Trap Filtering Is On

```
RFC1213-MIB::sysUpTime.0 = Timeticks: (10411) 0:01:44.11
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "0
"SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE
"SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "DUPLEX
"SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "PRIMARY
"SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN
"SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051207144108.553784-300"

RFC1213-MIB::sysUpTime.0 = Timeticks: (10511) 0:01:45.11
SNMPv2-MIB::snmpTrapOID.0 = OID:
SRA-ftLinux-MIB::ftcTrapGenericInformationTrap.0
SRA-ftLinux-MIB::ftcTrapDevicePathID.0 = STRING: "2
```

```
"SRA-ftLinux-MIB::ftcTrapAlertType.0 = STRING: "OPSTATE_CHANGE
"SRA-ftLinux-MIB::ftcTrapGenName.0 = STRING: "DUPLEX
"SRA-ftLinux-MIB::ftcTrapGenDetailInfo.0 = STRING: "SECONDARY
"SRA-ftLinux-MIB::ftcTrapGenAction.0 = STRING: "UNKNOWN
"SRA-ftLinux-MIB::ftcTrapGenEventTimeStampWithOffsetFromUTC.0 = STRING:
"20051207144109.557164-300"
```

Chapter 9
Troubleshooting ftServer Systems

This chapter discusses the following topics:

- ‘‘LED and Visual Diagnostics”
- ‘‘LED and Visual Diagnostics”
- ‘‘System Log Messages”

This chapter provides information that will help you use available ftServer system and Linux operating system features to diagnose system problems. In many cases, you will be able to identify the source of the problem. If you cannot, contact the Stratus Customer Assistance Center (CAC) or your authorized Stratus service representative.

LED and Visual Diagnostics

Stratus ftServer systems have a number of light-emitting diodes (LEDs) and indicator lamps that can provide information of diagnostic value. For a complete explanation of the location and interpretation of LEDs in your system, see the troubleshooting chapter in the operation and maintenance guide for your system.

System Boot Problems

If you experience problems in booting the system, the following information may help you diagnose the problem.

- “Normal Boot Sequence”
- “Possible Boot Problems”
- ‘‘Error and Log Messages Regarding Keyboard and Mouse”

Also, refer to ‘‘Booting in Linux Rescue Mode’’ on page 2-15 for a procedure to boot the operating system after a failed installation.

Normal Boot Sequence

The active CPU-I/O enclosure (that is, the primary enclosure whose power switch is lit green) initiates the boot by starting the BIOS. The BIOS on the booting enclosure scans the list of bootable devices (as configured in the BIOS) looking for a device to boot. When the search finds the disks, they are analyzed from bottom to top in the boot enclosure. When a disk with a boot partition is found, it is booted.

The boot sequence is a multi-stage process:

- The GRUB boot loader is loaded and started.
- The GRUB boot loader loads the second stage GRUB, which selects the Linux kernel to boot, loads, and starts it. The boot process uses the raw disk; RAID is not involved.

NOTE

For the boot sequence to work, a bootable medium must be found in the active CPU-I/O enclosure. For example, if you try to boot the top CPU-I/O enclosure, the boot sequence fails if sda is missing and there is no other bootable disk. If a boot partition is found in sdb or sdc, it is booted. If the boot fails, the system switches the active CPU-I/O enclosure and tries again (where it will find and boot sdd, if present).

- When the kernel starts, it loads the SCSI driver, which in turn loads the SATA driver.
- The SATA driver spins up all the disks it discovers. It also collects a list of all of the type 0xfd (Linux RAID autodetect) partitions.
- When the RAID-1 module is loaded, it processes this list of mirrors and starts RAID arrays. It processes mirrors from disks in the following order: sda, sdd, sdb, sde, sdc, and sdf. If a RAID array was not cleanly shut down, a resync is started.

NOTES

1. The RAID arrays that are started are started as described above, regardless of what kernel was booted (or which disk contained the kernel).
2. The recommended configuration of sda and sdd system disks results in the expected boot sequencing. The kernel is found in sda or sdd, and the system file systems are also on sda and sdd.

- Later in the boot sequence, `/etc/rc.sysinit` runs. It finds and starts any RAID arrays in `/etc/mdadm.conf` that were not already started and that are required by mounts in `/etc/fstab`.

NOTE

If a RAID array fails to start, the boot stops and enters a debug shell. This is almost always because of a configuration error in `/etc/fstab` or `/etc/mdadm.conf`. Exiting the shell forces a reboot.

Depending on your system's RAID configuration, you may see one or more error messages similar to the following:

```
md: could not bd_claim sdar1
md: error, md_import_device() returned -16
```

These messages indicate that md is refusing to start an array that has already been started. You can safely ignore them.

The recommended configuration has all RAID-1 arrays marked as type 0xfd (Linux RAID autodetect) so they start early, and all RAID-0 arrays in `/etc/mdadm.conf`, so they start later.

- The operating system checks the file systems.

NOTE

In the case of crash recovery, the file check (`fsck`) may take a long time, and it may fail. If it fails, the boot stops and enters a debug shell. The administrator must manually repair the problem file systems. Exiting the shell forces a reboot.

Possible Boot Problems

A problem in booting the system may be associated with missing or corrupt fault-tolerant drivers, the GRUB boot loader, or RAID.

Missing Stratus Drivers Prevent Booting

If required fault-tolerant drivers are not present at boot time, and if the system's fault-tolerant policy is set to prevent booting when Stratus drivers are missing (the default setting), the following prompt appears at the console:

```
This system is not fault tolerant because reason
Type "NON-FT-BOOT" to allow login for repair:
```

In this output, *logfile* is the name of a file that contains relevant details and *reason* is one of the following:

- `ERROR building Stratus kernel objects -- see logfile`
- `ERROR: missing Stratus kernel objects -- see logfile`
- `ERROR: incorporating Stratus kernel objects -- see logfile`

To override the system's fault-tolerant policy and allow the system to boot to a non-fault-tolerant state, at the console, type `NON-FT-BOOT` and press **Enter**.

If you provide any other response three times, the system starts the boot process again.

GRUB Problem

If the system boots and hangs before the operating system is loaded, it may be a problem with the GRUB boot loader. Reinstall GRUB in the master boot record (MBR) on the problem disk.

To manually run GRUB

1. Boot the system from your Linux operating system bootable CD by performing the procedures in "System Log Messages" on page 9-6.
2. Follow the prompts instructing you to run the `chroot` command.
3. Run GRUB as follows:

```
# /sbin/grub
grub> device (hd0) /dev/sda
grub> root (hd0,0)
grub> setup (hd0)
grub> device (hd0) /dev/sdb
grub> root (hd0,0)
grub> setup (hd0)
grub> quit
```

   Both system disks are now bootable.

CAUTION

Incorrect GRUB parameters can also cause problems in booting. Do not change the parameters from the defaults set when the operating system was installed. In particular, specifying the GRUB `noapic` option can make the operating system unbootable.

4. Shutdown the system and eject the CD.

RAID Problem

If a RAID-1 array has one type 0xfd (Linux RAID autodetect) mirror and one 0x83 (Linux) mirror, at boot, the RAID array is started in degraded mode using the type 0xfd mirror, and the type 0x83 mirror is not automatically added. You can add the mirror with `mdadm`. To fix this problem, just change the partition type with `fdisk`.

The system supports RAID-1 arrays that consist of type 0x83 mirrors.

You can create partitions of type 0x83, create RAID-1 arrays with them, and then create a RAID-0 array from the RAID-1 arrays. If you want to start the RAID-0 arrays automatically, add entries for them to `/etc/mdadm.conf`. Otherwise, the RAID-0 arrays are not started.

Automatic Reboot After Boot Monitoring Timeout

When the system is booted into certain modes, such as RAID repair mode, the system heartbeat is not enabled. After a defined period (the default is 10 minutes), the system is automatically rebooted if a heartbeat has not been received. If your troubleshooting and repair requires more than the defined period, you must disable boot monitoring in the BIOS during the boot sequence.

To disable boot monitoring

1. When the system is booting and the progress bar has started to fill, press **F2** to enter the BIOS Setup program. An `Entering Setup` message appears, but it may take several minutes for the BIOS Setup program to run.
2. Use the right arrow (→) key to select the `Stratus` tab.
3. On the `Stratus` menu, use the down arrow (↓) key to select `Monitoring Configuration`. Press **Enter**.
4. On the `Monitoring Configuration` menu, use the down arrow (↓) key to select `Boot Monitoring.`
5. Select `Boot Monitoring` and use the plus sign **(+)** key to change the value to `Disabled`.
6. Press **ESC** to exit from the submenu.
7. In the `Setup Confirmation` dialog box, select `Yes` and press **Enter** to save the new settings and exit from the BIOS Setup program.

After resolving the problem, reenable boot monitoring during the next boot by following the same procedure, but in step 5, change the value to `Enabled`.

CAUTION

Boot monitoring is a fault-tolerant feature of your ftServer system. You must reenable it for full fault tolerance.

System Log Messages

System log messages contain information on the operation state of the system. The file `/var/log/messages` contains system log messages. You can find logs that are specific to ftServer systems in the directory `/var/opt/ft/log`.

See “OpState:State Definitions” on page 8-30 and “OpState:Reason Definitions” on page 8-32 for explanations of some of the terminology you may see in these messages.

Error and Log Messages Regarding Keyboard and Mouse

In the system log or at system boot, you may see `stderr` messages such as those shown in Example 9-1. These and similar messages may occur multiple times. They are not a cause for concern if the system boots without undue delay, and when the operating system presents you with the logon prompt, it is appropriately presented on your display device and your input devices are supported to interact with the system. The messages are an unavoidable result of the order in which drivers need to be loaded during the Linux operating system distribution boot process.

Example 9-1. Possible Keyboard and Mouse Error Messages at Boot Time

```
Mar 15 07:15:00 ftlx rc.sysinit: Initializing USB keyboard: failed
Mar 15 07:15:00 ftlx modprobe: modprobe: Can’t locate module mousedev
Mar 15 07:15:00 ftlx rc.sysinit: Initializing USB mouse: failed
Mar 15 07:15:07 ftlx ifup: SIOCSIFADDR: No such device

Mar 15 07:15:51 ftlx kernel: Keyboard timed out[1]
kernel: keyboard timed out[1]
Mar 15 07:15:52 ftlx kernel: keyboard: Timeout - AT keyboard not present? (f4)
```

Appendix A
Linux Packages

The ftServer System Software for the Linux Operating System CD-ROM provided with your system includes Red Hat Package Manager (RPM) files that allow you to manage Stratus ftServer systems running a supported Linux distribution together with ftSSS.

On the CD-ROM, RPM files are located in the following subdirectories of the `/CD_ROOT` directory (where `CD_ROOT` is the location at which the CD-ROM is mounted: for example, `/media/cdrom`): `RPMS` and `SRPMS`.

- The `RPMS` directory contains software packages that must be installed in order to provide fault-tolerant capabilities on your ftServer system.

  CAUTION

  Be careful not to replace any of these packages with versions other than those included in your Stratus Linux Operating System distribution. Doing so can impair the functionality of your system.

- The SRPMS directory provides sources for all packages controlled by open-source licenses.

To view information about an RPM file, go to one of the RPM directories (`RPMS` or `SRPMS`) and type the **`rpm`** command (see *rpm*(8)).

In Example A-1, the `rpm -qpi` command displays information about the ASN RPM file, `lsb-ft-asn-4.0-77.x86_64.rpm` which is a proprietary package (as indicated in the output by `License: Stratus Proprietary`).

Example A-1. Displaying Information About a Stratus-Proprietary RPM

```
[pubs@ariapp]$ rpm -qpi RPMS/lsb-ft-asn-4.0-77.x86_64.rpm
Name        : lsb-ft-asn      Relocations: (not relocatable)
Version     : 4.0             Vendor: Stratus Technologies Be
muda Ltd.
Release     : 77              Build Date: Thu 27 Apr 2006 02:36:4
 AM EDT
Install Date: (not installed) Build Host: linuxbuild3.sw.stratus.com
Group       : Applications/System   Source RPM: lsb-ft-asn-4.0-77.src.rm
Size        : 248965                         License: Stratus Proprietary
Signature   : DSA/SHA1, Thu 27 Apr 2006 03:08:00 AM EDT, Key ID
ab60ce4cbb49da4
Packager    : ftlbuild
URL         : http://www.stratus.com
Summary     : Stratus ASN implementation
Description :
Software components for communication between the ftServer and the Stratus
ActiveService Network Hub.

Configuration file is to be obtained via the ActiveService Manager (ASM)
from:

http://www.stratus.ecacsupport.com
[pubs@ariapp RPMS]$
```

NOTE

Source files for proprietary ftSSS files are not available.

In Example A-2, the `rpm -qpil` command displays information about the `lsb-ft-lsscsi-4.0-77.src.rpm` file, which is an open-source Linux distribution package (as indicated in the output by `License: GPL`). In this example, the `l` option lists the files in the package.

Example A-2. Displaying Information About an Open-Source RPM

```
[pubs@ariapp SRPMS]$ rpm -qpil lsb-ft-lsscsi-4.0-77.src.rpm
Name        : lsb-ft-lsscsi       Relocations: (not relocatable)
Version     : 4.0                 Vendor: Stratus Technologies Ber
muda Ltd.
Release     : 77                  Build Date: Thu 27 Apr 2006 02:40:46
 AM EDT
Install Date: (not installed)     Build Host: linuxbuild3.sw.stratus.com
Group       : Applications/System           Source RPM: (none)
Size        : 179804                           License: GPL
Signature   : DSA/SHA1, Thu 27 Apr 2006 03:08:09 AM EDT, Key ID
ab60ce4cbb49da4e
Packager    : ftlbuild
URL         : http://www.stratus.com
Summary     : List SCSI devices (or hosts) and associated information
Description :
Uses information provided by the sysfs pseudo file system in Linux kernel
2.6 series to list SCSI devices or all SCSI hosts. Includes a "classic"
option to mimic the output of "cat /proc/scsi/scsi" that has been widely
used prior to the lk 2.6 series.

Author: Doug Gilbert <dgilbert at interlog dot com>
lsb-ft-lsscsi.spec
lsscsi-0.15.tgz
```

Index

A

ASN (ActiveService Network), 7-1, 7-14
- architecture, 1-1
- optional software, 1-3

B

backing up a system, 4-9, 5-25
- creating disk, 4-9

BIOS
- changing settings of Setup program, 9-5
- firmware, 2-7, 3-1
- updating, 3-1

BMC firmware, 2-7
- updating, 3-5

bonding. *See* channel-bonding interfaces

boot
- media, 2-3
- monitoring
  - disabling, 2-11, 9-5
  - enabling, 2-14
- problems, 9-1

booting the system
- disabling Boot Monitoring, 9-5
- error message, 9-3
- Linux rescue mode, 2-15
- normal sequence, 9-2
- problems, 9-1, 9-3

bring-down
- safe, 8-28

build number
- determining, 7-5

C

CD-ROMs
- Red Hat Linux, 2-2
- with ftSSS distribution, 1-2, 2-2
- with Linux distribution, 1-2, 2-2

channel-bonding interfaces, 5-27
- configuring, 5-28
- determining device names, 5-29
- monitoring, 5-27

checklists
- pre-installation, 2-7
- pre-upgrade, 4-3
- restoration, 4-3

clock, 5-33

configuring
- DNS resolution, 5-31
- Ethernet devices, 2-5
- ftlSNMP, 8-2
- IP adress for bond interfaces, 5-31
- Linux operating system, 1-3
- RAID arrays, 5-8
- static routes, 5-31
- system host name, 5-31
- system time zone, 5-31

console log, 5-2

controlling system devices, 7-14

Coordinated Universal Time (UTC), 5-33

CPU frequency changed message, 5-33

creating a file system, 5-13

D

data rate of serial ports, 2-6

default
- configuration notes, 2-16
- Ethernet configuration, 2-5
- Linux operating system, 2-3
- SATA drive
  - configuration, 5-6
  - settings, 2-5
- system initialization, 2-6

default password, 2-2

device enumeration, 7-5, 8-25

device names, 6-2

disaster recovery, 5-25

disk drives
inserting, 5-15
removing, 5-14
replacing, 5-18
distribution
Linux operating system, 2-2
CD-ROMs, 1-2, 2-2
separately released and optional components, 2-6
documentation, 1-3, 4-10, 5-33
ftServer systems, 1-4
Linux operating system, 1-5
UNIX, 1-5
drivers, missing, 9-3
dumps, 7-16
system, 7-16
duplex, 8-31
LED indicator, 8-28

E

error log messages
keyboard, 9-6
mouse, 9-6
Ethernet configuration, 2-5
channel bonding, 5-27
MAC addresses, 5-30
naming Ethernet devices, 5-25
PCI adapters, adding, 5-30
testing Ethernet ports, 8-23

F

failed disk, replacing, 5-18
fault tolerance, 7-1
ftSSS software, 1-3
hardened drivers, 5-1
hardware, 1-1
RAID disk arrays, 1-1, 5-8
file systems
availability, 5-8
RAID, 5-8
files and upgrade procedure, 4-2
firmware, 2-7
BIOS, 2-7
BMC, 2-7
SATA drives, 2-5
updating, 2-3, 3-1, 3-5
floppy drives, 6-4
flow control of serial ports, 2-6
ftlSNMP, 8-1
agents, 8-17
configuration files, 8-4
configuring, 8-2
for remote service management, 8-16
for Service Management, 8-6
description, 8-2
extensions, 8-15
GET operations, 8-33
installing, 8-1, 8-2, 8-3
inventory, 8-2
management commands, 8-8
managing, 8-18
MIBs, 8-11, 8-24
objects defined by, 8-12
prerequisites, 8-4
removing, 8-30
SET operations, 8-33
SRA-ftLinux-MIB, 8-26
subagents, 8-17
testing the configuration, 8-19
uninstalling, 8-30
upgrading, 8-3
verifying traps, 8-18
ftlsubagent
description, 8-2
initial testing, 8-29
ftltrapsubagent
description, 8-2
initial testing, 8-28
ftServer systems
documentation, 1-4
firmware, 3-1
system administration, 1-1
troubleshooting, 1-3, 9-1
`ftsmaint` command, 7-1
`acSwitch`, 7-2
`bringDown` *`path`*, 7-3
`bringUp` *`path`*, 7-3
`burnProm` *`fw_file path`*, 7-3
`clearMtbf` *`path`*, 7-3
`dump` *`path`*, 7-3
examples, 7-11
`identify [start|stop]` *`path`*, 7-3
`ls` *`path`*, 7-2
`lsLong`, 7-2
`lsPeriph`, 7-2
`lsVND`, 7-2
`powerOff modem`, 7-3

`ftsmaint` command (*Continued*)
`powerOn modem`, 7-3
`reset modem`, 7-3
`resetMtbf` *`path`*, 7-4
`runDiag` *`path`*, 7-4
`setMtbfThresh` *`value path`*, 7-4
`setMtbfType` *`policy path`*, 7-4
`setPriority` *`level path`*, 7-4
`setSensorThresh` *`th_name value path`*, 7-5
task arguments, 7-1
`-version`, 7-5
ftSSS
operational states, 8-25
recovering from an upgrade, 4-9
restoring, 4-7
upgrading, 1-2
using, 1-3

G

GRUB boot loader, 9-2
documentation, 2-19
problems, 9-4
shell, 2-6

H

hardware
fault-tolerant, 1-1
managing, 8-20
self checking, 1-1
supported, 2-7
host name, configuring, 5-31
hyperthreading, disabling, 5-31

I

installing
ftlSNMP, 8-1, 8-2
ftSSS
after a failed attempt, 2-15
overview, 2-2
restoring an installation, 4-7
Linux operating system, 1-2, 2-1, 4-1
installer interfaces, 2-7
overview, 2-2
pre-installation checklist, 2-7
remote network management services, 8-16

J

jumpswitch, 3-4

K

kernel memory dump file management, 7-16

L

LEDs, 2-6, 9-1
simplex/duplex indicator, 8-28
legacy devices, 5-1
USB storage device, 6-3
Linux operating system, 1-1, 2-7
configuring, 1-3
default setup, 2-3
distribution, 2-2
installation, 1-2, 2-1
network administration, 1-2
packages, A-1
pre-installation checklist, 2-7
recovering from an upgrade, 4-9
restoring, 4-1
system administration, 1-2
upgrading, 1-2
version information, 2-4
`logrotate`, 8-6
logs
console, 5-2
message, 9-6
lspci command
and fault tolerance, 5-30
`lspci` command, 5-29
`lsusb` command, 6-2

M

MAC addresses, 5-30
management
data storage devices, 1-3, 6-1
partitions, 5-3
SNMP, 8-1
memory dumps, 7-16
MIB (management information base), 8-11, 8-24
objects defined by, 8-12
SRA-ftLinux-MIB, 8-11
missing drivers, 9-3
mounting a file system, 5-13

N

naming devices, 6-2
Net-SNMP, 8-1, 8-26
    basic commands, 8-9
    description, 8-2
network administration, 1-2
network management stations (NMS), 8-26
Network Time Protocol (NTP)
NTP. *See* Network Time Protocol

O

operating system upgrade, 4-2
operational states
    managing, 8-25
    state definitions, 8-30
OpState
    reason definitions, 8-32
    state definitions, 8-30
optional software components, 2-6

P

partition tables
    creating, 5-5
    displaying, 5-4
partitions
    adding, 5-5
    managing, 5-4
password
    default root, 2-2
    valid, 2-2
pedestal system orientation, 2-4

R

rack system orientation, 2-4
RAID, 5-8
    arrays, 5-2, 5-8, 5-9
        checking the current state, 5-13
        creating and mounting a file system, 5-13
    problems, 9-5
    RAID-0, 5-8
    RAID-1, 5-8
        configured drives during installation, 2-9
        fault tolerance, 1-1
        system pair, 2-17
        two CPU-I/O enclosures, 5-2
    resynchronization, 5-17
README file for SNMP, 8-3
reason definitions, 8-32
recovering
    from a disaster, 5-25
    from an upgrade, 4-9
Red Hat Linux
    CD-ROMs, 2-2
    documentation, 1-4
    upgrading, 4-5
reinstalling
    Linux operating system, 4-6
    ftSSS, 2-15
remote network management services, 8-16
removing ftlSNMP, 8-30
rescue mode for Linux, 2-15
restoring the ftSSS, 4-7
restoring the Linux operating system
    preparation checklist, 4-3
    recovery from, 4-9
resynchronization of RAID mirrors, 5-17
root password (default), 2-2
`rpm` command, A-1
RPM files, 8-3, A-1
runlevel-controlled process configuration, 2-6

S

safe bring-down, 8-28
SATA drives
    checking, 5-7
    default configuration, 2-5, 5-6
    firmware, 2-5
    RAID, 2-5
SCSI drives
    disk names, 5-3
    management
        partitions of disks, 5-3
    storage allocation, 5-6
    supported SCSI disks, 5-2
    system-console messages, 5-2
SCSI subsystem errors, 5-16
serial ports
    data rate, 2-6
    flow control, 2-6
setting up the Linux operating system, 2-3

shells
debug, 9-3
GRUB, 2-6
simplex state
defined,, 8-31
LED indicator, 8-28
SNMP (Simple Network Management Protocol), 8-1
See also ftlSNMP *and* Net-SNMP
basic Net-SNMP commands, 8-9
concepts, 8-8
configuring for remote service management, 8-16
configuring for service management, 8-6
configuring to start at system initialization, 8-6
fault-tolerant operation, 8-15
initial testing, 8-27
managing, 8-18
managing hardware, 8-20
testing a configuration, 8-19
testing Ethernet ports, 8-23
traps.*See* traps
view of a network, 8-14
`snmpset` command, 8-20
`snmptrapd`, 8-18
`snmpusm` command, 8-6
`snmpwalk`, 8-19
software
optional components, 2-6
packages, A-1
SRA-ftLinux-MIB, 8-26
description, 8-2
GET operations, 8-33
OID values and properties, 8-33
SET operations, 8-33
storage devices, definition, 5-8
subagents, log files, 8-17
system
administration, 1-1, 1-2
backup, 5-25
boot
Linux rescue mode, 2-15
normal sequence, 9-2
problems, 9-1, 9-3
devices, controlling, 7-14
initialization, 2-6
log messages, 9-6
pedestal and rack orientation, 2-4
system clock, 5-33
system dumps, 7-16
system firmware, 2-3, 2-7

T

testing
Ethernet ports, 8-23
SNMP configuration, 8-19
ticks (clock), losing message, 5-33
time zone, configuring, 5-31
traps, filtering
activating, 8-34
deactivating, 8-34
description, 8-34
examples, 8-35
traps, verifying, 8-18
troubleshooting ftServer systems, 1-3, 1-3, 9-1

U

`udev` command, 6-2
`udevinfo` command, 6-3
`umount` command, 6-3
`uname` command, 2-4
up2date agent, 4-5
updating
BMC firmware, 3-5
system BIOS, 3-1
Linux operating system, 4-1
upgrading
ftSSS, 1-2, 4-7
Linux operating system, 1-2, 4-1, 4-2
pre-upgrade checklist, 4-3
recovery from an upgrade, 4-9
USB devices
during system installation, 2-5
floppy drives, 6-4
restoring after enclosure failure, 6-4
solid-state, 6-4
storage, 6-2
UTC. *See* Coordinated Universal Time

V

version, operating system identification, 2-4
visual diagnostics, 9-1
VND (Virtual Network Devices), 5-27
group configuration, 5-27
default, 5-27