



**SGI® InfiniteStorage Server 3000 (ISS3000)
User's Guide**

007-5721-001

COPYRIGHT

© 2010, 2011 SGI. All rights reserved; provided portions may be copyright in third parties, as indicated elsewhere herein. No permission is granted to copy, distribute, or create derivative works from the contents of this electronic documentation in any manner, in whole or in part, without the prior written permission of SGI.

LIMITED RIGHTS LEGEND

The software described in this document is “commercial computer software” provided with restricted rights (except as to included open/free source) as specified in the FAR 52.227-19 and/or the DFAR 227.7202, or successive sections. Use beyond license provisions is a violation of worldwide intellectual property laws, treaties and conventions. This document is provided with limited rights as defined in 52.227-14.

The electronic (software) version of this document was developed at private expense; if acquired under an agreement with the USA government or any contractor thereto, it is acquired as “commercial computer software” subject to the provisions of its applicable license agreement, as specified in (a) 48 CFR 12.212 of the FAR; or, if acquired for Department of Defense units, (b) 48 CFR 227-7202 of the DoD FAR Supplement; or sections succeeding thereto. Contractor/manufacturer is SGI, 46600 Landing Parkway, Fremont, CA 94538.

TRADEMARKS AND ATTRIBUTIONS

Silicon Graphics, SGI, the SGI logo, Roamer, Supportfolio, and XFS are trademarks or registered trademarks of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries worldwide.

3ware is a registered trademark of 3ware, Inc.

AMD Opteron is a trademark or registered trademark of Advanced Micro Devices Corporation.

Intel and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds.

LSI is a registered trademark of LSI Corporation.

PCIe and PCI-X are registered trademarks of PCI SIG.

Red Hat and all Red Hat-based trademarks are trademarks or registered trademarks of Red Hat, Inc. in the United States and other countries.

SUSE LINUX and SUSE logo are registered trademarks of Novell, Inc.

Windows is a registered trademark of Microsoft Corporation in the United States and/or other countries.

All other trademarks mentioned herein are the property of their respective owners.

Record of Revision

Version	Description
001	March 2011 Original printing.

Contents

	About This Guide	ix
	Audience.	ix
	Related Publications.	x
	Product Support	xi
	Reader Comments	xii
1	Overview	1
	The ISS3000 3U Server Series	1
	General Features.	2
	Configurable Options	2
	Support Services.	3
	Key Components of ISS3000 3U Servers	4
	Power Supply Options	8
	AC	8
	DC	8
2	Pre-Installation Guidelines	9
	Server and Rack Positioning	9
	Electrical Power	10
	Operating Environment	10
	Rack Stability	11
	Cables	11
	Security	11
3	Installation Guidelines	13
	Getting Ready to Install	14
	Determine Server Positioning in the Rack	14
	Installing into a Rack or Cabinet	15
	Installing into a Telco Rack	15

	Cabling	16
4	Installing Hard Drives	17
	Drive Insertion Order.	17
	Installing a Hard Drive into the Hard Drive Carrier	20
5	Using the Server.	21
	Powering the Server On	21
	Powering the Server Off	22
	Working with Roamer Remote Management	22
6	RAID Configuration	23
	Configuring Hardware RAID	23
	LSI Hardware RAID (SAS or SATA)	23
	3ware-Based Hardware RAID (SATA)	24
	Configuring Software RAID	24
7	Software Installation and Tuning	25
	Linux Server Configuration	25
	Operating System Installation	25
	Filesystem Selection.	26
	Data Drive Configuration (Hardware RAID)	26
	Exporting Filesystems	27
	Tuning Linux NFS Server Performance.	27
	Tuning Linux Client Performance	27
	Using Smartmontools Software to Improve Availability	28
	Windows Server Configuration	28
A	Technical Specifications	29
B	Safety Guidelines	31
	Electrical Safety	31
	General Safety	32
	Handling Your Server	33
C	RAID Basics	35
	RAID Implementations	35
	RAID 0	36

JBOD (Just a Bunch of Disks) 36
RAID 1 36
RAID 5 37
RAID 6 (Dual Parity) 37

About This Guide

This guide provides an overview of the SGI® InfiniteStorage Server 3000 (ISS3000) along with instructions for system installation and general operation. The appendices include technical specifications, safety practices, and an overview of RAID.

Audience

This guide is written for owners, installers, system administrators, and users of ISS3000 servers. It is written with the assumption that the reader has a good working knowledge of computers and computer systems.

Related Publications

Depending on your choice of system components and operating system, the following SGI documents may be relevant to your ISS3000 server:

- *SGI InfiniteStorage Server 3500 (ISS3500) User's Guide* (007-5653-00x)
- SGI InfiniteStorage series documentation
- Man pages (online)

You can obtain SGI documentation, release notes, or man pages in the following ways:

- Refer to the SGI Technical Publications Library at <http://docs.sgi.com>. Various formats are available. This library contains the most recent and most comprehensive set of online books, release notes, man pages, and other information.
- Refer to the SGI Supportfolio webpage for documents whose access require a support contract. See “[Product Support](#)” on page xi.
- You can also view man pages by typing `man <title>` on a command line.

Product Support

SGI provides a comprehensive product support and maintenance program for its products. SGI also offers services to implement and integrate Linux® applications in your environment.

- Refer to <http://www.sgi.com/support/>
- If you are in North America, contact the Technical Assistance Center at +1 800 800 4SGI or contact your authorized service provider.
- If you are outside North America, contact the SGI subsidiary or authorized distributor in your country.

Be sure to have the following information before you call Technical Support:

- Product serial number
- Product model name and number
- Applicable error messages
- Add-on boards or hardware
- Third-party hardware or software
- Operating system type and revision level
- Motherboard BIOS revision

Reader Comments

If you have comments about the technical accuracy, content, or organization of this document, contact SGI. Be sure to include the title and document number of the manual with your comments. (Online, the document number is located in the front matter of the manual. In printed manuals, the document number is located at the bottom of each page.)

You can contact SGI in any of the following ways:

- Send e-mail to the following address: techpubs@sgi.com
- Contact your customer service representative and ask that an incident be filed in the SGI incident tracking system.
- Send mail to the following address:

SGI
Technical Publications
46600 Landing Parkway
Fremont, CA 94538

SGI values your comments and will respond to them promptly.

Overview

The ISS3000 3U server series is one of the best and most reliable servers on the market today. It is designed with some of the most state-of-the-art features available, using non-proprietary components to provide the most scalable and cost-effective server solution possible.

This chapter describes the ISS3000 3U server series, its hardware and service features, as well as key components of the server.

The ISS3000 3U Server Series

The ISS3000 3U server series consists of the following models:

Model	Description
3006	A maximum of 6 hard drives
3009	A maximum of 9 hard drives
3012	A maximum of 12 hard drives
3118	A maximum of 18 hard drives

If your server was ordered from SGI and was not constructed in one of the standard 3U form factors that are detailed in this manual, then this manual will provide a general guideline to your product. Details that are specific to your design can be found in the datasheet or custom documentation that came with your product. For example, standard-depth models like the ISS3500 server have separate user's guides.

General Features

Among the distinctive features of the ISS3000 3U servers are the following:

- Open architecture
- Front-facing I/O for easiest serviceability
(The ISS3118 has two hard disk drives accessible from the rear.)
- Leading thermal management & power efficiency
- Ideal for multi-server deployments
- Best-in-breed software options
- Chassis mountable as half-depth back-to-back for double density, two-post front or center mount
- Auto-switching 100–240VAC or –48VDC (–36 to –73VDC) power supply options
- Three rack units (3RU) high

Configurable Options

- Supports up to two Intel®, Xeon®, or AMD Opteron™ processors
- Configurable number of hard drives:
 - Up to 6 removable hard drives (ISS3006 model)
 - Up to 9 removable hard drives (ISS3009 model)
 - Up to 12 removable hard drives (ISS3012 model)
 - Up to 18 removable hard drives (ISS3118 model)
- Supports SAS and SATA hard drives
- CD/CW-RW and DVD/DVD-RW
- Rackable Roamer™ remote management with LCD display option
- PCIe® expansion card slots available

Support Services

Among the support services included with the ISS3000 servers are the following:

- World-class customer support and services
- Availability of our professional services team to help with solutions outside traditional support packages
- Flexibility of server configurations to accommodate customer-specific operating system solutions
- Warranty options with full telephone and online support

Key Components of ISS3000 3U Servers

The configurations of the various ISS3000 servers are quite similar. Figure 1-1 shows the components of the front and rear panels of the ISS3012 server.




Figure 1-1 Front and Rear Panel Components

Table 1-1 describes the various components illustrated in Figure 1-1.

Table 1-1 Front Panel Components

Component	Description
Disk drives	These removable hard drives can be placed in drive carriers and can be hot-swappable. The drive carriers are connected to the system internally via a custom backplane that interfaces with a RAID storage card. Specifics of the RAID card used depend on the system specifications of your order.
DVD/CD drive	Depending on the configuration of your server, a DVD/DVD-RW or CD/CD-RW may be present. Also, in this area is space for an LCD display that is used with a Roamer remote management module if one was purchased with the server.
Front panel I/O	<p>The I/O interface to the motherboard. The configuration of the various I/O ports is dependent on the motherboard used in your configuration. For detailed I/O specifications, refer to the motherboard manual included with the documentation and accessory kit. Typical I/O includes the following:</p> <ul style="list-style-type: none"> – serial port (RS-232, 9600 baud) – USB ports – PS-2 connector for mouse and keyboard – VGA port – Ethernet ports (RJ-45, eth0/eth1, 10/100/1000Mb)
LCD	<p>(Optional)</p> <p>The LCD screen is used to display messages from the Roamer remote management solution. The Roamer module has the ability to display information relating to the system that can be customized according to your need.</p>

Table 1-1 Front Panel Components **(continued)**

Component	Description
LED indicator lights	<p>These LEDs provide you with critical information related to different parts of the system.</p> <p>Green (PWR): The LED farthest to the left indicates that power is applied to the server and the system is on. When the LED is off, the system is off, even if the AC power cord is plugged into the server.</p> <p>Blue (HDD): The LED in the middle indicates that the hard drive(s) are currently being accessed and activity is occurring.</p> <p>Red/Tri-color (ID2): The LED to the far right may have several functions depending on the configuration of your server. If your server was purchased with the Roamer remote management module, then this LED can be turned on or set to blink via Roamer so that a technician can easily identify a unit that may need service.</p>
PCI expansion slot	This slot is to accommodate add-on PCIe cards.
Power button	This single-push, push button is a power on/off button. When the system is in an off state, pushing the button once powers the system on. When the system is in an on state, pushing and releasing the button power the system off. This button does not have a reset function.
Power supply connector	<p>AC: Connect only 100-240 Volt AC, 5.5A maximum, 50-60 Hz. Use a standard AC connector to connect power to the server. An AC power cord is included in your documentation and accessory kit.</p> <p>DC: Connect only.</p>
Power switch	This power switch turns the server on and off. It controls the power supply that is inside the chassis.

Table 1-1 Front Panel Components **(continued)**

Component	Description
Rear fans	Keep the area behind the fans clear and unobstructed. This will ensure proper airflow to the components inside the server chassis.
Serial connector (RJ-45 management port)	This port is accessed with an RJ-45 connector. When the Roamer remote management module is installed, this port functions as an interface to the module. When the Roamer remote management module is not installed, the pinout for the serial connector is as shown in Table 1-2 .

[Table 1-2](#) describes the pinouts for the serial connector when Roamer is not installed.

Table 1-2 Pinouts for the Serial Connector J3

Pin	Name	Description
1	RTS	Request To Send (JP1 "A")
	CTS	Clear To Send (JP1 "B")
2	DSR	Data Set Ready
3	RxD	Receive Data
4	GND	Ground Return
5	GND	Ground Return
6	TxD	Transmit Data
7	DTR	Data Terminal Ready
8	CTS	Clear To Send (JP1 "B")
	RTS	Request To Send (JP1 "B")

Power Supply Options

SGI offers both DC and AC options for its servers. The make and model of the power supply that is present in your system can vary depending on what was ordered. SGI does have a standard power supply for both AC and DC configurations. In addition to single AC or DC power supplies, SGI offers a Redundant Power Supply (RPS) configuration option for some configurations.

AC

For AC powered servers, SGI typically uses a 250W or 450W power supply (or a redundant supply for RPS systems). The power supply provides a small footprint and can handle a peak of 450W DC output. The input is 100-240VAC, 43-63Hz.

DC

For DC powered servers, SGI uses a 250W or 450W power supply. The power supply provides a small footprint and can handle a peak of 450W DC output. The DC input threshold is +/-38 – +/-75VDC.

Pre-Installation Guidelines

To ensure safe and efficient operation of your server, proper planning is essential. Factors such as the proper location of the server in its rack, adequate power to the components in the rack, and the appropriate operating environment for the rack. The guidelines described in this chapter are categorized as follows:

- “Server and Rack Positioning” on page 9
- “Electrical Power” on page 10
- “Operating Environment” on page 10
- “Rack Stability” on page 11
- “Cables” on page 11
- “Security” on page 11

Server and Rack Positioning

Follow these guidelines to ensure that the server is safely and appropriately positioned for efficient operation and service.

- When determining the location for the server, position it so that there is adequate space for airflow and servicing.
- When positioning the server in a rack, be sure to allow adequate space for airflow and servicing from the front and the rear.
- Clearance in the front of the server should be a minimum of 36” to allow the chassis to easily slide in and out of the rack.
- Clearance in the rear should be approximately 30” to allow for sufficient space for airflow and ease of servicing.
- Do not cover the front or rear of the server or any of the openings on the server front and rear panels.

- If you are installing multiple servers or other components in the rack, place the components and servers in positions that they can be easily opened and serviced. For example, the heaviest components—such as a uninterruptible power supply—are placed at the bottom of the rack; servers are often in the middle of the rack.

Electrical Power

If installing this server in a rack that contains multiple components, be sure that the power connections to the rack are sufficient for the combined power requirements of all the components. Ensure adequate and safe power by following these guidelines:

- Check documentation of all the components in the rack to obtain the total power requirement that the rack will use. Cross-reference this need with information on the total available power to see if the power supply for the rack is sufficient for the planned components.
- When planning electrical power arrangements, make sure you have more power available than specified for all components. Also, make sure the power load is distributed evenly among circuits to the location of the rack.
- Make sure power connections for the server and all other components are grounded.
- See the [Appendix A, “Technical Specifications”](#) for specific data on the electrical power requirements for this server.

Operating Environment

The operating environment where the server is installed must meet the following temperature and ventilation requirements:

- Verify the temperature range of the datacenter or room where the server is installed. Ensure that it is within the limits established for the server and all other components. This server is designed to function within the temperature range of 10-35°C / 50-90°F.
- Adequate ventilation must be present to maintain the specified temperature range. This is particularly important for a server or rack enclosed in a cabinet.

Rack Stability

The rack must be strong enough to hold the installed components.

- Ensure that the leveling jacks at the bottom of the rack are fully extended and are secure.
- If using a two-post (telco) rack, be sure that the rack is securely fastened to the building at the top and the bottom.
- In single rack installations, stabilizers should be attached to the rack.
- When working with components in the rack, never slide out more than one unit at a time.

Cables

For optimal efficiency in server operation and ease of maintenance, handling cables properly is required.

- Arrange all component cables so they do not interfere with access to the rack.
- To ensure full signal strength for Ethernet, serial, and other connections, make sure the cables do not exceed established length limits.

Security

The following are security considerations:

- If using a server cabinet that is not in a secure room, be sure the cabinet is adequately locked and access is limited to authorized staff.
- Control access to the server environment with a plan for distributing and controlling keys and access information.
- Store a copy of server access information at a safe location away from the server site.

Installation Guidelines

This chapter describes how to install your server into the rack.

Warning: Your server is designed for rack mounting. It is not designed for use as a desktop system. Do not place a monitor on the server or place anything on top of the enclosure. Additional weight on top of the server could cause damage to internal components.

Your server can be installed in these types of racks:

- Open four-post rack, 19" wide, 24-26" deep or 29-36" deep
- Cabinet with four-post rack inside, 19" wide, and 29-36" deep
- Two-post Telco rack, 19" wide.

The server is three rack units (3RUs) tall.

Getting Ready to Install

You can install the server into any four-post rack or two-post (Telco) rack. Preparations for installing are the same regardless of what type of rack you using for the server. The ISS3000 servers do not require convention rack or server rails to be installed. The servers are designed to slide easily between the mounting rails of a rack.

Important: Check the documentation for your rack for any special instructions or requirements.

Before beginning work with the server and rack, make the following preparations:

- If possible, arrange to work with another person for safety.
- Assemble the tools, brackets, and connectors you will need for installation.
- Note the screw sizes that you will need to use with your rack. Typically, racks use one of the following two:
 - English size 10-32 or 12-24
 - Metric size M5 for racks with metric holes

Check the documentation for your rack and if screws are provided, use those.

- Clear a table, cart, or other flat surface near the rack. You will need a space to temporarily place the server during installation.

Determine Server Positioning in the Rack

Review the guidelines for positioning the server in the rack explained earlier. Then follow these steps to determine its positioning.

1. Determine the exact position where you want to attach the server and mark it on the side of the rack.

Some racks have marks already to aid in positioning a server. If your rack does not have such aids, measure and count holes to determine desired positioning. The distance between holes on a rack varies on racks made by different manufacturers.
2. Once positioning is correctly determined in the rack you are ready to connect the server to the rack. The server can be connected to standard Telco two- or four-post racks.

Installing into a Rack or Cabinet

A rack can be either open or enclosed in a cabinet. Follow the same instructions to attach the server to either type of rack. For a cabinet, it helps to remove the door before installing the server. The server does not require conventional rack or server rails to be installed. The server is designed to slide easily between the mounting rails of the rack. Once the exact position of the server is determined on the rack, installation of the server is as follows:

1. Position the server so that the two “ears” on the side of the server are positioned equally and level in the rack.
2. Install four screws through the mounting ears and into the rails of the rack.

Typical screw sizes are 10-32 and 12-24. Check the accessory kit for spare screws.

If your rack does not have threaded holes, you will be required to provide threaded inserts to snap into the rails.

Installing into a Telco Rack

Installing the server into a Telco rack is exactly the same procedure as just described. The only difference in the installation procedure will be the positioning of the server mounting ears to the rack. The ears should be moved to the middle of the server to accommodate installation into a Telco rack.

Cabling

Once the server is secured in the rack, you can connect the ethernet cables and power cord directly to its front and back panels, respectively.

Gather the cables you will connect to the server. Make certain each cable has the proper connector and that it is designed for use in a high-capacity server.

Label each cable so that you can locate a specific cable quickly.

Ethernet ports on the server are numbered 0 and 1. 0 is the primary ethernet port, and 1 is the secondary.

Warning: Because the server is being installed into a rack with other equipment, be certain that the power used with the rack is designed to carry the electrical load of all of the devices in the rack. Refer to [Appendix A, “Technical Specifications”](#) for details on the power requirements for this server.

Installing Hard Drives

You must install your hard drives into the chassis. Install them only after you have securely installed the server in a rack or cabinet. The order in which you populate the hard drive slots is significant. This chapter describes the drive insertion order and the procedure for inserting hard drives.

Drive Insertion Order

Table 4-1 shows the drive insertion order for model ISS3006.

Table 4-1 Drive Insertion Order—ISS3006

0	4
1	5
2	6

Table 4-2 shows the drive insertion order for model ISS3009.

Table 4-2 Drive Insertion Order—ISS3009

0	4	8
1	5	9
2	6	10

The chassis for the ISS3006 and ISS3009 use a 1x3 backplane. Drive slots 3, 7, and 11 (where applicable) are skipped.

Table 4-3, Table 4-4, and Table 4-5 show the drive insertion order for several configurations of the ISS3012.

Table 4-3 Drive Insertion Order—ISS3012 (6Gb Backplane, SAS/SATA with One 12-port SAS/SATACard)

0	1	2	3
4	5	6	7
8	9	10	11

Table 4-4 Drive Insertion Order—ISS3012 (6Gb Backplane, SAS/SATA with Two 8-port SAS/SATA Cards)

0-0	0-1	0-2	0-3
0-4	0-5	0-6	0-7
1-0	1-1	1-2	1-3

Table 4-5 Drive Insertion Order—ISS3012 (3Gb Backplane)

3	2	1	0
7	6	5	4
11	10	9	8

Table 4-6 and Table 4-7 show the disk insertion order for configurations of the ISS3118 with 3Gb or 6Gb backplanes.

Table 4-6 Drive Insertion Order—ISS3118 (3Gb Backplane)

3	2	1	0
7	6	5	4
11	10	9	8
15	14	13	12

Table 4-7 Drive Insertion Order—ISS31118 (6Gb Backplane)

0	1	2	3
4	5	6	7
8	9	10	11
12	13	14	15

Installing a Hard Drive into the Hard Drive Carrier

To install a hard drive into the hard drive carrier, perform the following steps after removing the hard drive from the chassis:

1. Place the hard drive carrier on a flat, stable surface such as a desk, table, or work bench.
2. Slide the hard drive into the carrier with the printed circuit board side facing down (see [Figure 4-1](#)).
3. Carefully align the mounting holes in the hard drive and the carrier. Make sure the bottom of the hard drive and bottom of the hard drive carrier are flush.
4. Secure the hard drive using all four screws (see [Figure 4-1](#)).
5. Replace the drive carrier into the chassis. Make sure to close the drive carrier using the drive carrier handle.




Figure 4-1 Inserting Hard Drive into Carrier

Using the Server

This chapter describes the following:

- “Powering the Server On” on page 21
- “Powering the Server Off” on page 22
- “Working with Roamer Remote Management” on page 22

Powering the Server On

After the server chassis is connected to AC power, the server can be turned on in any of the following ways:

- Press the power-control button on the front of the server.
While the server is powering up, the power LED on the front of the chassis is lit.
- If a power failure occurs, the server can be started automatically when power is restored.
This is only available if the Rackable Roamer remote management module was ordered for the server.
- If the operating system you installed on the server supports the Wake on LAN feature, the Wake on LAN feature can turn on the server.

The Wake on LAN feature only works if it was not disabled through the system BIOS, and if the server was previously turned on and the operating system was shut down properly.

Powering the Server Off

When the server is turned off, it remains connected to AC power through the server chassis. The server can respond to requests to turn on remotely or locally through Roamer. To remove all power from the server, power must be disconnected from the server chassis.

Shut down your operating system before turning off the server. Refer to the operating system documentation in reference to proper operating system shutdown. Improper shutdown of a server will not allow the server to be restarted using Wake on LAN.

The server can be turned off in the following ways:

- Press the power-control button on the front of the server.
This begins the shutdown of the operating system, if this feature is supported by your system.
- If the operating system stops functioning, you can press and hold the power-control button to turn off the server.
- The Roamer remote management module (if one is present in your server) can be used to turn off the server.

Working with Roamer Remote Management

If the server you purchased has the optional Roamer remote management hardware present, then the server can be managed remotely using this option. The Roamer option provides standard motherboards with enhanced remote control and serial redirection capabilities. The Roamer is accessed locally using an RJ-45 serial connection (shown in [Figure 1-1](#)).

With the Roamer option comes a LCD display. This display can be directly programmed within the OS. The LCD display can also be used to display server status, and other information as needed.

RAID Configuration

This chapter describes how to configure your server to use hardware RAID and to use software RAID. For an overview of RAID, refer to [Appendix C, “RAID Basics”](#).

Configuring Hardware RAID

The hardware RAID setup procedures vary depending on the manufacturer of the RAID controller card(s). SGI typically ships the ISS3000 configurations with 3ware™ or LSI® RAID controller cards. This chapter describes the setup of 3Ware and LSI RAID controller card configurations. If your server does not have a 3ware or LSI card, refer to the documentation that accompanied your RAID controller card for instructions.

LSI Hardware RAID (SAS or SATA)

1. During the boot-up sequence, access the LSI controller configuration screen by pressing `Alt-3`.
The system displays instructions to guide you.
2. Follow the instructions to create the RAID partition of choice on each RAID controller.
3. (Optional) Once the server is fully operational, install the LSI management software (available at <http://www.lsi.com>).

The LSI storage management software allows administrators to perform RAID diagnostics and provide automatic notification of drive failures.

3ware-Based Hardware RAID (SATA)

1. During the boot-up sequence, access the 3ware controller configuration screen by pressing **Alt-3**.
The system displays instructions to guide you.
2. Follow the instructions to create the RAID partition of choice on each RAID controller.
3. (Optional) Once the server is fully operational, install the 3ware 3DM™ management software (available at <http://www.3ware.com>).

The 3ware storage management software allows administrators to perform RAID diagnostics and performs automatic notification of drive failures.

Configuring Software RAID

Software RAID arrays are configured through the server operating system. To configure a RAID array using software RAID, do the following:

1. Follow the hardware RAID setup instructions above for the RAID card that is installed on the server.
2. Enter the RAID card BIOS screen and erase any RAID arrays that exist.
This should leave all drives in a JBOD state.
3. Repeat for every RAID card in the server.
4. After you remove all arrays, create new RAID arrays within the operating system.

Software RAID array management is performed using operating system level utilities.

Software Installation and Tuning

The ISS3000 series of servers can be configured as one of the following:

- Linux® server
- Windows® server

This chapter describes installation and tuning guidelines for each of these.

Linux Server Configuration

Any variant of Linux can be installed to create a Linux NAS server on the ISS3000 series. SuSE® or Red Hat® versions are recommended due to overall stability. The following sections will assist with the installation and tuning of Linux systems. This section assumes a moderate level of Linux system administration skills.

Operating System Installation

The system may include separate boot drives that are isolated from the RAID array(s). Install the operating system on the drives that are separate from the drives in the array. When using software RAID on the data drives, isolate the boot drive from the array.

If there are not any separate boot drives, then install a boot partition (for `/boot`) on a small part of one of the RAID volumes (100-200MB is sufficient). Leave the remaining capacity for data. To create partitions, use a partitioning tool such as `parted` or `fdisk`.

Filesystem Selection

There are several production-quality filesystems available for Linux. Choose the filesystem that suits the needs of the server application. Table 7-1 lists some examples of common and useful filesystems.

Table 7-1 Popular Linux Filesystems

Filesystem	Description
Ext4	The widely supported, default Linux filesystem. Supports filesystem journaling. Average performance with small- and large-file streaming.
XFS®	Supports filesystem journaling. Average small-file transaction performance and high-performance large-file streaming.

The Red Hat distribution installer does not support XFS by default. Newer releases of RHEL and Fedora™ support XFS if filesystem support is enabled in the installer. Keep `/boot` as an Ext4 filesystem when installing Red Hat Linux.

Data Drive Configuration (Hardware RAID)

When using hardware RAID, one SCSI device will be visible to the operating system for each hardware RAID partition that is created during RAID card BIOS configuration. In Linux, these devices will be visible as `/dev/sdx` (where *x* is a letter). Identify the correct device name that is to be configured.

Create a filesystem using the method applicable for the filesystem that was selected earlier. For example, to create an XFS filesystem on `/dev/sda`, enter the following command:

```
mkfs.xfs /dev/sda
```

After the filesystems have been created, create mount points. Add the mount points to the `/etc/vfstab` file so that they are mounted on bootup.

Mount the mount points to mount the filesystems. Command `df` will show the available capacity.

Exporting Filesystems

To export the filesystems via NFS, entries need to be added to the `/etc/exports` file. Add `no_subtree_check` to increase NAS performance. Add `async` to the export options to cache disk writes in system memory prior to being written to disk.

Windows clients will need to set up Samba to export filesystems via CIFS protocol. For more information on Samba, visit the Samba website <http://samba.org>.

Tuning Linux NFS Server Performance

Increase the size of the data read-aheads to increase the large-file streaming performance of the server. To accomplish this, the following lines need to be added to `/etc/sysctl.conf`:

```
vm.max-readahead = 256
vm.min-readahead = 128
```

Enter `sysctl -e -p /etc/sysctl.conf` to apply the changes.

Tuning Linux Client Performance

Client-side optimization is required to achieve the most performance possible with a Linux server configuration. Varying the mount options is the most effective way to optimize the Linux client.

Using NFS v3 over UDP with the `RSIZE` and `WSIZE` mount options set to 32768 will optimize the client for streaming environments. These settings will provide the highest throughput rates provided the network does not show excessive packet retransmits. For older networking equipment, try smaller block sizes (4K or 8K) for more optimized results.

To improve large-file single-stream performance, modify the `/etc/sysctl.conf` file with the following lines:

```
net.core.rmem_max = 262143
net.core.wmem_max = 262143
net.core.rmem_default = 262143
net.core.wmem_default = 262143
```

Enter `sysctl -e -p /etc/sysctl.conf` to apply the changes.

Using Smartmontools Software to Improve Availability

Note: This section applies if the system is using a 3ware SATA RAID controller card in either hardware or software RAID modes. The Smartmontools package (available at the website <http://smartmontools.sourceforge.net>) should be used in conjunction with the RAID management package that accompanied the RAID controller card. Refer to [Chapter 6, “RAID Configuration.”](#)

The Smartmontools utility provides a mechanism to run long integrity tests on individual hard drives through a 3ware controller card. It is useful to detect SATA drives that are likely to fail in the future but have not resulted in a failure yet. This allows for preventative drive replacement prior to a significant RAID failure. Running Smartmontools tests on drives periodically will allow for improved server and data availability. You can schedule the tests to run without manual intervention.

Verify that the test has completed by entering the `smartctl` command with the appropriate port number and device specified, as shown in the following example:

```
smartctl -a -d3ware,4 /dev/sda
```

After test completion, you can analyze the test data. If any tests return errors, submit the test data to SGI Support to determine the appropriate support action. The drive may need to be replaced to prevent future RAID array failure.

Windows Server Configuration

Windows 2003 or 2008 Server operating systems can easily be installed on the ISS3000 series of servers to create a standard Windows file or application server. Use the following guidelines:

- Install the server operating system onto drives that are separate from the RAID array if possible.
- Ensure that the RAID card drivers are compatible with the RAID card firmware.
- A reboot is necessary after operating system and driver installation.

Use the standard Windows administration tools to create new partitions and drive letters for the RAID volumes on the array.

Technical Specifications

Table A-1 lists the physical and environmental specifications for the ISS3000 series of 3U servers.

Table A-1 Technical Specifications

Attribute	ISS3006	ISS3009	ISS3012	ISS3118
Processor options	Intel Xeon (2) or AMD Opteron (2)	Intel Xeon (2) or AMD Opteron (2)	Intel Xeon (2) or AMD Opteron (2)	Intel Xeon (2) or AMD Opteron (2)
Drive options	Up to 6 hot-swap SAS or SATA	Up to 9 hot-swap SAS or SATA	Up to 12 hot-swap SAS or SATA	Up to 18 hot-swap SAS or SATA
Expansion cards	Up to 5 PCIe slots	Up to 5 PCIe slots	One PCIe slot (dedicated to RAID)	Up to 7 PCIe slots
Controller support	SAS or SATA	SAS or SATA	SAS or SATA	SAS or SATA
RAID levels supported	0,1,5,6,10,50	0,1,5,6,10,50	0,1,5,6,10,50	0,1,5,6,10,50
Chassis mount	Half-depth back-to-back or two-post front or center	Half-depth back-to-back or two-post front or center	Half-depth back-to-back or two-post front or center	Full-depth with four-post rack
Power supply (typical power supply configurations listed)	1+1 redundant autoswitching (RPS) Output: 250W/450W Input: 100-240VAC (standard), -48VDC (optional)	Output: 250W/450W Input: 100-240VAC (standard), -48VDC (optional)	Output: 250W/450W Input: 100-240VAC (standard), -48VDC (optional)	1+1 redundant autoswitching (RPS) Output: 250W/450W Input: 100-240VAC (standard), -48VDC (optional)

Table A-1 Technical Specifications **(continued)**

Attribute	ISS3006	ISS3009	ISS3012	ISS3118
Dimensions	3U chassis 5.25"x17.6"x15.5" (rack ears mounted)	3U chassis 5.25"x17.6"x15.5" (rack ears mounted)	3U chassis 5.25"x17.6"x15.5" (rack ears mounted)	3U chassis 5.25"x17.3"x27.3"
Gross weight	39 lbs.	45 lbs.	51 lbs.	84 lbs
Operating temperature	10-35° C/50-90° F	10-35° C/50-90° F	10-35° C/50-90° F	10-35° C/50-90° F
Operating humidity	8-80%	8-80%	8-80%	8-80%
Max. thermal sustained output	1600BTU/hr	1800BTU/hr	2000BTU/hr	2000BTU/h
Compliance	US (FCC, CFR47 Part 15, Class B) Europe (CE, EN55022 & EN55024) Australia (C-TICK) Japan (VCCI)	US (FCC, CFR47 Part 15, Class B) Europe (CE, EN55022 & EN55024) Australia (C-TICK) Japan (VCCI)	US (FCC, CFR47 Part 15, Class B) Europe (CE, EN55022 & EN55024) Australia (C-TICK) Japan (VCCI)	US (FCC, CFR47 Part 15, Class B) Europe (CE, EN55022 & EN55024) Australia (C-TICK) Japan (VCCI)

Safety Guidelines

The safety guidelines categorized as follows:

- “Electrical Safety” on page 31
- “General Safety” on page 32
- “Handling Your Server” on page 33

Electrical Safety

Observe the following basic electrical safety precautions to protect yourself from harm and the server from damage:

- Be aware of the locations of the power on/off switch on the chassis as well as the room's emergency power-off switch, disconnection switch, or electrical outlet. If an electrical accident occurs, you can then quickly remove power from the system.
- Do not work alone when working with high voltage components.
- Power should always be disconnected from the system when removing or installing main system components—such as the motherboard, memory modules, and drives. When disconnecting power, you should first power down the system with the operating system first and then unplug the power cords of all the power supply units in the system.
- When working around exposed electrical circuits, another person who is familiar with the power-off controls should be nearby to switch off the power if necessary.
- Use only one hand when working with powered-on electrical equipment. This is to avoid making a complete circuit, which will cause electrical shock. Use extreme caution when using metal tools, which can easily damage any electrical components or circuit boards they contact.
- Do not use mats designed to decrease static electrical discharge as protection from electrical shock. Instead, use rubber mats that have been specifically designed as electrical insulators.

- The power supply power cords must include a grounding plug and must be plugged into grounded electrical outlets.
- Verify that the electrical outlets and power strips used by the server have reliable earthing.
- Verify that the server is plugged into an electrical outlet with appropriate circuit overloading protection.

General Safety

The following are general safety guidelines:

- Keep the area around the server clean and free of clutter.
- The server weighs approximately 20 lbs when fully loaded. When lifting the system, two people at either end should lift slowly with their feet spread out to distribute the weight. Always keep your back straight and lift with your legs.
- Place the chassis top cover and any system components that have been removed away from the system or on a table so that you will not step on them.
- While working on the system, do not wear loose clothing such as neckties and unbuttoned shirt sleeves, which can come into contact with electrical circuits or be pulled into a cooling fan.
- Remove any jewelry or metal objects from your body. They are excellent metal conductors that can create short circuits and harm you if they come into contact with printed circuit boards or areas where power is present.
- After accessing the inside of the system, close the system back up and secure it to the rack unit with the retention screws after ensuring that all connections have been made.
- If you replace the motherboard battery, ensure that you install it properly. This battery must be replaced only with the same or an equivalent type recommended by the manufacturer. Dispose of used batteries according to the manufacturer's instructions.

Caution: There is a danger of explosion if the onboard battery is installed upside down, which will reverse its polarities

- To ensure proper cooling when the server is operating, take care to ensure that all chassis covers are in place, that all drive bays either have a drive or a blank installed, and that the room ambient temperature is within the specified operating temperature range.

Handling Your Server

When handling the server, use the following guidelines:

- When the server is removed from the rack, set it on a sturdy and flat surface. Do not set any heavy objects on top of the server. Any additional weight could possibly damage essential internal components.
- When connecting or disconnecting a cable, always hold the cable by its connector.
- Never force a connector into a port. If the connector and port do not join with ease, check the matching of the connector with the port. Correct positioning when joining a connector to a port is also essential.
- Protect the server and its components from direct sunlight and rain or other moisture.
- Keep all ventilation openings clear and unobstructed. Without proper air circulation, components can overheat, causing damage or unreliable operation.

RAID Basics

To assist you in building RAID arrays, this chapter gives a general overview of RAID and describes various levels of RAID along with JBOD.

RAID Implementations

RAID can be implemented in either dedicated hardware or software running on standard hardware. Occasionally, you might want hybrid RAID systems—combining software- and hardware-based solutions.

With software implementations of RAID, the operating system manages the disks through the normal drive controller (SAS/SATA). Software RAID can be faster than hardware-based RAID at the cost of using CPU power. One exception to this is when a dedicated hardware RAID system has a write cache dedicated to managing the RAID.

A hardware implementation of RAID requires a RAID controller. This may be a PCIe expansion card or may be a card onboard the system motherboard in some systems. Disks may be of any type: SAS or SATA. The controller links to the system with a high-speed SCSI, or iSCSI connection. The controller handles the management of disks, including parity calculations (error checking) which is suited for larger arrays of disks. Hardware RAID tends to have better performance and makes operating system support easier. Hardware RAID typically supports *hot-swapping* of drives as well. Both hardware and software versions support the use of a hot spare, a preinstalled drive which is used to immediately replace a failed drive.

Some software RAID systems allow building arrays from partitions instead of whole disks.

RAID 0

RAID 0 (also known as a stripe set) splits data evenly across two or more disks with no parity information for redundancy. RAID 0 is not redundant. RAID 0 is used to increase performance, and can be used as a way to create a small number of large virtual disks out of a large number of small physical ones.

RAID 0 can be created with disks of differing sizes, but the storage that is added to the array by each disk is limited to the size of the smallest disk.

Example: A 120GB drive is striped with a 100GB drive. The size of the array will be 200GB (100+100).

RAID 0 is useful for read-only NFS servers where mounting many disks is time-consuming and redundancy is irrelevant.

JBOD (Just a Bunch of Disks)

JBOD is a popular method for combining multiple physical disk drives into a single virtual one. Disks are merely concatenated, beginning to end. JBOD is commonly used to turn several odd-sized drives into one useful drive.

One advantage JBOD has over RAID 0 is in the case of drive failure. When a single drive of a RAID 0 array fails, usually it will result in the loss of all data in the array. When a drive failure occurs in a JBOD array, only the data on the affected drive is lost.

RAID 1

RAID 1 creates an exact copy (or mirror) of all the data on two or more disks. This is useful where redundancy is more important than using all of the disk storage capacity. The array in a RAID 1 can only be as big as the smallest member disk. An ideal RAID 1 set contains two disks. Since each member can be addressed independently if the other fails, reliability is a linear multiple of the number of disks in the array. RAID 1 benefits most when an independent disk controller is used for each disk.

Seek time with a RAID 1 remains low because each drive can be accessed independently, and the requested sectors can be split evenly between disks. When writing, the RAID 1 array acts like a single disk because all writes must be written to all disks.

RAID 1 has administrative advantages: it allows 24/7 environments to do backups on the array. Mirroring is useful when critical data needs to be stored and backed up.

RAID 5

A RAID 5 array uses block-level striping with parity data distributed across all the member disks. RAID 5 is the most popular RAID level and is used in both hardware and software RAID. You can create a RAID 5 array using disks of different sizes, but the storage space added to the array by each disk will be limited to the size of the smallest disk in the array.

Example: If a 120GB disk is used to build a RAID 5 array with two 100GB disks, each disk will donate 100GB to the array. 200GB will be used for storage while 100GB will be used for parity information.

Every time a data block is written on a disk, a parity block is generated within the same stripe. If another block is written on that same stripe, the parity block is recalculated and rewritten. The disk used for the parity block is staggered from one stripe to the next.

The parity blocks are not read on data reads. The parity blocks are used when a read of a data sector results in a CRC (cyclic redundancy check) error. In a case like this, the sector in the same relative position within each of the remaining data blocks in the stripe and within the parity block are used to reconstruct the sector with the error. Also, if a disk were to fail in a RAID 5 array, the parity blocks from the surviving disks are combined mathematically with the data blocks from the surviving disks to reconstruct the data on the failed drive.

The maximum number of drives in a RAID 5 array are theoretically unlimited. However, with a greater number of disks in an array, there is a greater chance of drives failing in succession.

RAID 6 (Dual Parity)

Striped set with dual distributed parity. Provides fault tolerance from two drive failures; the array continues to operate with up to two failed drives. With dual parity, it gives time to rebuild the array without the data being volatile while the failed drive is being recovered.

