IBM @server xSeries 450 Planning and Installation Guide

- Describes the technical details of the new 64-bit server
- Covers Windows Server 2003 and SuSE Linux Enterprise Server
- Helps you prepare for and perform an installation

David Watts
Gerry McGettigan
Michael L. Nelson
Lubos Nikolini
Jose Rodriguez Ruibal

Redbooks

ibm.com/redbooks

International Technical Support Organization

IBM @server xSeries 450 Planning and Installation Guide

July 2003

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

First Edition (July 2003)

This edition applies to the IBM @server xSeries 450, machine type 8688.

This document was updated on July 1, 2003.

Contents

Notices vii
Trademarks viii

Preface ix
The team that wrote this redbook ix
Become a published author xii
Comments welcome xii

Chapter 1. Technical description 1
1.1 The x450 product line 2
1.2 IBM XA-64 chipset 3
1.3 Extensible Firmware Interface 5
1.3.1 GUID Partition Table disk 7
1.3.2 EFI System Partition 8
1.3.3 EFI and legacy-free concept 10
1.4 Intel Itanium 2 processor 10
1.5 System assembly 13
1.5.1 Memory-board assembly 14
1.5.2 Processor-board assembly 15
1.6 IBM XceL4 Server Accelerator Cache 17
1.7 System memory 17
1.8 PCI subsystem 21
1.9 Redundancy 23
1.10 Light path diagnostics 25
1.11 Remote Supervisor Adapter 26
1.12 Operating system support 27

Chapter 2. Positioning 29
2.1 x450 application solutions 30
2.1.1 Database applications 30
2.1.2 Business logic 32
2.1.3 e-business and security transactions 33
2.1.4 In-house developed compute-intensive applications 34
2.1.5 Science and technology industries 35
2.2 Why choose x450? 35

Chapter 3. Planning 39
3.1 System hardware 40
3.1.1 Processors 40

3.1.2 Memory 41
3.1.3 PCI slot configuration 44
3.1.4 Broadcom Gigabit Ethernet controller 47
3.2 Cabling and connectivity 48
3.2.1 Remote Supervisor Adapter connectivity 50
3.2.2 Remote Expansion Enclosure 51
3.2.3 Serial connections 55
3.3 Storage considerations 56
3.3.1 xSeries storage solutions 56
3.3.2 Tape backup 61
3.4 Rack installation 62
3.5 Power considerations 63
3.6 Solution Assurance Review 64

Chapter 4. Installation 67
4.1 Using the Extensible Firmware Interface 68
4.1.1 The EFI shell 70
4.1.2 Flash update 80
4.1.3 Configuration and Setup utility 83
4.1.4 Diagnostic utility 84
4.1.5 Boot maintenance menu 85
4.2 Installing Windows Server 2003 86
4.2.1 Overall process 87
4.2.2 Microsoft Reserved Partition 88
4.2.3 Windows installation 89
4.2.4 Pre-installation phase 89
4.2.5 Starting the installation 90
4.2.6 Text-mode setup 91
4.2.7 GUI setup 93
4.2.8 Post-setup phase 93
4.3 Installing SuSE Linux Enterprise Server 98
4.3.1 Background 99
4.3.2 Linux IA-64 kernel overview 100
4.3.3 Choosing a Linux distribution 101
4.3.4 Installing SuSE Linux Enterprise Server 103
4.3.5 Linux boot process 107
4.3.6 Information about the installed system 108
4.3.7 Using the serial port for the Linux console 112
4.3.8 RXE-100 Expansion Enclosure 113
4.3.9 Upgrading drivers 114

Chapter 5. Management 117
5.1 The Remote Supervisor Adapter 118

5.1.1 The Remote Supervisor Adapter . . . 119
5.1.2 Connecting via a Web browser . . . 119
5.1.3 Configuring a static IP address . . . 121
5.1.4 Connecting via the ASM interconnect . . . 124
5.1.5 Installing the device driver . . . 125
5.1.6 Configuring the remote control password . . . 125
5.2 Management using the Remote Supervisor Adapter . . . 126
5.2.1 Configuring which alerts to monitor . . . 126
5.2.2 Configuring SNMP . . . 127
5.2.3 Sending alerts directly to IBM Director . . . 128
5.2.4 Creating a test event action plan in IBM Director . . . 130

Related publications . . . 135
IBM Redbooks . . . 135
Referenced Web sites . . . 135
How to get IBM Redbooks . . . 137
IBM Redbooks collections . . . 137

Index . . . 139

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:
This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

Redbooks (logo)™ 	@server ™	Redbooks™
iSeries™	@server™	ServerProven®
pSeries™	ESCON®	ServeRAID™
xSeries™	FlashCopy®	Summit®
zSeries™	FICON™	ThinkPad®
Chipkill™	IBM®	Tivoli®
DB2 Universal Database™	Lotus®	TotalStorage™
DB2®	Netfinity®	Wake on LAN®
Enterprise Storage Server™	Predictive Failure Analysis®	X-Architecture™
eServer™	PS/2®	

The following terms are trademarks of other companies:

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

Other company, product, and service names may be trademarks or service marks of others.

Preface

The IBM @server xSeries™ 450 is IBM®'s new 64-bit Itanium Processor Family (IPF) Architecture server and is the first implementation of the 64-bit IBM XA-64 chipset, as part of the Enterprise X-Architecture™ strategy. This IBM Redbook is a comprehensive resource on the technical aspects of the server, and is divided into five key subject areas:

- Chapter 1, "Technical description" introduces the server and its subsystems and describes the key features and how they work. This includes the new Extensible Firmware Interface, which provides a powerful replacement to the BIOS facility found on the IA-32 platform.
- Chapter 2, "Positioning" examines the types of applications that would be used on a server such as the x450.
- Chapter 3, "Planning" describes the considerations when planning to purchase and planning to install the x450. It covers such topics as configuration, operating system specifics, scalability, and physical site planning.
- Chapter 4, "Installation" covers the process of installing Windows Server 2003, Enterprise Edition and SuSE Linux Enterprise Server on the x450.
- Chapter 5, "Management" describes how to use the Remote Supervisor Adapter to send alerts to an IBM Director management environment.

The team that wrote this redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization, Raleigh Center.

David Watts is a Consulting IT Specialist at the International Technical Support Organization in Raleigh. He manages residencies and produces IBM Redbooks™ on hardware and software topics related to IBM xSeries systems and associated client platforms. He has authored over 20 Redbooks; his most recent books include the *IBM @server xSeries 440 Planning and Installation Guide*. He has a Bachelor of Engineering degree from the University of Queensland (Australia) and has worked for IBM for over 14 years. He is an IBM @server Certified Specialist for xSeries and an IBM Certified IT Specialist.

Gerry McGettigan is an Advanced Technical Support engineer in EMEA working in Greenock, Scotland. Before joining ATS in the summer of 2002, he provided pre-sale and post-sale technical training on high-end xSeries systems and solutions as part of the xSeries Academy. He holds a Bachelor of Science in Multimedia Technology from Glasgow Caledonian University and has been with IBM for over five years. His area of expertise is xSeries systems and hardware.

Michael L. Nelson is a Performance Consultant in IBM @server Solutions Engineering working in Raleigh. He was previously an architect for many of the xSeries Web-based sizing tools. He is a Microsoft Certified Systems Engineer on both Windows NT 4.0 and 2000, a Lotus® Certified Professional, a Cisco Certified Network Associate, a Citrix Certified Administrator, and an IBM @server Certified Expert for xSeries. He also holds a Bachelor of Arts degree from Whittier College. He has five years of experience with Intel-based server hardware, and has worked for IBM for six years. His areas of expertise include xSeries hardware, Microsoft Windows, and Citrix.

Lubos Nikolini is a Systems Engineer working for HT Computers a.s. in Bratislava, Slovak Republic. He has six years of experience with Microsoft operating systems, Intel-based hardware (including IBM xSeries and Netfinity®), networking, and storage technologies. He is a Microsoft MCSE, Citrix CCEA and a Compaq Proliant ASE. His areas of expertise include Active Directory deployments, thin-client technology, and groupware solutions.

Jose Rodriguez Ruibal is a Linux Specialist in the Advanced Technical Support (ATS) Products and Solutions Support Center (PSSC) for EMEA in Montpellier, France. Prior to joining the ATS team, he worked for Red Hat in Spain and for the INRIA (National French Institute for Research in Computing and Robotics) in Linux-related activities. He has four years of experience in Linux and UNIX-like network operating systems. He holds a Bachelor's degree with Honors in Computer Science from the Universidad Antonio de Nebrija in Madrid, Spain, and he is also Red Hat Certified Engineer. His areas of expertise include xSeries hardware and storage, and Linux. He has also authored a white paper about Oracle 9i RAC installation under Linux on xSeries x360 and x440 machines.

The redbook team (l-r): David, Jose, Gerry, Mike, Lubos

Thanks to the following IBM employees for their contributions to this project:

Henry Artner, Service Education Curriculum Manager, Raleigh
Charlotte Brooks, ITSO Project Leader for Tivoli® Storage, San Jose
Pat Byers, Program Director, Linux xSeries Alliances & Marketing
Alex Candelaria, IBM Center for Microsoft Technologies, Seattle
Zeynep Dayar, xSeries Systems Management, Raleigh
Peter Figueroa, Technical Project Manager, Software Development, Raleigh
Amy Freeman, Manager Brand Communications, Linux on xSeries, Raleigh
Maynard Ferguson, x450 Product Development Team Lead, Raleigh
Mauro Gatti, EMEA xSeries Solution Architects, Italy
Barney Hallman, Server Development, Austin
Jim Hanna, xSeries Server Development, Austin
Roger Hellman, xSeries Global Product Manager, Raleigh
William L. Jones, IBM xSeries Preloads, Raleigh
Michael S Lee, IBM Center for Microsoft Technologies, Seattle
Grace Lennil, IBM Center for Microsoft Technologies, Seattle
Cecil Lockett, Senior Engineer, Engineering Software, Raleigh
John McAbel, Datacenter/Cluster Product Manager, Beaverton
Chris McDermott, Linux Technology Center, IA-64 Enablement, Beaverton
Willie Nathan, Systems Management Firmware Development, Raleigh
Charles Perkins, Course Developer, Service and Support Education, Raleigh
David Pfeffer, Global Service Planner, Beaverton
Steve Powell, Technical Advisor, Service & Support Education Team, Raleigh
Ajit Ravani, Firmware Project Manager, Austin
Mike Schiskey, Technical Support, x450 Development, Raleigh
Al Thomason, Senior Product Marketing Manager, Beaverton
Damon A. West, IBM @server Technology Enablement Center
Bob Zuber, x450 World Wide Product Manager, Raleigh

Also, thanks to the following people from Intel:

Scott Diaz
Tim Farrell
Dave Myron
John F Stewart

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our Redbooks to be as helpful as possible. Send us your comments about this or other Redbooks in one of the following ways:

- Use the online **Contact us** review redbook form found at:

  ibm.com/redbooks

- Send your comments in an Internet note to:

  redbook@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HZ8 Building 662
  P.O. Box 12195
  Research Triangle Park, NC 27709-2195

1

Technical description

The IBM @server xSeries 450 is the latest IBM top-of-the-line server and is the first full implementation of the 64-bit IBM XA-64 chipset, code named “Summit”, as part of the Enterprise X-Architecture strategy. The x450 completes the xSeries product family, leveraging the proven Enterprise X-Architecture used in the x440 to deliver commercially viable 64-bit systems.

The following are the key features of the x450:

- One or two-way Intel Itanium 2 models, upgradable to four-way
- 4U rack-dense design
- 64 MB XceL4 Server Accelerator Cache providing an extra level of cache
- 1 GB or 2 GB RAM standard, up to 40 GB total using 512 MB, 1 GB, and 2 GB ECC SDRAM DIMMs
- Memory enhancement such as memory mirroring, Chipkill™, and Memory ProteXion
- Six Active PCI-X slots: two 64-bit 133 MHz, two 64-bit 100 MHz, two 64-bit 66 MHz
- Connectivity to an RXE-100 external PCI-X enclosure for an additional 12 PCI-X slots
- Integrated dual-channel Ultra 320 SCSI controller
- Two hot-swap 1” drive bays

- Support for major storage subsystems, including Fibre Channel and ServeRAID™
- Light path diagnostics and the Remote Supervisor Adapter for systems management
- Integrated dual 10/100/1000 Mbps Ethernet controller

1.1 The x450 product line

Powered by XA-64 Enterprise X-Architecture and 64-bit Itanium 2 “Madison” processors, xSeries 450 servers bring the future of 64-bit processing and production-level reliability to your data centers today. Featuring mainframe-inspired, advanced mission-critical function, you can depend on these four-way-capable enterprise servers to run your complex business applications around the clock.

The models of the x450 that will be available on July 1 are listed in Table 1-1.

Table 1-1 Initial x450 models

Model	Standard processors	Max SMP	L2 cache	L3 cache	Std memory
8688-4RX	1x 1.3 GHz Intel Itanium 2	4-way	256 KB	3 MB	1 GB (2x 512 MB)
8688-5RX	1x 1.4 GHz Intel Itanium 2	4-way	256 KB	4 MB	2 GB (4x 512 MB)
8688-6RX	1x 1.5 GHz Intel Itanium 2	4-way	256 KB	6 MB	2 GB (4x 512 MB)

IBM had ealier announced models that were based on the Itanium 2 “McKinley” processor, but these have been withdrawn due to a bug in these processors. These models are listed in Table 1-2.

Table 1-2 Withdrawn “McKinley” processor-based models

Model	Standard processors	Max SMP	L2 cache	L3 cache	Std memory
8688-1RX	1x 900 MHz Intel Itanium 2	4-way	256 KB	1.5 MB	1 GB (2x 512 MB)
8688-2RX	2x 1.0 GHz Intel Itanium 2	4-way	256 KB	1.5 MB	2 GB (4x 512 MB)
8688-3RX	2x 1.0 GHz Intel Itanium 2	4-way	256 KB	3 MB	2 GB (4x 512 MB)

The x450 models support one, two, three, and four processors.

The attachment of a single RXE-100 Remote Expansion Enclosure is also supported using either a single or two remote I/O cables. The RXE-100 has six PCI-X slots standard, upgradable to 12 PCI-X slots, giving the customer up to a total of 12 PCI-X or 18 PCI-X slots respectively.

1.2 IBM XA-64 chipset

The IBM XA-64 chipset is the product name describing the chipset developed under the code name “Summit” and implemented on the IA-64 platform. A product of the IBM Microelectronics Division, the XA-64 chipset leverages the proven Enterprise X-Architecture chipset used in the x440 and applies those same technologies to the IA-64 architecture. The XA-64 chipset is composed of the following components:

- Memory controller: a single memory controller, code named “Cyclone”, located within the memory-board assembly.
- Processor/cache controller: a single processor and cache controller, code named “Tornado”, located within the processor-board assembly.
- PCI bridges: two PCI bridges, code named “Winnipeg”, per x450 located on the PCI-X board and the I/O board that control both the PCI-X and Remote I/O.

Figure 1-1 on page 4 shows the various IBM XA-64 components in an x450 configuration.

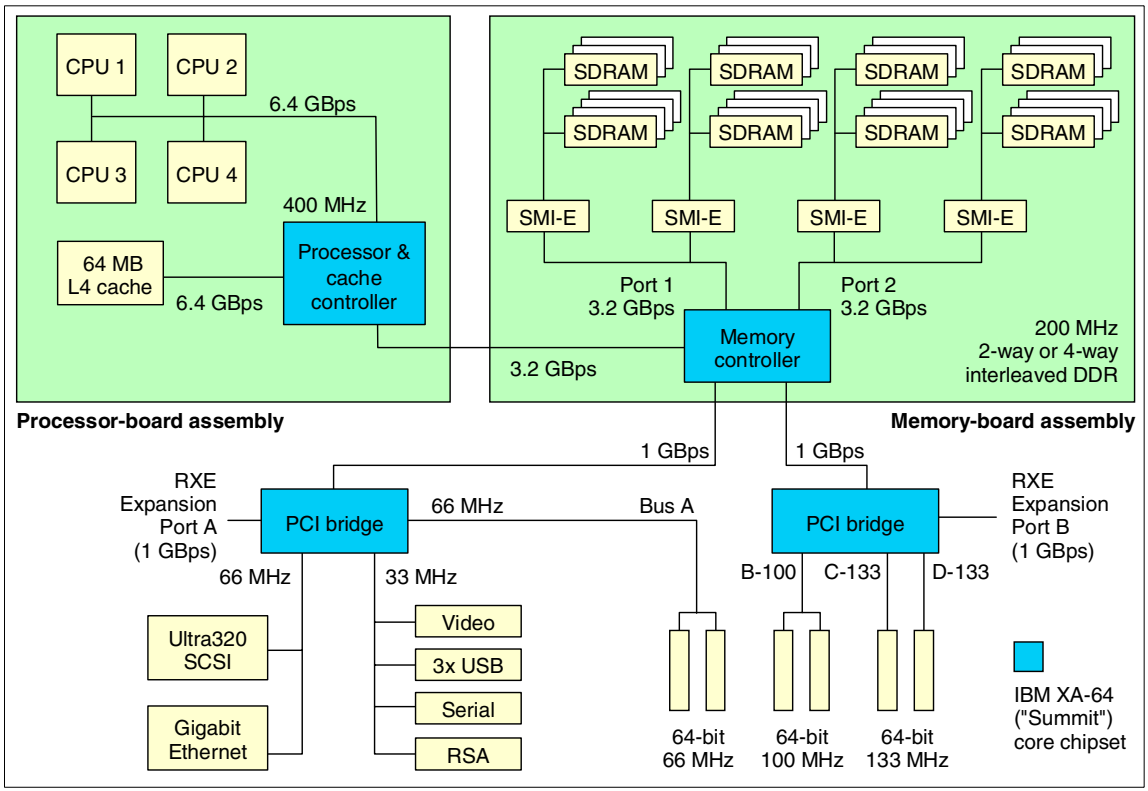

Figure 1-1 xSeries 450 system block diagram

What was called the SMP Expansion module in the x440 has been divided into two components in the x450. The component that contains the CPUs, processor/cache controller, and cache is called the *processor-board assembly*. The component that contains the memory controller and memory is called the *memory-board assembly*.

The CPUs are connected together with a 200 MHz frontside bus, but supply data at an effective rate of 400 MHz using the “dual-pump” design of the Intel Itanium 2 architecture as described in 1.4, “Intel Itanium 2 processor” on page 10. To ensure the processors are optimally used, the x450 has a 64 MB XceL4 Server Accelerator Cache, comprised of 200 MHz DDR memory. This L4 system cache services all CPUs.

Memory used in the x450 is standard PC2100 ECC DDR SDRAM DIMMs. With 2 GB DIMMs, up to 40 GB can be installed using 20 of the 28 DIMM sockets. The memory is two-way interleaved; however, four-way interleaving is also supported in order to ensure that the memory subsystem can supply data fast enough to match the throughput of the CPUs. Two-way interleaving means that DIMMs must be installed in matched pairs and in specific DIMM sockets (see 3.1.2, “Memory” on page 41)

The two PCI bridges in the XA-64 chipset provide support for 33, 66, 100, and 133 MHz devices using four PCI-X buses (labeled A-D in Figure 1-1 on page 4). This is discussed further in 1.8, “PCI subsystem” on page 21.

The PCI bridge also has two 1 GBps bi-directional Remote Expansion I/O (RXE) ports for connectivity to the RXE-100 enclosure. The RXE-100 provides up to an additional 12 PCI-X slots and can be connected by a single cable to port A.

The rear panel of the x450, indicating the location of the RXE Expansion Ports, is shown in Figure 1-2.

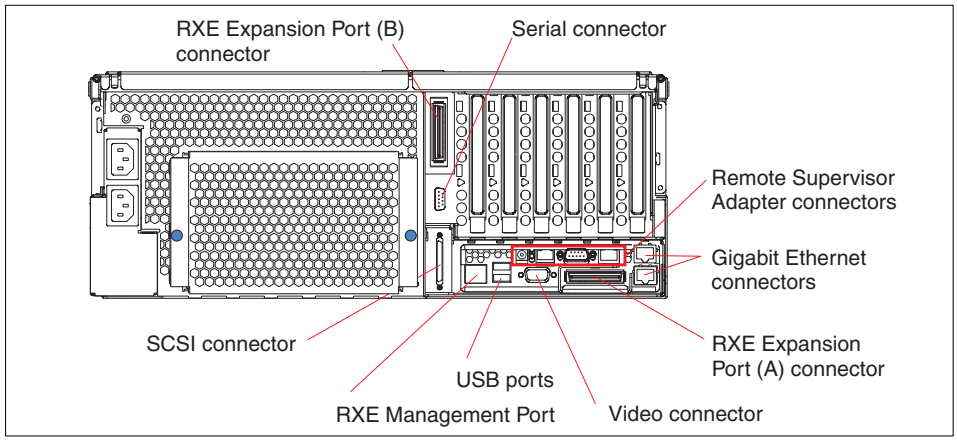

Figure 1-2 Rear view of the x450

1.3 Extensible Firmware Interface

The *Extensible Firmware Interface* (EFI) specification describes an interface between the operating system and platform firmware as shown in Figure 1-3 on page 6. The interface offers platform-related information to the operating system as well as boot and runtime service calls that are available to the operating system and OS loader. Together, these makes a well-defined environment for booting the operating system and running pre-boot applications, such as diagnostics and system setup.

In comparison to a BIOS-based, legacy system, the EFI is an additional layer between the operating system and the firmware. In a legacy system, the OS loader calls BIOS functions directly. Consequently, to provide a stable boot environment, changes in the OS loader and the platform firmware must go hand-in-hand.

Note: All operating systems supported on x450 are EFI-aware. The OS loader communicates with the firmware and hardware through the EFI interface.

The primary goal of this specification is to provide an abstract model both for operating system and hardware developers. With such a model in a place, OS loader customizations are not required if there are changes in the platform hardware or firmware (added new boot or input devices for instance). The EFI breaks up a tight dependency between the operating system and the firmware, thus speeding up the process of releasing the new products and introducing the new features and functionality (both operating system and hardware).

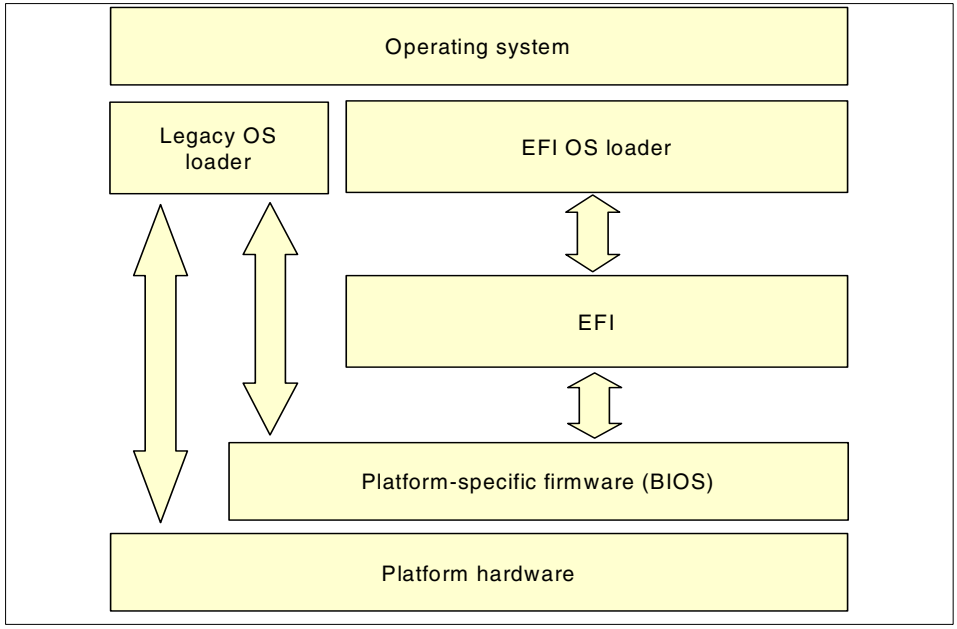

Figure 1-3 The EFI concept

Consider, for example, the situation where a new type of boot device, for example a USB key, is to be implemented. First the BIOS would have to offer an option to choose this new device for booting, then new USB key-specific functions would have to be added to the firmware to support booting from a USB device, and finally, the OS loader would have to be modified to use these functions.

The same situation with the EFI would be dramatically simplified. The OS loader calls unified (not vendor-specific) EFI API functions for booting. These functions are not dependent on the boot device used, so when a new boot device type is added to the platform and the firmware is modified to recognize it, the operating system can immediately boot.

The EFI architecture is modular, extensible and offers backward compatibility for the older systems by default. This means there is a way for non-EFI-aware operating systems to communicate directly with system BIOS as shown in Figure 1-3 on page 6.

Note: The EFI concept was originally introduced with Itanium Architecture-based computers, but is not restricted to 64-bit platforms. There is a gradual transition from BIOS to the EFI expected on the IA-32 platform as well.

1.3.1 GUID Partition Table disk

The *GUID Partition Table* (GPT) was introduced as part of the EFI initiative. Every disk is assigned a global unique identifier (GUID) to allow self-identification of the disks. GPT replaces the older Master Boot Record (MBR) partitioning scheme that has been common to PCs.

There are several reasons for introducing a new partitioning scheme:

- MBR disks support only four partition table entries. If more partitions are wanted, a secondary structure — an extended partition — is necessary. Extended partitions are then subdivided into one or more logical disks. On any given drive, only one extended partition can be present.

  In theory, a GPT disk can have an unlimited number of partitions. The number of partitions is limited only by the amount of space reserved for making partition entries.
- GPT disks use primary and backup partition tables for redundancy and CRC32 fields for improved partition data structure integrity.
- GPT disks can grow to a very large size. In theory, a GPT disk can be up to 2^{64} logical blocks in length (logical blocks are typically 512 bytes). In practice, the maximum is less. For example, Windows Server 2003 supports GPT disks up to approximately 18 Exabytes in size.

For backward compatibility with legacy MBR disk tools, all GPT disks contain a *protective MBR*. The protective MBR, beginning in sector 0, precedes the GUID Partition Table on the disk and contains only one partition that appears to span the disk. The legacy tools are not aware of GPT and do not know how to properly access a GPT disk. The benefit of protective MBR is that these tools will view a

GPT disk as having a single encompassing (possibly unrecognized) partition, rather than mistaking the disk for one that is unpartitioned. That is why GPT-partitioned disk appears to have MBR.

Note: GPT disks can be converted to MBR disks and vice versa only if all existing partitioning is first deleted, with associated loss of data.

In general, the structure of any GPT disk is as shown in Figure 1-4. The protective MBR is followed by a theoretically unlimited number of data partitions.

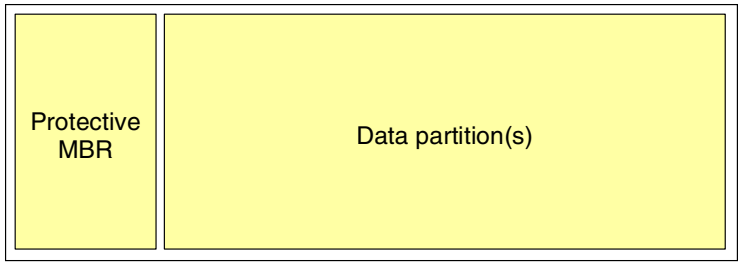

Figure 1-4 General GPT disk structure

Note: Currently, only 64-bit operating systems have the ability to read, write, and boot from GPT disks. 32-bit operating systems do not have built-in support for GPT disks.

The specification for GPT disk partitioning can be found in Chapter 16 of the Extensible Firmware Interface (EFI) specification. This document is available at:

`http://developer.intel.com/technology/efi/download.htm`

1.3.2 EFI System Partition

A special partition on the GPT disk is *EFI System Partition* (ESP). It contains the OS loader files of all installed operating systems. These files are stored in the EFI directory. The ESP may also contain other files necessary to boot the system, such as drivers.

Note: The EFI System Partition is shareable among all installed operating systems. To support multiple operating system installations, create multiple data partitions.

An example directory structure for an EFI System Partition present on a hard disk with SuSE and Windows Server 2003 installed is as follows:

```
\EFI
      \Microsoft
            \WINNT50
            \EFIDrivers
      \SuSE
\MSUtils
```

There can be only one ESP on a single disk. The size of the ESP is determined using the following algorithm:

```
ESP = max(100 MB, min(1% of physical disk, 1GB))
```

In other words, the size of the ESP must be the larger of these two numbers, 100 MB or 1% of the physical disk size (up to 1 GB). For example, for an 18 GB disk, the size of the ESP is 184 MB. The value 1% of the physical disk is calculated at the time that the ESP is created and does not change if the disk is extended later (for example, via RAID).

Note: Each bootable GPT disk must contain an EFI System Partition.

The ESP should be the first partition on the disk, right after protective MBR as shown in Figure 1-5.

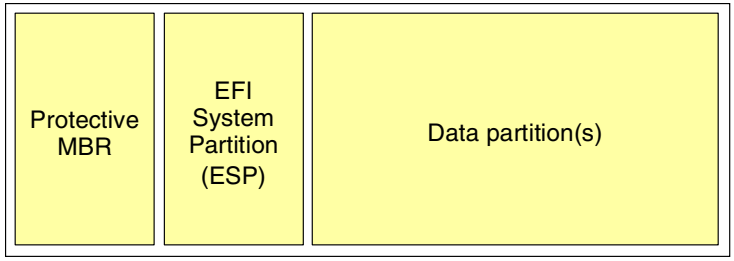

Figure 1-5 Boot GPT disk structure

The EFI specification supports only FAT or FAT32 on the ESP partition.

Note: The ESP is not visible to the operating systems users by default but can be accessed for read/write operations from within the operating system by special commands. For Windows-specific information, see “Accessing EFI System Partition from Windows” on page 96. For SuSE information, see “Partitions on IA-64 Linux” on page 111.

1.3.3 EFI and legacy-free concept

The EFI complements legacy-free concept of PCs. *Legacy-free* refers to PC system designs that eliminate certain hardware and firmware elements of the original PC architecture while advancing the PC's stability and usability. Specifically, we are talking about a set of I/O options that have been part of the PC architecture for a very long time, for example parallel, serial and game port, ISA slots or devices, floppy disk controller (FDC), PS/2® mouse, and keyboard.

BIOS interfaces require the OS loader to have a specific knowledge of the workings of hardware devices. The EFI abstract concept makes it possible to build code that works on a range of hardware devices without having explicit knowledge of the specifics of these devices. This EFI feature allows a replacement of legacy devices and adding new boot devices. The legacy devices are replaced by USB devices in x450.

Tip: One key interface that is no longer supported is INT 13, disk I/O.

For more information about the EFI specification, see the following:

http://www.intel.com/technology/efi/index.htm
http://www.microsoft.com/hwdev/platform/firmware/EFI/default.asp

1.4 Intel Itanium 2 processor

The Itanium 2 processor used in the x450 (code named “Madison”) uses a ZIF socket design, although the socket is designed differently from the one on the x440, for example. This small form factor is what permits the x450 to have up to four processors in a 4U node.

Table 1-3 outlines some of the differences between the Itanium and Itanium 2 processors (both the “Madison” and the earlier “McKinley” processor):

Table 1-3 Itanium vs Itanium 2 processors

Feature	Itanium	Itanium 2 “McKinley”	Itanium 2 “Madison”
Processor core speed	733 or 800 MHz	900 MHz or 1.0 GHz	1.3, 1.4 or 1.5 GHz
L3 Cache	2 or 4 MB	1.5 or 3 MB	3, 4 or 6 MB
Frontside bus	266 MHz	400 MHz, 128 bit	400 MHz, 128 bit
Frontside bus bandwidth	2.1 GBps	6.4 GBps	6.4 GBps
Pipeline stages	10	8	8

Feature	Itanium	Itanium 2 “McKinley”	Itanium 2 “Madison”
Issue ports	9	11	11
on-board registers	328	328	328
Integer units	3	6	6
Branch units	3	3	3
Floating point units	2	2	2
SIMD units	2	1	1
Load and store units	2 (total)	2 load and 2 store	2 load and 2 store

The Itanium 2 processor has three levels of cache, all of which are on the processor die:

- Level 3 cache is equivalent to L2 cache on the Pentium III Xeon, or the L3 cache on the Pentium Xeon MP processor. Itanium 2 processors in the x450 models contain either 3, 4 or 6 MB of L3 cache. Unlike the design of the original Itanium processor, this L3 cache is now on the processor die, greatly improving performance, up or 2 times greater than that of the original Itanium.
- Level 2 cache is equivalent to L1 cache on the Pentium III Xeon and is 256 KB in size.
- A new level 1 cache, 32 KB in size, is “closest” to the processor and is used to store micro-operations (that is, decoded executable machine instructions) and serves those to the processor at rated speed. This additional level of cache saves decode time on cache hits.

The x450 also implements a Level 4 cache as described in 1.6, “IBM XceL4 Server Accelerator Cache” on page 17.

Intel has also introduced a number of features associated with its Itanium micro-architecture. These are available in the x450, including:

- 400 MHz frontside bus

  The Pentium III Xeon processor had a 100 MHz frontside bus that equated to a burst throughput of 800 MBps. With protocols such as TCP/IP, this had been shown to be a bottleneck in high-throughput situations. The Itanium 2 processor improves on this by using a single 200 MHz clock and using both edges of each clock to transmit data. This is shown in Figure 1-6 on page 12.

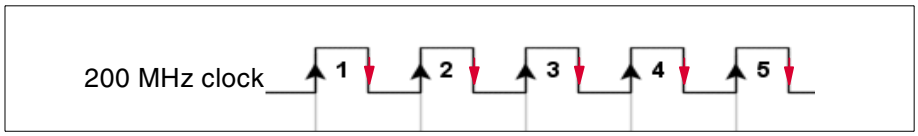

Figure 1-6 Dual-pumped frontside bus

This increases the performance of the frontside bus. The end result is an effective burst throughput of 6.4 GBps (128-bit wide data path running at 400 MHz), which can have a substantial impact, especially on TCP/IP-based LAN traffic. This is opposed to the Itanium processor, which had a burst throughput of only 2.1 GBps (64-bit wide data path running at 266 MHz).

- Explicitly Parallel Instruction Computing (EPIC)

  EPIC technology, developed by Intel and HP, leads to more efficient, faster processors because it eliminates numerous processing inefficiencies in current processors and attacks the perennial data bottleneck problems by increasing parallelism, rather than simply boosting the raw “clock” speed of the processor.

  Specifically, in today's 32-bit processors, much of the instruction scheduling--the order in which computing instructions are executed--is done on the chip itself, leading to a great deal of overhead and slowing down overall processor performance. Moreover, today's processors are plagued by instruction flow problems since the processor often has to stop what it's doing and reconstruct the instruction flow due to inherent inefficiencies in instruction handling.

  EPIC makes the instruction scheduling more intelligent and handles much of the scheduling off-chip, in the compiler program, before feeding “parallelized” instructions to the Itanium 2 processor for execution. The parallelized instructions allow the chip to process a number of instructions simultaneously, increasing performance. A compiler prepares instructions for execution on the processor.

  The Itanium 2 architecture is based on EPIC technology and has the following features:

  – Provides faster online transaction processing
  – Has the capability to execute multiple instructions simultaneously, processing more data and allowing more users
  – Enables faster calculations and data analysis
  – Allows for faster storage and movement of large models (CAD, CAE)
  – Speeds up simulation and rendering times

For more information about the features of the Itanium 2 processor, go to:

http://www.intel.com/design/itanium2

1.5 System assembly

Unlike the x440, then x450 does not use a single SMP Expansion Module, which contains the processors, memory, XceL4 Server Accelerator Cache, and respective controller. Instead, there are now two separate assemblies:

- A single memory-board assembly, located at the top of the system, which contains the physical DIMM's as well as the memory and I/O controller
- A single processor-board assembly, located at the bottom of the system, which contains the processors, XceL4 Server Accelerator Cache, and L4 Cache/Scalability Controller

See Figure 1-7 on page 14 for the location of the memory-board and processor-board assemblies.

Tip: Be careful when removing or installing either the memory-board assembly or the processor-board assembly, because it is possible to damage the midplane.

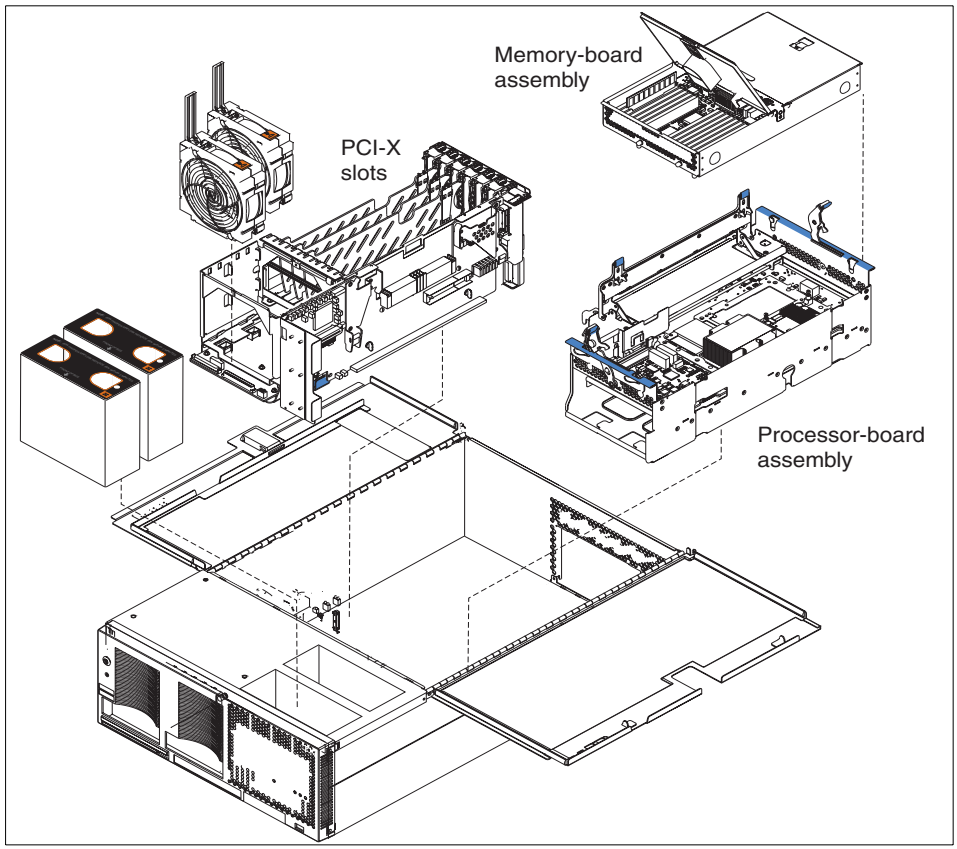

Figure 1-7 Memory-board and processor-board assembly locations

1.5.1 Memory-board assembly

The x450 memory-board assembly is installed from the top of the server and mounts to the side of the midplane using two levers on the top. This location allows for easy access to all memory DIMMs without having to remove any components from the system.

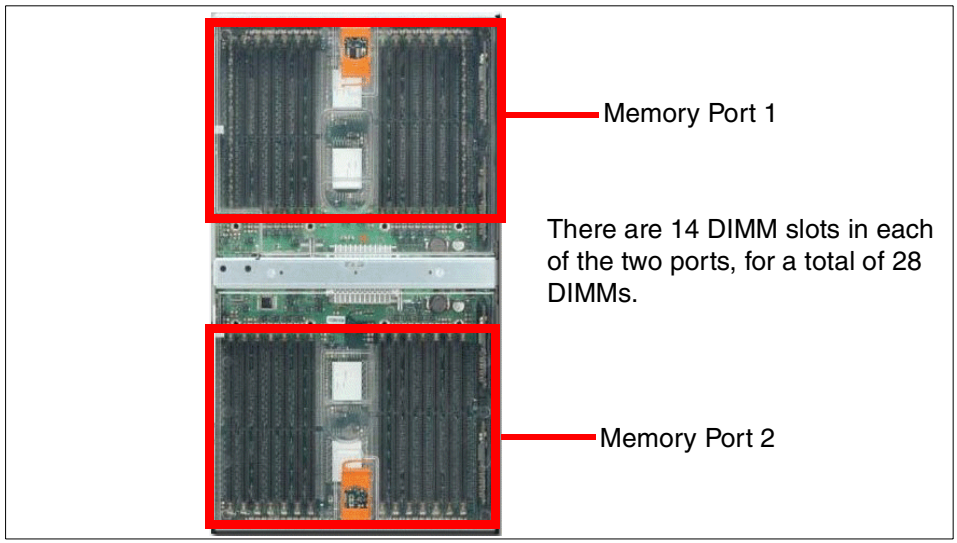

Figure 1-8 Memory-board assembly, showing the two memory ports

The memory-board assembly contains 28 DIMM slots. All DIMM slots can be used when 512 MB or 1 GB DIMMs are used. If 2 GB DIMMs are used, up to 20 DIMMs slots can be used. The maximum installable memory is 40 GB (using 2 GB DIMMs).

The memory-board assembly is also equipped with LEDs for light path diagnostics for each DIMM. In addition, the assembly is equipped with LEDs for the following:

- Power to memory port 1
- Power to memory port 2
- Hot-plug memory enabled

Restriction: The ability to hot-add or hot-replace memory is not available in the x450.

For information about installing memory in the memory-board assembly, see 3.1.2, “Memory” on page 41.

1.5.2 Processor-board assembly

The x450 processor-board assembly is located under the memory-board assembly, as shown in Figure 1-7 on page 14. It is installed from the top of the server and mounts to the side of the midplane using two levers on the side, as shown in Figure 1-9 on page 16.

Tip: Be careful when removing or installing either the memory-board assembly or the processor-board assembly, since it is possible to damage the midplane.

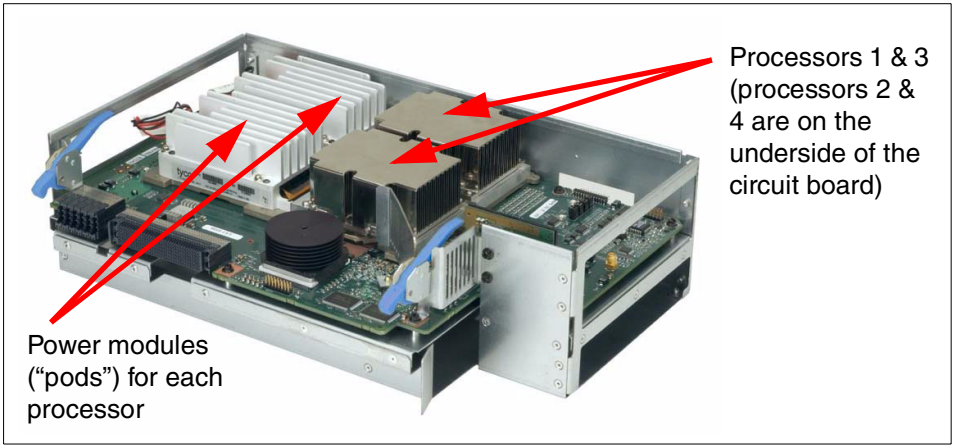

Figure 1-9 Processor-board assembly

x450 models have either one or two processors installed. The unused CPU sockets will hold metal air baffles. The power pods shown in Figure 1-9 supply power to the processors and are equivalent to VRMs in other systems.

Important: While processors should be installed in the order listed in Figure 1-9, the bootstrap processor (BSP) may not necessarily be the processor located in Processor Socket 1. The Intel Itanium Architecture processors are initialized and tested in parallel. The first processor to complete initialization becomes the BSP.

The processor-board assembly is also equipped with LEDs for light path diagnostics for the following components:

- Each processor
- Each power module (“pod”)

In addition, a “remind” button is located on the upper side of the processor-board assembly. Pressing this button while the processor-board assembly is not attached to AC power will illuminate any light path LEDs that had been lit while the system was under power for a total of 10 seconds.

1.6 IBM XceL4 Server Accelerator Cache

Integrated into the processor-board assembly is 64 MB of Level 4 cache, which is shown in Figure 1-1 on page 4. This XceL4 Server Accelerator Cache provides the necessary extra level of cache to maximize CPU throughput by reducing the need for main memory access under demanding workloads, resulting in an overall enhancement to system performance.

Cache memory is two-way interleaved 200 MHz DDR memory and is faster than the main memory because it is directly connected to the memory controller and does not have additional latency associated with the large fan-out necessary to support the 28 DIMM slots. Since the data interface to the controller is 400 MHz, peak bandwidth for the XceL4 cache is 6.4 GBps.

1.7 System memory

The x450 has 1 GB or 2 GB of RAM standard, depending on the model. Memory packaging is PC2100 ECC DDR SDRAM DIMMs, and standard memory is either two or four 512 MB DIMMs. Memory options are 512 MB, 1 GB, or 2 GB DIMMs.

There are a total of 28 DIMM sockets (two ports of 14). All 28 DIMM sockets can be used to install DIMMs, with the exception of the 2 GB DIMM option. If 2 GB DIMMs are installed, the total number of DIMM sockets that can be used is limited to 20. A maximum of 40 GB of system memory is supported by populating 20 DIMM sockets each with a 2 GB DIMM.

DIMMs must be installed in matched pairs, since the DIMMs are two-way interleaved. However, if memory is installed in matched fours (a matched pair in each port), the system automatically detects this and will enable four-way interleaving. With this, memory access is performed simultaneously from both ports (two separate paths into the memory controller as shown in Figure 1-1 on page 4), leading to improved memory performance.

See 3.1.2, “Memory” on page 41 for a further discussion of how memory is implemented in the x450 and what you should consider before an x450 installation.

There are a number of advanced features implemented in the x450 memory subsystem, collectively known as *Active Memory*:

Restriction: The ability to hot-add or hot-replace memory is not available in the x450.

- Memory ProteXion

  Memory ProteXion, also known as “redundant bit steering”, is the technology behind using redundant bits in a data packet to provide backup in the event of a DIMM failure.

  Currently, other industry-standard servers use 8 bits of the 72-bit data packets for ECC functions and the remaining 64 bits for data. However, the x450 uses an advanced ECC algorithm that is based not on bits but on memory symbols. Symbols are groups of multiple bits, and in the case of the x450, each symbol is 4 bits wide. With two-way interleaved memory, the algorithm needs only three symbols to perform the same ECC functions, thus leaving one symbol free (2 bits on each DIMM). See Figure 1-10.

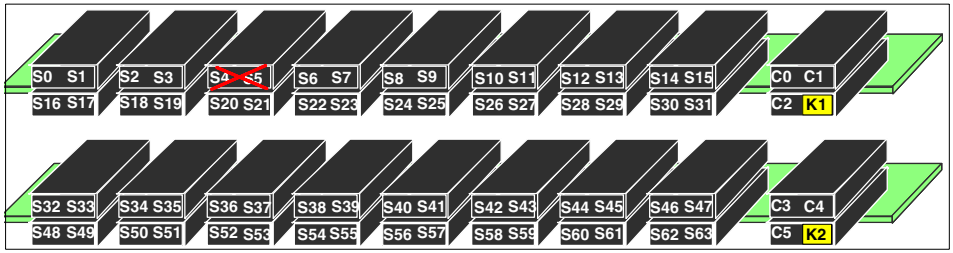

Figure 1-10 Memory ProteXion

  In the event that a chip failure on the DIMM is detected by memory scrubbing, the memory controller can re-route data around that failed chip through the spare symbol (similar to the hot-spare drive of RAID array). It can do this automatically without issuing a Predictive Failure Analysis® (PFA) or light path diagnostics alert to the administrator. After the second DIMM failure, PFA and light path diagnostics alerts would occur on that DIMM as normal.

- Memory scrubbing

  Memory scrubbing is an automatic daily test of all the system memory that detects and reports memory errors that might be developing before they cause a server outage.

  Memory scrubbing and Memory ProteXion work in conjunction with each other, but they do not require memory mirroring (as described below) to be enabled to work properly.

  When a bit error is detected, memory scrubbing determines if the error is recoverable or not. If it is recoverable, Memory ProteXion is enabled and the data that was stored in the damaged locations is rewritten to a new location. The error is then reported so that preventative maintenance can be performed.

  As long as there are enough good locations to allow the proper operation of the server, no further action is taken other than recording the error in the error

logs. Errors from scrubbing are not reported unless bit steering has already been invoked.

If the error is not recoverable, then memory scrubbing sends an error message to the light path diagnostics, which then turns on the proper lights and LEDs to guide you to the defective DIMM. If memory mirroring is enabled, then the mirrored copy of the data in the damaged DIMM is used until the system is powered down and the DIMM replaced.

- Memory mirroring

  Memory mirroring is roughly equivalent to RAID-1 in disk arrays, in that memory is divided in two ports and one port is mirrored to the other half (see Figure 1-11). If 8 GB is installed, then the operating system sees 4 GB once memory mirroring is enabled (it is disabled in BIOS by default). All mirroring activities are handled by the hardware without any additional support required from the operating system.

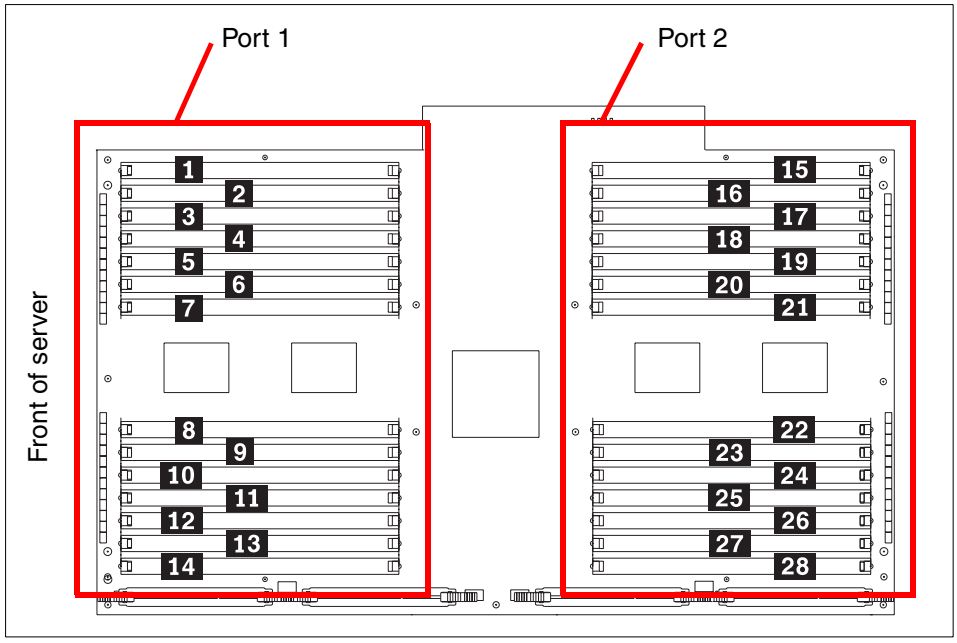

Figure 1-11 Memory DIMMs are divided into two ports

When memory mirroring is enabled (see “Enabling memory mirroring” on page 83), the data that is written to memory is stored in two locations. One copy is kept in the port 1 DIMMs, while a second copy is kept in the port 2 DIMMs.

During the execution of the read command, the data is read simultaneously from both ports, and error-free data from either port is forwarded. This

provides an extra level of error recovery capability. (In the x440, the read command is read from the DIMM with the least amount of reported memory errors through memory scrubbing).

If memory scrubbing determines the DIMM is damaged beyond use, read and write operations are redirected to the partner DIMM in the other port. Memory scrubbing then reports the damaged DIMM and the light path diagnostics display the error. If memory mirroring is enabled, then the mirrored copy of the data in the damaged DIMM is used until the system is powered down and the DIMM replaced.

Certain restrictions exist with respect to placement and size of memory DIMMs when memory mirroring is enabled. These are discussed in “Memory mirroring” on page 42.

- Chipkill memory

  Chipkill is integrated into the XA-64 chipset and does not require special Chipkill DIMMs. Chipkill corrects multiple single-bit errors to keep a DIMM from failing. When combining Chipkill with Memory ProteXion and Active Memory, the x450 provides very high reliability in the memory subsystem. Chipkill memory is approximately 100 times more effective than ECC technology, providing correction for up to 4 bits per DIMM, whether on a single chip or multiple chips.

  If a memory chip error does occur, Chipkill is designed to automatically take the inoperative memory chip offline while the server keeps running. The memory controller provides memory protection similar in concept to disk array striping with parity, writing the memory bits across multiple memory chips on the DIMM. The controller is able to reconstruct the “missing” bit from the failed chip and continue working as usual.

  Chipkill support is provided in the memory controller and implemented using standard DIMMs, so it is transparent to the operating system.

In addition, to maintain the highest levels of system availability, if a memory error is detected during POST or memory configuration, the server can automatically disable the failing memory bank and continue operating with reduced memory capacity. You can manually re-enable the memory bank after the problem is corrected via the Setup menu in BIOS.

Memory mirroring, Chipkill, and Memory ProteXion provide multiple levels of redundancy to the memory subsystem. Combining Chipkill with Memory ProteXion enables up to two memory chip failures per memory port (14 DIMMs) on the x450. An x450 with its two memory ports could sustain up to four memory chip failures. Memory mirroring provides additional protection with the ability to continue operations with memory module failures.

1. The first failure detected by the Chipkill algorithm on each port doesn't generate a light path diagnostics error, since Memory ProteXion recovers from the problem automatically.
2. Each memory port could then sustain a second chip failure without shutting down.
3. Provided that memory mirroring is enabled, the third chip failure on that port would send the alert and take the DIMM offline, but keep the system running out of the redundant memory bank.

1.8 PCI subsystem

As shown in Figure 1-12 on page 22, there are six PCI-X slots internal to the x450:

- Two 133 MHz slots, which accept 32-bit or 64-bit, 3.3 V, PCI or PCI-X adapters, from 33-133 MHz
- Two 100 MHz slots, which accept 32-bit or 64-bit, 3.3 V, PCI or PCI-X adapters, from 33-100 MHz
- Two 66 MHz slots, which accept 32-bit or 64-bit, 3.3 V, 33 or 66 MHz, PCI or PCI-X adapters

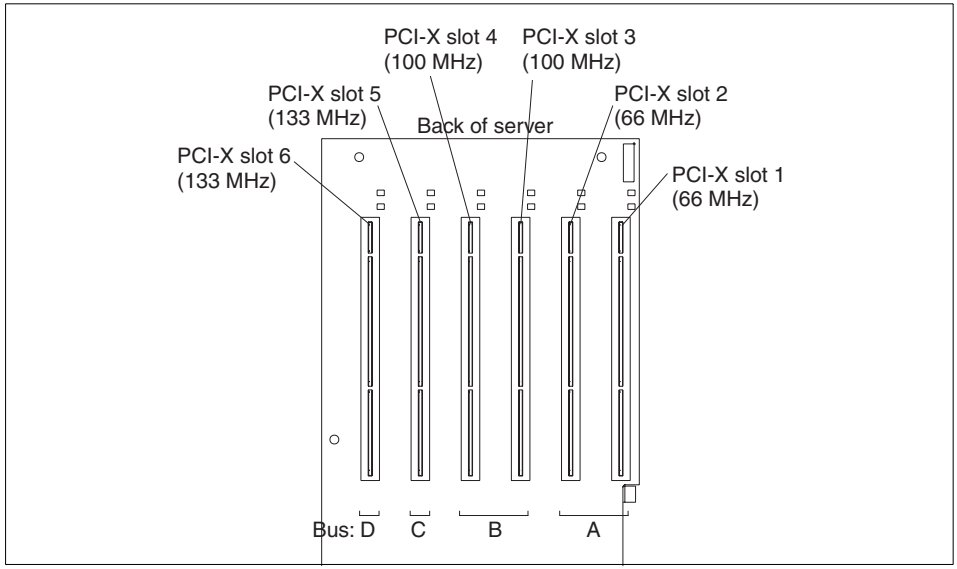

Figure 1-12 PCI slot information

See 3.1.3, “PCI slot configuration” on page 44 for details on what adapters are supported and in what combinations.

The PCI subsystem also supplies these I/O devices:

- Two Wide Ultra320 SCSI ports, one internal and one external (LSI LSI53C1030 chipset). This SCSI controller supports both single disks and RAID-1 mirrored pairs of disks.
- Dual Gigabit Ethernet ports (Broadcom 5704 chipset).

  The x450 includes a dual-port Broadcom BCM5704 10/100/1000 Base-T MAC (Media Access Controller) on a PCI 64-bit 66 MHz bus.

  The BCM5704 supports full and half-duplex performance at all speeds (10/100/1000 Mbps, auto-negotiated) and includes integrated on-chip memory for buffering data transmissions to ensure the highest network performance and dual onboard RISC processors for advanced packet parsing and backwards compatibility with today's 10/100 network. The Broadcom controller also includes software support for failover, layer-3 load balancing, and comprehensive diagnostics.

  Category 5 or better Ethernet cabling is required with RJ-45 connectors. If you plan to implement a Gigabit Ethernet connection, ensure your network infrastructure is capable of the necessary throughput to match the server’s I/O capacity.
- SVGA with 8 MB video memory (ATI RageXL chipset).

- Three USB ports (one on the front panel, two on the rear). All USB ports are 2.0 compliant.
- One RS-232 serial port, located on the rear of the machine.
- Remote Supervisor Adapter (RS-485 ASM interconnect bus, 10/100 Ethernet and serial ports).

Note: There are no PS/2 keyboard or mouse ports on the x450. USB keyboard and mice are supported, as well as serial connections via the integrated serial port.

If you require KVM support the 1.5 M USB Conversion Option (UCO) (part number 73P5832) enables the x450 to be attached to one of the Advanced Connectivity Technology (ACT) switches for common management within the rack. This smart cable is plugged into the USB and video ports on the server. It converts KVM signals to CAT5 signals for transmission over a CAT5 cable to either a Remote Console Manager (RCM) or Local Console Manager (LCM). USB servers can be managed on the same set of switches as legacy PS/2-based or C2T-based KVM servers.

With the addition of an RXE-100 Remote Expansion Enclosure, you can connect an additional six or 12 PCI-X adapters to the x450. See 3.2.2, “Remote Expansion Enclosure” on page 51 for details.

Note: Currently, only one RXE-100 can be connected to an x450 configuration.

1.9 Redundancy

The x450 has the following redundancy features to maintain high availability:

- Four hot-swap multi-speed fans

  With four hot-swap redundant fans, the x450 has adequate cooling for each of its major component areas. There are two fans located at the front of the server that direct air through the memory-board assembly and processor-board assembly. These fans are accessible from the top of the server without having to open the system panels. In the event of a fan failure, the other fan will speed up to continue to provide adequate cooling until the fan can be hot-swapped by the IT administrator.

  The other two fans are located just behind the power supplies and provide cooling for the I/O devices. Similar to the SMP Expansion Module fans, these fans will speed up in the event that one should fail to compensate for the reduction in air flow. In general, failed fans should be replaced within 24 hours following failure.

The four fans are shown in Figure 1-7 on page 14.

Important: Due to airflow requirements, fans should not be removed for longer than two minutes. The fan compartments need to be fully populated even if the fan is defective. Therefore, remove a defective fan only when a new fan is available for immediate replacement.

- Two hot-swap power supplies with separate power cords.

  Note: For large configurations, redundancy is achieved only when connected to a 220 V power supply. See 3.5, “Power considerations” on page 63 for details.

Note: To ensure adequate power, a UPS with a rating of RMB 5000 or more is recommended.

- Two hot-swap hard disk drive bays. Using the onboard LSI chipset, these drives can be configured to form a RAID-1 disk array for the operating system.
- The memory subsystem has a number of redundancy features, including memory mirroring, as described in 1.7, “System memory” on page 17.

The layout of the front panel of the x450, showing the location of the drive bays, power supplies and fans, is shown in Figure 1-13.

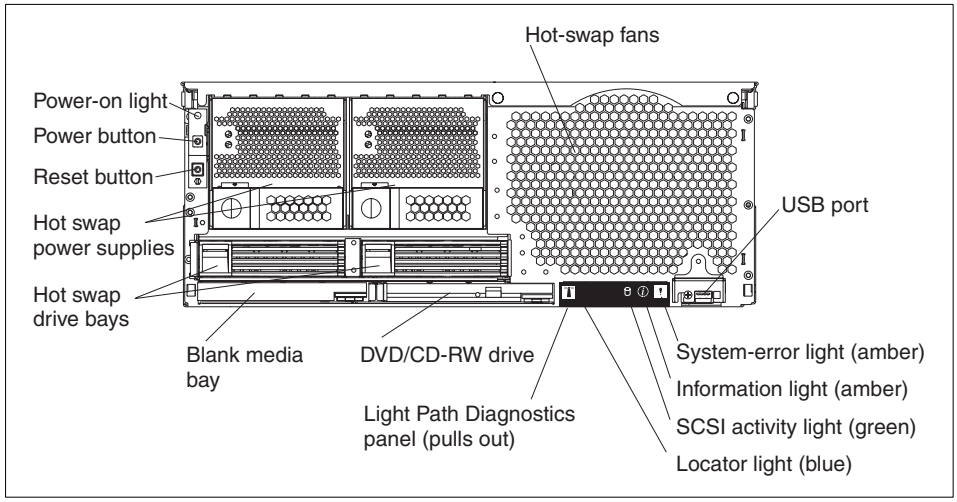

Figure 1-13 Front panel of the xSeries 450

1.10 Light path diagnostics

To limit the need to slide the server out of the rack to diagnose problems, a light path diagnostics panel has been incorporated in the front of the x450, as shown in Figure 1-14. This panel can be ejected from the server to view all light path diagnostics-monitored server subsystems. In the event that maintenance is then required, the customer can slide the server out from the rack and, using the LEDs, find the failed or failing component.

As illustrated in Figure 1-14, light path diagnostics is able to monitor and report on the health of CPUs, main memory, hard disk drives, PCI-X and PCI slots, fans, power supplies, VRMs, and the internal system temperature.

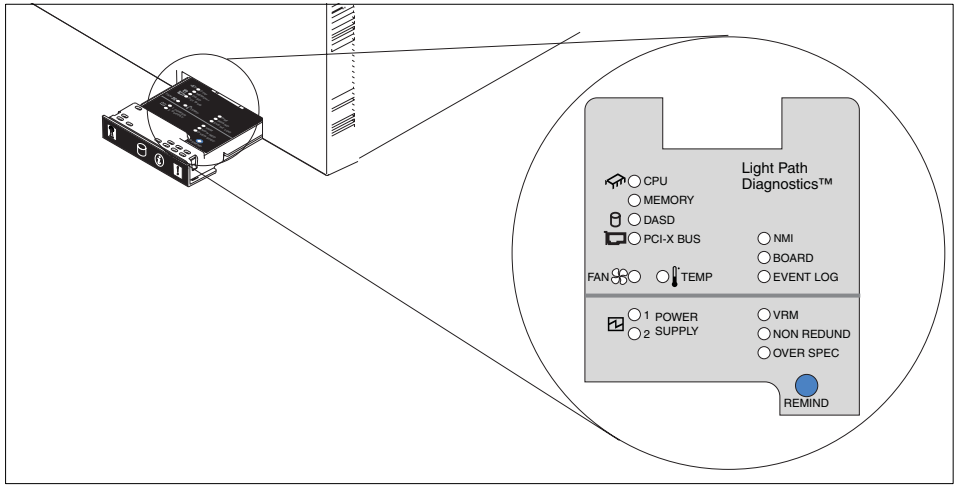

Figure 1-14 Light path diagnostics panel on the x450

Important: If a light path diagnostics LED has been illuminated and system power is removed, there is no way to redisplay the LEDs on the system tray without re-applying AC power. If the fault has not been rectified when power is restored, the LED will re-light.

The light path diagnostics on the x450 has four levels:

1. The first level is the front panel fault LED.
2. Level 2 is the pop-out panel as shown in Figure 1-14.
3. For further investigation, there are light path diagnostics LEDs visible through the top of the server. This requires the server to be slid out of the rack.

4. For the fourth level of diagnostics, LEDs on major system components indicate the component causing the error.

As the processor-board assembly is not visible during normal operation, a light path diagnostics button has been incorporated into it to assist with diagnosing errors. You can light up the LEDs for a maximum of 2 minutes. After that time, the circuit that powers the lights is exhausted.

The pop-out panel (Figure 1-14 on page 25) also has a remind button. This places the front panel system-error LED into remind mode, which means it flashes briefly every 2 seconds. By pressing the button, you acknowledge the failure but indicate that you will not take immediate action. If a new failure occurs, the system-error LED will turn on again and no longer blink. The system-error LED remains in the remind mode until one of the following situations occurs:

- All known problems are resolved
- The system is restarted
- A new problem occurs, at which time it then is illuminated continuously

Tip: The remind button on the pop-out LPD panel does not function when AC power has been removed from the system. The button is just used to acknowledge a system error as described above.

1.11 Remote Supervisor Adapter

The x450 includes a Remote Supervisor Adapter (RSA), which is positioned horizontally in a dedicated PCI slot beneath the PCI-X adapter area of the system.

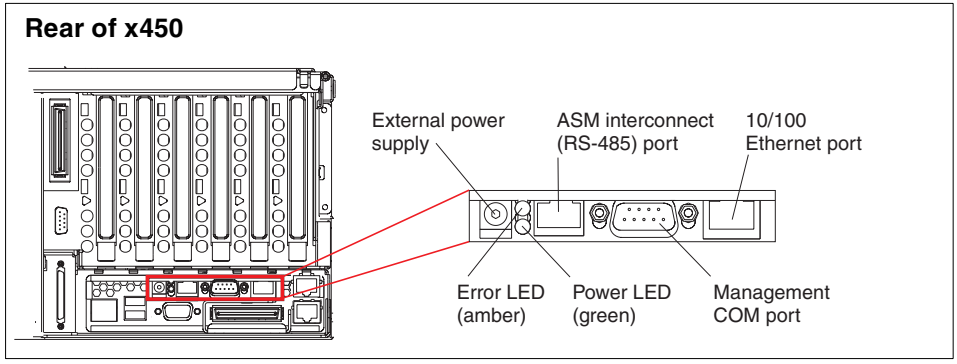

Figure 1-15 Remote Supervisor Adapter connectors

The Remote Supervisor Adapter offers the following capabilities:

- In-band and out-of-band remote server access and alerting through IBM Director
- Full Web browser support with no other software required
- Enhanced security features
- Graphics/text console redirection for remote control
- Dedicated 10/100 Ethernet access port
- ASM interconnect bus for connection to other service processors
- Serial dial in/out
- E-mail, pager and SNMP alerting
- Event log
- Predictive Failure Analysis on memory, power, hard drives, and CPUs
- Temperature and voltage monitoring with settable threshold
- Light path diagnostics
- Automatic Server Restart (ASR) for operating system and POST
- Wake on LAN®
- Remote firmware update
- LAN access
- Alert forwarding

See the IBM Redbook *Implementing Systems Management Solutions using IBM Director*, SG24-6188 for more information on the Remote Supervisor Adapter.

1.12 Operating system support

In line with the overall message of providing application flexibility to meet the varying needs of our enterprise customers, the x450 is optimized for several operating system and application solutions. Table 1-4 lists the supported operating systems for the x450. For the latest operating system support information, go to:

http://www.pc.ibm.com/us/compat/nos/matrix.shtml

Note: Hyper-Threading is not available on systems with the Itanium 2 processor.

Table 1-4 x450 operating system support

Description	Release	SMP support
Windows Server 2003 Enterprise (64-bit)		Supports up to four-way
SuSE Linux Enterprise Server	8.x	Supports up to four-way

Note: While operating systems may support large systems, scalability is a function of both the operating system and the application/workload. Few applications are designed to take advantage of larger SMP systems.

2

Positioning

In this chapter we discuss topics that help you to understand how the x450 can be useful for your business. The topics covered are:

- x450 application solutions
- Why choose x450?

2.1 x450 application solutions

With the x450 in the IA-64 environment, you are ready to deploy even larger implementations of enterprise solutions.

As companies' performance demands grow, 64-bit technology becomes an increasingly attractive option, due to increased memory addressability and true parallel architecture. There are a number of ways the x450 can be deployed in specific application solution environments. These include:

- Database applications
- Business logic
- e-business and security transactions
- In-house developed compute-intensive applications
- Science and technology

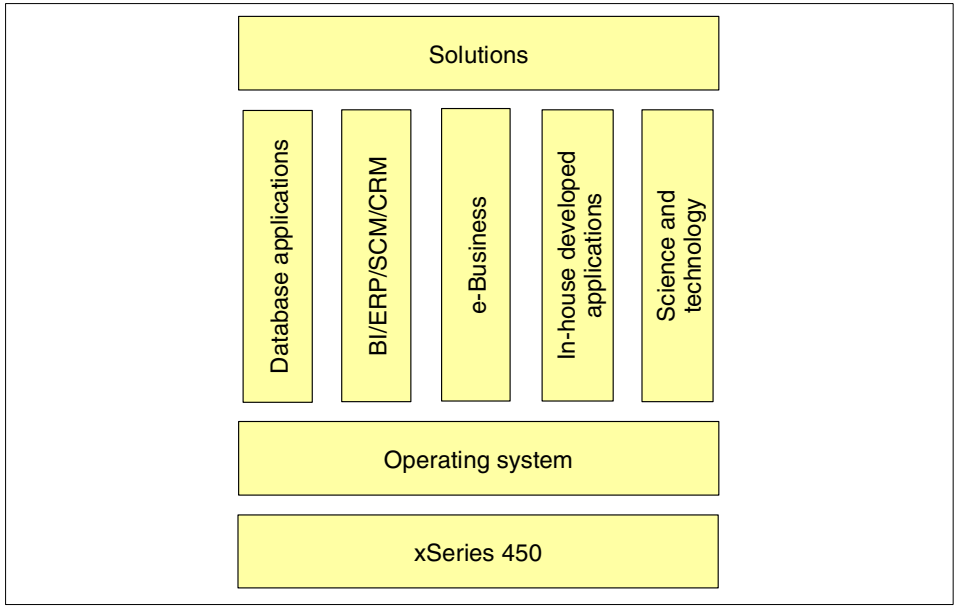

Figure 2-1 xSeries 450-based solutions

2.1.1 Database applications

Four-way x450 configurations can be used as database servers and application servers providing a highly available platform. These configurations require an external storage enclosure or SAN, depending on the size of the database that is driven by the number of users.

Database applications with memory-sensitive workloads that require working data sets larger than 4 GB to be loaded in memory will benefit from the larger memory support of the 64-bit platform.

The following is an example from the field. Microsoft SQL Server Enterprise Edition uses AWE memory only for the buffer pool. The AWE (Advanced Windowing Extensions) API allows applications that are written to use the AWE API to access more than 4 GB of RAM (basically anything between 4 GB and 64 GB). However, due to the AWE mapping overhead, it is not practical to try to use it for sort areas, procedure cache, or any other type of work area. Many applications do make heavy use of these areas and may not benefit by having the large buffer pool. The most efficient solution in such cases is to move the applications on to a 64-bit database server, which can access memory area above 4 GB much more efficiently without AWE's overhead.

Users will see a reasonable performance improvement on a 64-bit implementation over a 32-bit implementation. Through the 64-bit memory paths, all memory operations move twice as much data at the same speed as the 32-bit platform.

The database server will also benefit from a larger 3 MB third-level and 64 MB XceL4 cache. With such large cache, the need to go to memory or disk for database transaction elements is greatly reduced and this directly implies a performance increase, faster access to data, and improved throughput. Itanium 2 systems are likely to be able to hold database transaction records in cache during the entire transaction, which enables the I/O portion of the transaction to occur at speeds faster than memory access.

In-memory databases

Architectures with 64-bit addresses can store reasonably large databases in memory and access them with little or no paging overhead. This is often done for databases that are constantly being accessed and for databases that serve as the basis for complex analysis. The theoretical maximum of 16 Exabytes for memory has not yet been tested, but multi-Gigabyte databases are frequently run on 64-bit machines.

A major challenge to providing high-performance access to database information is the time it takes to access disk drives. When disk access is required, disk access times add what can be an intolerable delay to efficient information access and utilization. Access to disk is typically hundreds to thousands times slower than access to memory.

Today, the disk access time challenge can be overcome. The price of random access memory has come down to affordable levels for many systems. This price

reduction means that an entire database can be stored in system memory if the system processors can provide a very large linear address space.

A processor that supports a 64-bit address space may provide access to in-memory databases that range from tens of Gigabytes to thousands of Terabytes. In contrast, traditional 32-bit processors most often only address a maximum of 4 Gigabytes.

Key database software available for Itanium 2 systems include IBM DB2® Universal Database™ 8.1 (both Linux and Windows version), Microsoft SQL Server 2000 (currently beta 2 version), and Oracle 9i Database (Linux/Windows).

2.1.2 Business logic

More and more enterprise applications such as ERP, SCM, CRM, and BI are released or announced to be released on a 64-bit platform. Such applications process large amounts of data and the large flat memory model means this processing will be more efficient. That combined with up to four Itanium 2 processors and a highly efficient cache system make the x450 an ideal choice.

Market leaders offer 64-bit optimized versions of their enterprise applications for use on the x450 now, including SAS Release 9.0 (Windows version) and SAP R/3 4.6C (Windows/Linux), among others.

Independent software vendors such as JD Edwards, Baan/Invensys, i2 Technologies, PeopleSoft, Veritas, Computer Associates, BMC Software and many others already offer a variety of products available in 64-bit version.

Business intelligence

Business intelligence (BI) is a broad category of applications and technologies for gathering, storing, analyzing, and providing access to data to help enterprise users make better business decisions. BI applications include the activities of decision-support systems, query and reporting, online analytical processing (OLAP), statistical analysis, forecasting, and data mining.

The x450 brings high I/O bandwidth and performance to handle compute-intensive BI applications. A world-class floating-point engine, capable of addressing much larger amounts of memory, helps speed up data-intensive BI applications that help companies to increase employee productivity.

Enterprise Resource Planning

ERP is an industry term for the broad set of activities supported by multi-module application software that helps manufacturers manage the important parts of their business: product planning, procurement, inventory maintenance, supplier

interaction, customer service, and order tracking. ERP can also include application modules for the finance and human resources aspects of a business. Typically, an ERP system uses or is integrated with a relational database system.

Key server attributes for ERP applications are availability, scalability, and performance. The x450, with its flat memory model, Itanium 2 processors, and Enterprise X-Architecture technology such as Active Memory and XceL4 Server Accelerator Cache, provides a robust base on which to build and implement successful ERP solutions.

Supply chain management

Supply chain management (SCM) is the oversight of materials, information, and finances as they move in a process from supplier to manufacturer to wholesaler to retailer to consumer. Supply chain management involves coordinating and integrating these flows both within and among companies.

The x450 is a preferred platform for 64-bit SCM management applications. The x450 offers a range of leading technologies that will help to deliver the uptime required for business-critical applications at the lowest price/performance ratio. The x450 covers all high-availability features for customers looking for servers to power their SCM solutions.

Customer relationship management

Customer relationship management (CRM) is an information-industry term for methodologies, software, and usually Internet capabilities that help an enterprise manage customer relationships in an organized way.

The x450 provides a performance-based foundation upon which customers can build and deploy CRM solutions in which the x450 will most likely be implemented as an application server and/or a database server.

2.1.3 e-business and security transactions

e-business is the use of Internet technologies to improve and transform key business processes.

This includes Web-enabling core processes to strengthen customer service operations, streamlining supply chains and reaching existing and new customers. In order to achieve these goals, e-business requires a highly scalable, reliable, and secure server platform.

The x450 is a strong candidate for an application integration server that integrates the back-end data with the servers containing end-user or client programs. This involves data transformation, process flow, and other capabilities,

thus allowing companies to integrate applications and other data sources. These types of servers benefit from the processing power offered by the x450. The performance of the server shows promise in Web servers that perform secure e-commerce transactions, which is a market segment currently dominated by Sun Microsystems and their SPARC machines.

To see Intel's performance comparison of Intel Itanium 2 and Sun UltraSPARC III processors, click **Compare processor performance** at:

http://www.intel.com/products/server/processors/server/itanium2/index.htm

When using x450 as a Web server for your e-commerce solutions, you will benefit from its highly parallel computation that can handle higher volumes of secure data transmissions using complex encryption/decryption techniques.

The Itanium 2 micro-architecture is a perfect match for the significant compute power requirements of Secure Sockets Layer (SSL), providing protection without incurring delay. SSL uses advanced public key encryption technology to safely move sensitive information across the Internet.

In today's complex environments, encryption can occur at two levels simultaneously. For example, the IPsec standard performs encryption on every packet sent over the network. These packets frequently contain data that is encrypted itself. As a result, enterprises need platforms that can encrypt and decrypt data very quickly. This is the time to deploy x450 with its 4-way SMP Itanium 2 processors, each with two (rather than one) floating-point units, two memory/ALU, two integer ALU units, and three branch execution units, 128 integer and 128 floating-point registers.

To compare the performance in integer and floating-point operations with other processors on the market, see SPECint_base2000 and SPECfp_base2000 benchmarks at:

http://www.spec.org

2.1.4 In-house developed compute-intensive applications

For developers, the 64-bit architecture allows creation of applications using a familiar programming model that encourages the development of a wide-ranging set of enterprise solutions. The experiences gained on the IA-32 platform can be reused when creating applications for Itanium 2-based systems. Thus developers are not required to start from scratch to make their transition to the 64-bit world almost seamless. In most cases, existing 32-bit code will not require a complete rewrite, but only recompilation.

******DAVID!!!*****The important message for developers is that IA-64 parallelism managed by the compiler itself, application development is uncharged of using

special techniques to employ it. Currently available compilers optimized for Itanium 2 include Intel C++ Compiler, Intel Fortran Compiler, and Microsoft SDK/C/C++ Compiler.

2.1.5 Science and technology industries

Science and technology industries (S&TC) require the processing of large and complex calculations to solve challenging problems.

While S&TC industries are characterized by compute-intensive workloads that require special server characteristics to meet their performance needs, each industry -— petroleum, automotive, aerospace, weather or research — also has its own set of S&TC applications, each demanding different computing solutions.

The Itanium 2-based systems are a solid base for compute-intensive applications. It is for this reason that the Itanium and Itanium 2 processors were chosen for the largest American-built supercomputer, the TeraGrid project sponsored by the National Center for Supercomputing Applications. When finally deployed with more than 3,000 Itanium 2 chips, TeraGrid will be capable of more than 14 billion floating-point operations per second (teraflops).

Itanium 2 features, such as increased system memory bandwidth, total number of 328 onboard registers (including 128 floating-point registers) and thus high floating-point performance, speed up calculations and data analysis in S&TC applications.

2.2 Why choose x450?

Modern enterprise applications require servers with significant processing power that are able to process large quantities of data and store numerous transactions in cache, and do it all with the highest possible availability and reliability. Systems that cannot deliver at this level are destined to serve as solutions in non-critical settings.

IBM, as a market leader, has been delivering top-of-the-line solutions for enterprise environment for years and xSeries 450 is evidence that this is not going to change with future IA-64 technology. With close cooperation between IBM and Intel, the xSeries 450 will continue to be an excellent platform for the most demanding 64-bit enterprise applications.

xSeries 450 combines the best of three worlds:

- The best of the xSeries servers family, the x450 server is based on the mainframe-inspired reliable components of the xSeries 440, IBM's most

scalable IA-32-based server, which has become a well-accepted and tested solution in medium to large enterprises.

- A replacement to the xSeries 380, the first xSeries server with a 64-bit architecture targeted, tested, and proven by software and hardware developers.
- Unlike competitors without such background, IBM put many years of invaluable experience with 64-bit technology (IBM pSeries™ and zSeries™ servers) into the development of the x450.

All of these advantages make the xSeries 450 production enterprise-level server a viable alternative to RISC-based architectures, while protecting many years of investments and knowledge gained from IA-32 platform development.

Intel expects the Itanium Architecture to be a product base for the next 25 years or more.

The Itanium 2 processor used in the x450 is based on the *Explicitly Parallel Instruction Computing* (EPIC) architecture, which incorporates a number of new technologies, features, and capabilities that make it ideal for the high-end server and workstation markets. EPIC allows users to take advantage of its large memory addressability and parallel execution capabilities. The chip also supports intelligent prediction and speculation of events to deduce redundancy and improve performance. The Itanium's floating point engine enhances performance for complex computations that are required for data-mining, scientific, and technical applications.

The Itanium 2 processor is the second in a family of Intel enterprise-class processors. For more information about the Intel Itanium 2 processor, see:

`http://www.intel.com/products/server/processors/server/itanium2/`

Intel Itanium Architecture-based microprocessors have the following features:

- Advanced parallelism

  High performance requires parallel execution, which is either very limited or hardly achieved in today's architecture. The traditional PC systems are not designed for parallelism, which is critical for current demanding applications (for example databases and application servers).

  Today’s processors using limited parallelism are often 60% idle. When source code is compiled on today’s systems, the result of the compilation is sequential machine code. A regular (non-Itanium processor family) compiler takes sequential code and examines and optimizes it for parallelism but then has to regenerate sequential code, but in such a way that the processor can re-extract the parallelization from it. The processor will then be required to read this implied parallelism from the machine code, re-build it, and execute it.

The parallelism is there, but it is not as obvious to the processor and more work has to be done by the hardware before it can be utilized.

Itanium 2 supports parallelism on multiple levels. Instruction-level parallelism (ILP) is the ability to execute multiple bundles (three instructions in a bundle) at the same time. The Itanium 2 micro-architecture can deliver faster performance by executing multiple bundles per clock cycle. Parallelism, both at the instruction level and at the SMP system level, permits more efficient use of virtually all system resources to enable improved scalability.

The Itanium processor's instruction-level parallelism helps ensure the scalability necessary to manage large data warehouses.

Note: Performance should no longer be measured by just the speed in MHz, but also by the degree of parallelism that the processor achieves.

- Large memory addressability

  Another key advantage is that 64-bit operating systems can support far more physical memory than a 32-bit operating system. The theoretical limit for directly addressable memory was 4 GB in 32-bit architectures and is now 2^{64} or 16 Exabytes.

  The Itanium 2 processor memory subsystem has a three-level cache structure consisting of first-level instruction cache, first-level data cache, second-level cache (L2), and third-level cache (L3). In addition, to ensure the processors are optimally used, the x450 has a fourth-level (L4) system cache dedicated to the CPUs (the 64 MB XceL4 Server Accelerator Cache).

  The increased physical memory includes the following benefits for applications:

  – Each application can support more users.

    For a comparison of the number of maximum connected users in SAP for various hardware platforms (including Itanium 2), visit the following link:

    http://www.sap.com/benchmark/

  – Each application has better performance. Increased physical memory allows more applications to run simultaneously and remain completely resident in the system's main memory. This reduces or eliminates the performance penalty of swapping pages to and from disk.

  – Each application has more memory for data storage and manipulation. Databases can store more of their data in the physical memory of the system. Data access is faster because disk reads are not necessary.

  – Applications can manipulate large amounts of data easily and more reliably. Video composition and modeling for scientific and financial

applications benefit greatly from memory-resident data structures that are not possible on 32-bit operating systems.

For the enterprise customer, a larger physical memory subsystem provides access to more data more quickly from system memory, since much more data can be held near the processor for faster calculations and data analysis. Large memory addressability allows larger file system caches for read-ahead and write-behind I/O operations, and also allows retention of large amounts of data in memory instead of repeatedly reading the data from disk. Combined with high memory bandwidth and a variety of performance optimization techniques, this solution provides the performance the enterprise market needs.

- Compatibility

  Another important aspect of Itanium 2 architecture is 32-bit compatibility. Understanding that many applications will never need the processing power of 64-bit computing, the chip will also support legacy software that relies on IA-32 architectures.

  The x86 binaries (applications running on IA-32), identical to those running on Pentium and Xeon processor-based systems, may execute on Itanium 2 without modification. In fact, both x86 and 64-bit binaries can run simultaneously on the same processor. This allows a system that is only partially converted to 64 bits to run on x450. Of course, unless the applications are recompiled, 32-bit applications will run slower on the Itanium because of emulation overhead and an inability to extract all the benefits the 64-bit platform offers.

  xSeries 450 is built to scale, process, and access larger amounts of data to yield significant performance increases of data compute-intensive applications than any other Intel-based server before.

3

Planning

In this chapter we discuss topics you need to consider before you finalize the configuration of your x450 system and before you begin implementing the system. The topics covered are:

- System hardware
- Cabling and connectivity
- Storage considerations
- Rack installation
- Power considerations
- Solution Assurance Review

3.1 System hardware

The x450 provides a scalable and flexible hardware platform. There are a number of important aspects of the system hardware to consider when planning your configuration. These are discussed in this section.

3.1.1 Processors

xSeries 450 servers combines copper-based, XA-64 Enterprise X-Architecture technologies with 64-bit, Intel Itanium 2 processors.

Table 3-1 lists the processors standard in each x450 model and the part number of additional processors.

Table 3-1 Processors in each x450 model

Model	Standard CPUs	L2 cache	L3 cache	Max SMP	Extra CPUs
8688-4RX	One 1.3 GHz	256 KB	3 MB	Four-way	73P7076
8688-5RX	Two 1.4 GHz	256 KB	4 MB	Four-way	73P7077
8688-6RX	Two 1.5 GHz	256 KB	6 MB	Four-way	73P7078

The L2 and L3 cache run at the full speed of the processor.

A total of four processors can be installed in the x450. One, two, three, or four processors can be installed. Each processor within a node must be the same speed and cache size. Each processor option includes the processor with heatsink (pre-assembled) and its associated power module (which is referred to as the *power pod*).

Processors must be installed in a specific order, as shown in Figure 3-1 on page 41. Special tools are required to install the processors and these are included in the option.

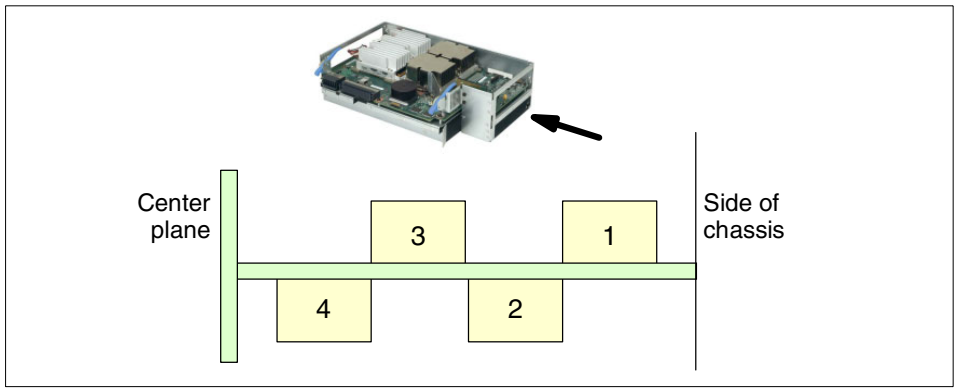

Figure 3-1 Processor installation order

3.1.2 Memory

A maximum of 28 DIMMs may be installed in the memory-board assembly of the x450, depending on the size of the DIMMs used. Supported DIMMs are listed in Table 3-2.

Table 3-2 Supported DIMMs

Size	Description	Part number	Max installable
512 MB	PC2100 CL2.5 ECC DDR SDRAM	73P2027	28
1 GB	PC2100 CL2.5 ECC DDR SDRAM	73P2028	28
2 GB	PC2100 CL2.5 ECC DDR SDRAM	73P2030	20

Go to the ServerProven® site for the latest information on supported memory modules:

http://www.ibm.com/pc/us/compat

When installing memory DIMMs, they must be installed in matched pairs (size and technology), in the order shown in Table 3-3.

Table 3-3 DIMM order

Pair	Port 1 DIMM slots	Pair	Port 2 DIMM slots
1	1 and 14	2	15 and 28
3	4 and 11	4	18 and 25
5	2 and 13	6	16 and 27

Pair	Port 1 DIMM slots	Pair	Port 2 DIMM slots
7	5 and 10	8	19 and 24
9	3 and 12	10	17 and 26
11	6 and 9	12	20 and 23
13	7 and 8	14	21 and 22

Restriction: When 2 GB DIMMs are used, a maximum of 20 DIMMs can be installed. In such a configuration, the last four pairs of sockets as shown in Table 3-3 must be left empty (that is, sockets 6, 7, 8, 9, and 20, 21, 22, 23).

Figure 3-2 shows the DIMM locations and the way the DIMMs are divided into the two memory ports.

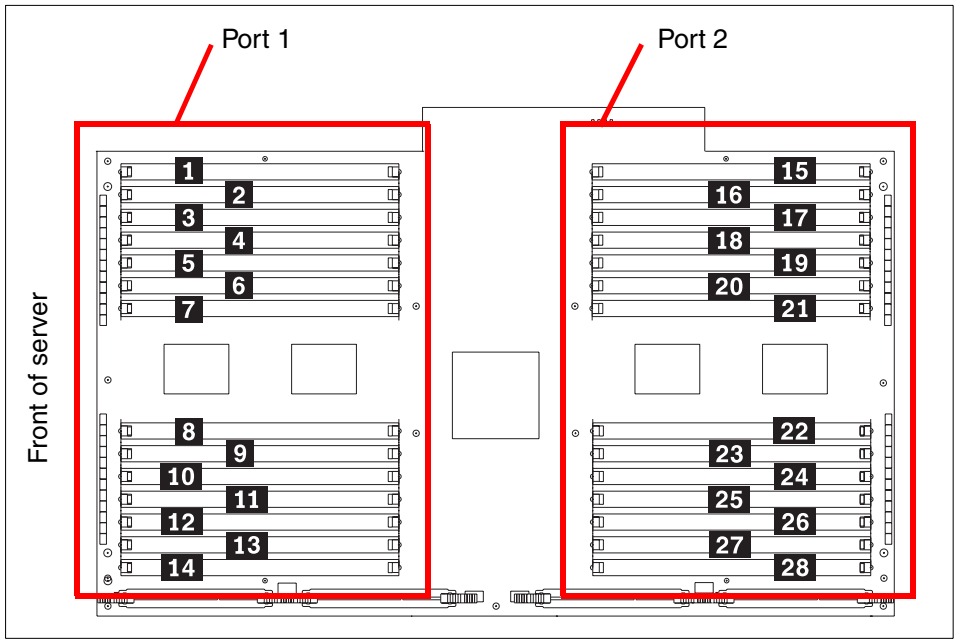

Figure 3-2 DIMM locations on the memory-board assembly

Memory mirroring

Memory mirroring is supported by the x450 for increased fault tolerance and high levels of availability.

Key configuration rules relating to memory mirroring:

- Memory mirroring must be enabled in the System Setup (it is disabled by default).
- Enabling memory mirroring halves the amount of memory available to the operating system.
- You must install two pairs of DIMMs at a time. These four DIMMs (known as a *bank*) must be identical. Table 3-4 shows the pairs that are in each bank.

Table 3-4 DIMMs that form a bank

Bank	DIMM pairs (see Table 3-3)
1	1 and 2 (DIMMs 1, 14, 15, 28)
2	3 and 4 (DIMMs 4, 11, 18, 25)
3	5 and 6 (DIMMs 2, 13, 16, 27)
4	7 and 8 (DIMMs 5, 10, 19, 24)
5	9 and 10 (DIMMs 3, 12, 17, 26)
6	11 and 12 (DIMMs 6, 9, 20, 23)
7	13 and 14 (DIMMs 7, 8, 21, 22)

Note: The ability to hot-replace a failed DIMM or hot-add additional DIMMs is currently not supported.

Memory performance considerations

As shown in the server block diagram in Figure 1-1 on page 4, there are two memory ports to the memory controller, each with a throughput of up to 3.2 GBps. These ports correspond to the ports as shown in Figure 3-2 on page 42. The front side bus of the processors is 6.4 GBps, so maximum performance is achieved when both memory ports are used to access memory simultaneously.

Consequently, for maximum performance, you must install four DIMMs of the same size at a time into a bank (see Table 3-4). In this configuration, all memory addresses are spread across all four DIMMs in the bank and, when accessed, both memory ports are used, resulting in 6.4 GBps burst transfer rates.

Maximum performance can also be achieved with DIMMs of different sizes, as long as the total memory in port 1 matches the total memory in port 2. For example, if you have 6 DIMMs (four 512 MB DIMMs and two 1 GB DIMMs for a total of 4 GB), install all four 512 MB DIMMs (2 GB) in port 1 and the two 1 GB DIMMs (also 2 GB) in port 2.

If there is a mismatch between the total memory in port 1 and the total memory in port 2, then there will be a delay when accessing the upper memory. For example if you have two 512 MB DIMMs and two 1 GB DIMMs, memory accesses to the first 2 GB of memory (1 GB from each port) will be at maximum rate of 6.4 GBps. However, the remaining 1 GB (the upper half of the 1 GB DIMMs) will only be at 3.2 GBps maximum.

3.1.3 PCI slot configuration

As shown in Figure 1-1 on page 4, there are six PCI-X slots internal to the x450. These six slots are implemented using four PCI buses:

- Bus A (slot 1 and slot 2): Two 64-bit 66 MHz slots
- Bus B (slot 3 and slot 4): Two 64-bit 100 MHz slots (133 MHz if only one slot is occupied)
- Bus C (slot 5): One 64-bit 133 MHz slot
- Bus D (slot 6): One 64-bit 133 MHz slot

These slots can accept adapters rated at speeds ranging from 33 MHz to 133 MHz.

You should also consider the following:

- Each adapter has a maximum rated speed. Each bus also has a maximum rated speed.
- Installed adapters in a single bus will operate at the slowest of three speeds:
  - The rated speed of adapter 1
  - The rated speed of adapter 2 (if the bus the adapter is installed in has two slots)
  - The rated speed of the bus
- Bus B supports one adapter at up to 133 MHz or two adapters at up to 100 MHz.
- 32-bit adapters can be installed in any of the slots and will run in 32-bit mode. 32-bit and 64-bit adapters can coexist in 64-bit slots in the same bus. The 32-bit adapters will run in 32-bit mode, and the 64-bit adapters will run in 64-bit mode.

Tip: Take the time to understand these rules and to select the best slots for your adapters. Incorrect choices can result in a loss of PCI adapter performance.

As extreme configuration examples, you could configure either of the following:

- Six 33 MHz PCI adapters, all operating at 33 MHz.
- Six 133 MHz PCI-X adapters, with two operating at 133 MHz (buses C and D), two at 100 MHz (bus B) and two at 66 MHz (bus A).

Important: A PCI-X and a PCI adapter can be installed in slots on the same bus. However, those two adapters will both operate in PCI mode.

In addition, if you have a PCI-X adapter installed, you cannot hot-add a PCI adapter to the same bus. This is because with just the PCI-X adapter installed, the bus is running in PCI-X mode, and you cannot hot-add a PCI adapter into a bus that is in PCI-X mode.

Table 3-5 summarizes the supported adapter speeds. Take into account the speed reductions when there are two adapters installed in a bus, as described above.

Table 3-5 Supported adapter speeds in each slot

Slot	Bus	Width (bits)	Supported adapter speed (MHz)
1	A	32 or 64	33 or 66
2	A	32 or 64	33 or 66
3	B	32 or 64	33, 66, or 100 (133 as long as no adapter is in slot 4)
4	B	32 or 64	33, 66, or 100 (133 as long as no adapter is in slot 3)
5	C	32 or 64	33, 66, 100 or 133
6	D	32 or 64	33, 66, 100 or 133

The physical location of these slots in the server is shown in Figure 3-3 on page 46.

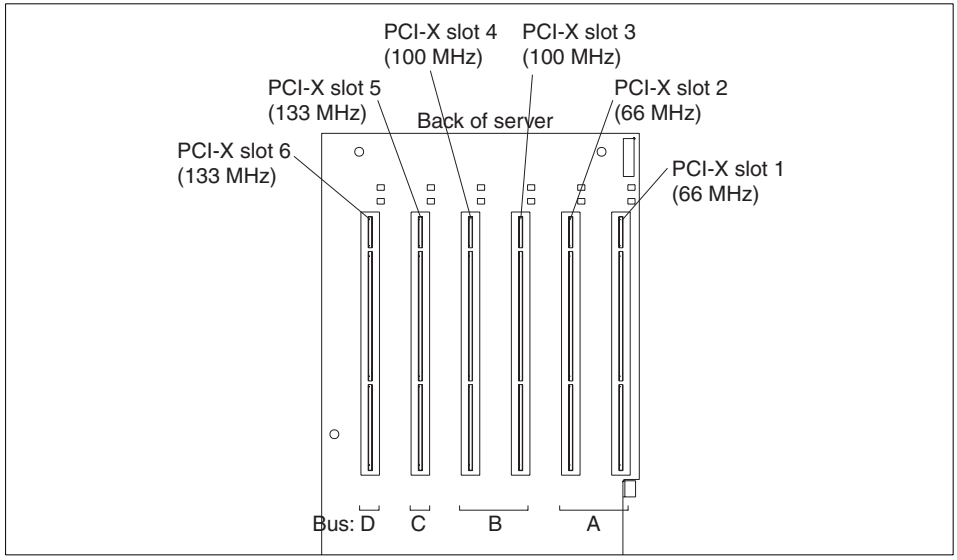

Figure 3-3 PCI-X slots in the x450

Other configuration information:

- The x450 server supports connection to a RXE-100.
- Video adapters are not supported.
- The PCI slots support 3.3 V adapters only.

> **Important:** 5.0 V adapters are not supported.

- Only the on-board SCSI/RAID controller is supported for connection to the internal drive bays. The use of ServeRAID or Fibre Channel adapters is only supported for connection to external disks.
- Do not install a ServeRAID card in slot 1. This is because there is little space between the top of the adapter and the cover when the covers are closed. This could damage the SCSI cable.
- Some long adapters have extension handles or brackets installed. Before installing the adapter, you must remove the extension handle or bracket.
- The system scans PCI-X slots to assign system resources. The system attempts to start the first device found. The search order is:

  a. CD-ROM
  b. Disk drives
  c. Integrated SCSI devices
  d. x450 PCI-X slots (in the order 1, 2, 6, 5, 3, 4)

e. Integrated Ethernet controller

If an RXE-100 is attached, the order is:

a. CD-ROM
b. Disk drives
c. Integrated SCSI devices
d. x450 PCI-X slots (1, 2, 6, 5, 3, 4)
e. RXE-100 slots (A5, A6, A3, A4, A1, A2, B6, B5, B3, B4, B1, B2)
f. Integrated Ethernet controller

3.1.4 Broadcom Gigabit Ethernet controller

The x450 offers a dual Gigabit Ethernet controller integrated standard in the system. The x450 includes a single-port Broadcom BCM5704 10/100/1000 BASE-T controller on a PCI 64-bit 66 MHz bus.

The BCM5704 supports full and half-duplex performance at all speeds (10/100/1000 Mbps, auto-negotiated) and includes integrated on-chip memory for buffering data transmissions, and dual onboard RISC processors for advanced packet parsing and backwards compatibility with 10/100 devices. The Broadcom controller also includes software support for Wake on LAN, failover, layer-3 load balancing, and comprehensive diagnostics.

Category 5 or better Ethernet cabling is required with RJ-45 connectors. If you plan to implement a Gigabit Ethernet connection, ensure your network infrastructure is capable of the necessary throughput to match the server's I/O capacity.

Adapter teaming

The Broadcom controller is capable of participating in an adapter team for the purposes of failover, load balancing, and port trunking. The choice of adapters to team with the onboard controller depends on whether you have a copper-only network or a mixed copper/fiber network. Our recommendations are:

- If you have a copper Gigabit environment, use the Broadcom-based NetXtreme 1000T Ethernet adapter, part 31P6301. Alternatively, use the Intel PRO/1000 XT Server adapter, part 22P6801. However, this may only be supported in specific slots. See the following for details:

  http://www.pc.ibm.com/us/compat/x450/ibm_22P6801.html

- If you have a mixed fiber/copper Gigabit server switch network, use the Broadcom-based 22P7801, NetXtreme 1000 SX Fiber Ethernet adapter.

You can also team any of the onboard Gigabit cards with 10/100 cards such as 06P3601 and 22P4901, but this is not a recommended configuration. You can also team with the older Gigabit fiber card, 06P3701.

Adapter teaming and failover works by using software additional to the adapter driver to provide the failover functionality. This software is operating system dependent. Detailed instructions for installing the individual driver and failover packages are available with the driver software.

For the latest network adapter drivers and software for the x450 server, go to the xSeries support page:

http://www.pc.ibm.com/support

For details about compatibility, see the ServerProven LAN adapter page:

http://www.pc.ibm.com/us/compat/lan/matrix.html

3.2 Cabling and connectivity

There are a number of unique factors to consider when cabling the x450 server:

- Remote Supervisor Adapter connectivity
- RXE-100 connectivity
- Serial connectivity

We discuss each of these in this section.

The rear panel of the x450 showing the locations of cable connectors is illustrated in Figure 3-4 on page 49. For more details about ports on the Remote Supervisor Adapter, refer to 3.2.1, “Remote Supervisor Adapter connectivity” on page 50.

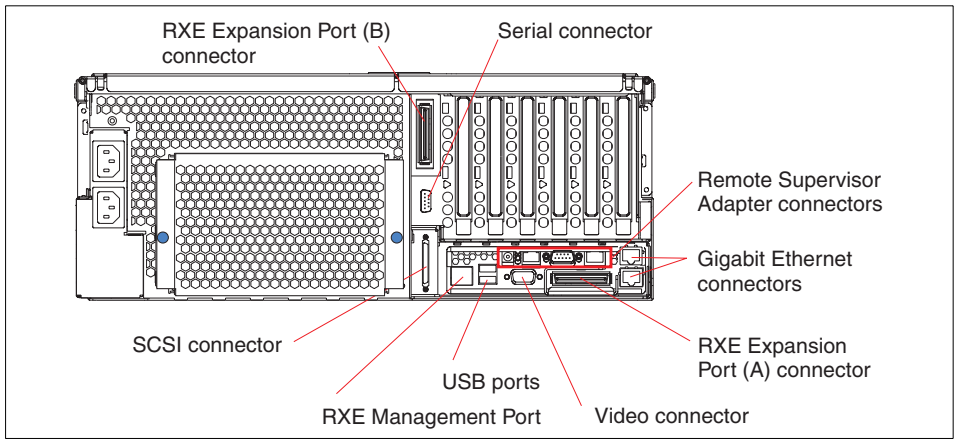

Figure 3-4 x450 rear view

Of note are the following items:

- There are no PS/2 keyboard or mouse ports on the x450. Only USB devices are supported.

  For attachment to an Advanced Connectivity Technology (ACT) KVM switch, a new USB Conversion Option (part number 73P5832) can be used. This smart cable is plugged into the USB and video ports on the server and it converts KVM signals to CAT5 signals for transmission over a CAT5 cable to either a Remote Console Manager (RCM) or Local Console Manager (LCM). The x450 can then be managed on the same set of switches as PS2- or C2T-based KVM servers.

- Unlike the x440, the x450 includes a dedicated external serial port. During the boot process, the port acts as an auxiliary console for the EFI, as described in 1.3, “Extensible Firmware Interface” on page 5. After the operating system has booted, the port is dedicated to the operating system as COM1.

- The RXE Expansion Ports provide connectivity to a RXE-100. Either a single cable in port A or two cables to both ports are supported. Using two cables provides redundancy as well as additional throughput.

- The external SCSI connector is for tape drive connectivity only.

- There are two RJ-45 ports providing Gigabit Ethernet connectivity, as shown in Figure 3-4. A third RJ-45 Ethernet connector is also on the Remote Supervisor Adapter. This connector is only used to connect to the Remote Supervisor Adapter for out-of-band management, as described in 5.1, “The Remote Supervisor Adapter” on page 118.

3.2.1 Remote Supervisor Adapter connectivity

The x450 features an integrated Remote Supervisor Adapter, one of the products in the Advanced System Management (ASM) family. It provides around-the-clock remote access and system management of your server and supports the following features:

- Remote management regardless of the status of the server
- Remote control of hardware and operating systems
- Web-based management with standard Web browsers (no other software is required)
- Text-based user interface terminal access

The configuration and use of the Remote Supervisor Adapter is discussed in Chapter 5, “Management” on page 117.

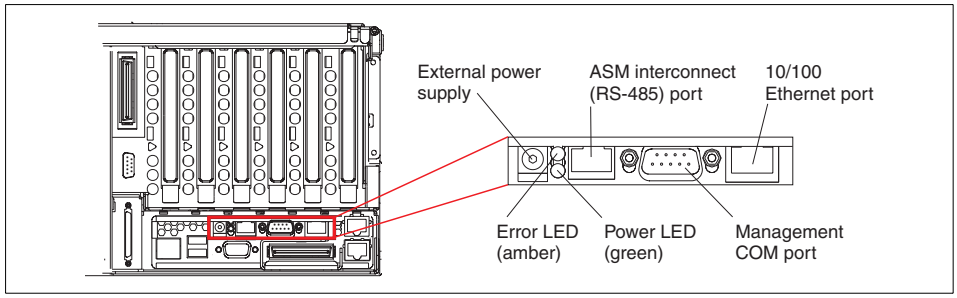

Figure 3-5 Remote Supervisor Adapter connectors

The following RSA connections need to be considered when cabling the x450 (see Figure 3-5):

- External power supply connector. This connector allows the RSA to be connected to its own independent power source. This external power supply is not included with the x450 and will need to be ordered as an option (order a ThinkPad® 56W AC Adapter with a suitable power cord for your country/region).

  If this power supply is not used, the RSA will draw power from the server as long as the server is connected to a functioning power source.

  Important: An external power supply is not officially supported in the x450.

- 9-pin serial port, which supports systems management functions through null modem or modem connections. This port is dedicated and can only be used for RSA purposes.

- Ethernet port, which provides system management functions over the LAN.
- Advanced Systems Management (ASM) RS-485 interconnect port to facilitate advanced systems management connections to other servers.

  For detailed instructions on cabling ASM interconnect networks, refer to the redbook *Implementing Systems Management Solutions using IBM Director*, SG24-6188.

  Note: The x450 does not include the necessary dongle to connect the Remote Supervisor Adapter to an ASM interconnect bus using the RS-485 port on the adapter. Consequently, you will need the Advanced System Management Interconnect Cable Kit (part number 03K9309) for connection to an ASM interconnect network.

3.2.2 Remote Expansion Enclosure

The RXE-100 can be connected to the x450 to provide an additional six or 12 PCI-X slots to the server. Only one RXE-100 can be connected to the x450, although this can be with either one or two data cables.

The RXE-100 has six 133 MHz 64-bit PCI-X slots as standard and can accept adapters with speeds ranging from 33 MHz to 133 MHz. With the optional six-slot expansion kit (part number 31P5998) installed, the RXE-100 has 12 slots. Each set of six adapter slots is divided into three buses of two slots each, as shown in Figure 3-6 on page 52.

Note: As described in “Connecting the RXE-100” on page 53, when connecting the RXE-100 to an x450 configuration using only one cable, the RXE-100 can have six or 12 PCI-X slots.

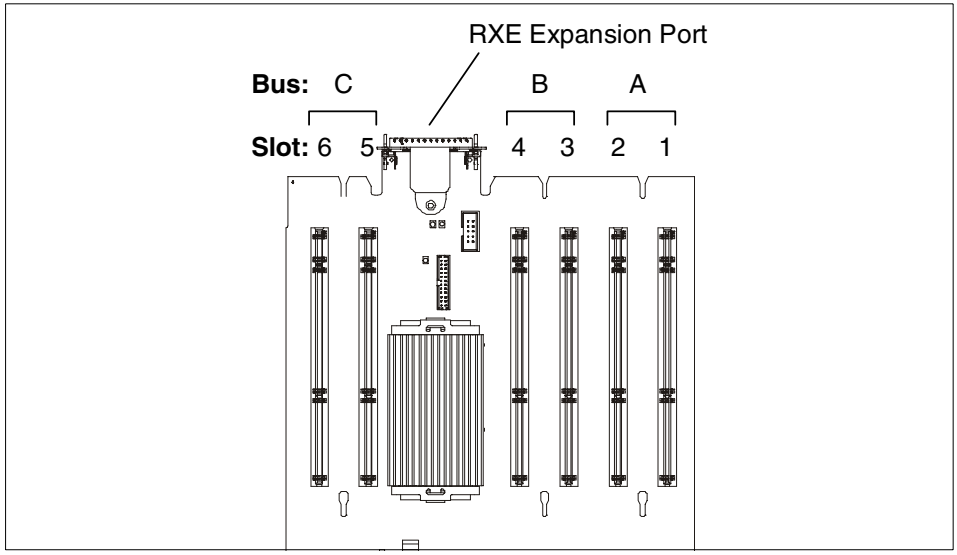

Figure 3-6 RXE-100 PCI-X expansion board (six slots)

For each of the three buses (A, B, C), one of the following can be installed:

- One 64-bit 3.3 V PCI-X 133 MHz adapter (in the odd-numbered slot), running at up to 133 MHz
- Two 64-bit 3.3 V PCI-X 133 MHz adapters running at up to 100 MHz
- Two 64-bit 3.3 V PCI or PCI-X, 33 or 66 MHz adapters

Note: The PCI slots support 3.3 V adapters only. 5.0 V adapters are not supported.

Like the x450, these slots can accept adapters rated at speeds ranging from 33 MHz to 133 MHz. When deciding which adapters to put in which slots, consider the following:

- Each adapter has a maximum rated speed. Each bus also has a maximum rated speed.
- Installed adapters will operate at the slowest of three speeds:
  - The rated speed of adapter 1 in the bus
  - The rated speed of adapter 2 in the bus
  - The rated speed of the bus
- 32-bit adapters can be installed in any of the slots and will run in 32-bit mode. 32-bit and 64-bit adapters can coexist in 64-bit slots in the same bus. The 32-bit adapters will run in 32-bit mode, and the 64-bit adapters will run in 64-bit mode.

- When installing a 133 MHz PCI-X adapter, it must be installed in the first or odd-numbered slot in the bus (that is in slots 1, 3 or 5).
- Like the x450, a PCI-X and a PCI adapter can be installed in slots on the same bus in the RXE-100. However, these two adapters will both operate in PCI mode.

  In addition, if you have a PCI-X adapter installed, you cannot hot-add a PCI adapter to the same bus. This is because with just the PCI-X adapter installed, the bus is running in PCI-X mode, and you cannot hot-add a PCI adapter into a bus that is in PCI-X mode.

Connecting the RXE-100

There are two types of cables used to connect the RXE-100 to the x450:

- Remote I/O cable, for data. Two lengths are available:
  - 3.5 m Remote I/O cable kit (part number 31P6102), like the one shipped with the RXE-100
  - 8 m Remote I/O cable kit (part number 31P6103)
- Interconnect management cable, for remote I/O management. This is a standard Ethernet cable. The RXE-100 ships with a 3.5 m cable.

  Two lengths are available:

  - 3.5 m interconnect management cable kit (part number 31P6087)
  - 8 m interconnect management cable kit (part number 31P6088)

  Use the 8 m versions of each of these cables if the distances between the two devices warrant the extra length (for example, in separate racks).

  When the RXE-100 has only six PCI-X slots installed, connect a single Remote I/O cable and a single management cable, as shown in Figure 3-7 on page 54.

Important: Power to the RXE-100 is controlled by the x450 via the interconnect management cable and under the control of the Remote Supervisor Adapter.

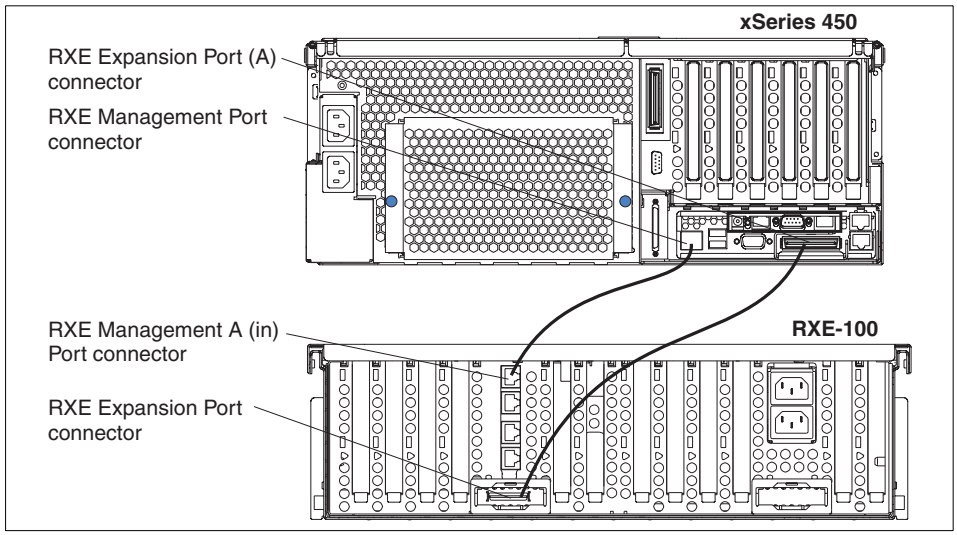

Figure 3-7 Connecting the RXE-100 to the x450 (six slots in the RXE-100)

If the RXE-100 has the second set of six PCI slots installed, connect the two devices as shown in Figure 3-8.

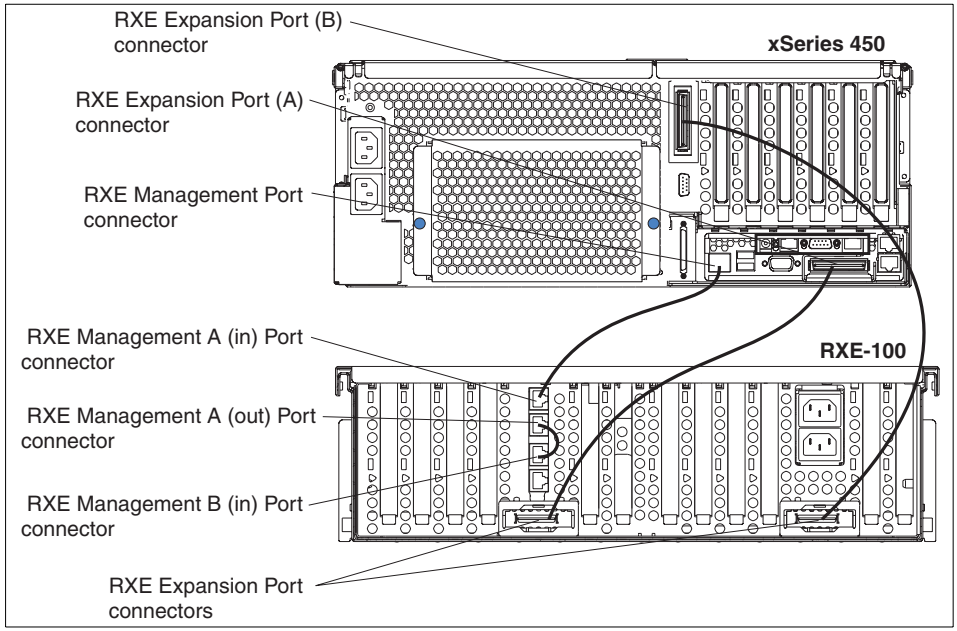

Figure 3-8 Connecting the RXE-100 to the x450 (12 slots in the RXE-100)

The short interconnect management cable to connect Management A (out) Port to Management B (in) Port is supplied with the second set of PCI slots.

The second Remote I/O data cable is optional. However, it is recommended because it provides redundancy if the other Remote I/O cable fails, and it also increases throughput from the PCI adapters in the RXE-100 to the x450.

Tip: As described in Chapter 4, “Installation” on page 67, you should disconnect the RXE-100 before installing an operating system.

3.2.3 Serial connections

The x450 has an integrated serial port as shown in Figure 3-4 on page 49. This port has two purposes:

- During the boot process (before the OS loader starts), the Extensible Firmware Interface (EFI) has control of the port and uses it as an auxiliary console where POST messages are transmitted, many even before the server’s video port is enabled. This is especially useful in performing problem determination on the system.
- Once the operating system loads, it is made available to the operating system as a standard (and dedicated) COM port.

To use the serial port, for use as an auxiliary console, you will need the following:

- A null modem cable
- A system running a terminal emulation program, such as HyperTerminal in Windows or minicom in Linux

First, connect the RS-232 cable between the two systems. The cable should be connected to the serial port of the server, not the serial port of the Remote Supervisor Adapter.

Once the cable is connected, start HyperTerminal or your emulation program. HyperTerminal should be set to the following settings:

Speed: 115200 bps
Data bits: 8
Parity: None
Stop Bits: 1
Flow Control: None

For HyperTerminal, you can leave emulation set to auto detect. If you are using another program that does not have auto detect, you may need to set emulation to ANSIW.

Once your session is configured, click **Connect** and, assuming there is no operating system currently running, you should see POST messages or the EFI Boot Manager menu.

Tip: If you are running Linux, you can redirect the console messages to the serial port. For more information about this feature, refer to the “Console messages over the serial port” section in Chapter 4 of the publication *Linux on Itanium 2*.

3.3 Storage considerations

When you are planning the storage configuration to accompany the x450, there are important performance and sizing issues that need to be considered.

The two internal hot-swap 1” drive bays will typically be used for operating system installation. We recommend these drives be configured as a two-drive RAID-1 array to provide a higher degree of system availability. Drives up to 15,000 RPM and the converged tray design are supported. To configure RAID-1, you must use the internal LSI controller that is provided onboard.

Restriction: ServeRAID adapter is supported on external disks only. For internal disks, you must use the internal LSI controller.

Typically the x450 will be attached to an external disk enclosure for data storage requirements. Some of the supported IBM storage options are:

- SCSI RAID adapters and storage enclosures
- Fibre Channel adapters and Storage Area Networks (SANs)
- Network Attached Storage (NAS)
- SCSI over IP (iSCSI)
- IBM Enterprise Storage Server™ (ESS)
- ESCON® connectivity to a zSeries server

Note: You can also add as an option an LS-120 diskette drive that can be used to upgrade the firmware or the EFI version. Please note that no diskette drive is provided with the default configuration, and that this solution is not supported.

3.3.1 xSeries storage solutions

This section discusses some of the available xSeries storage solutions and related technologies, as well as tape backup and performance considerations.

This section discusses some of the available xSeries storage solutions and related technologies, as well as tape backup and performance considerations.

Note: All the IBM Storage Solutions are supported under Linux, as well as under Windows. For more information see:

http://www.storage.ibm.com/linux

ServeRAID with external storage enclosures

The current ServeRAID-4 family of adapters includes the ServeRAID 4Mx and 4Lx. These 64-bit, Active PCI controllers include advanced features such as Logical Drive Migration, nine RAID levels including RAID 1E, 1E0 and 5E, as well as adapter and cluster failover.

Restriction: At the time of writing, the ServeRAID adapters were not supported due to the unavailability of a 64-bit version of the ServeRAID configuration CD-ROM.

- ServeRAID-4Mx features two Ultra160 SCSI channels, 64 MB of battery-backed ECC cache memory, and an Intel i80303 processor. Up to 28 Ultra160 and Ultra2 SCSI devices are supported.
- ServeRAID-4Lx features one Ultra160 SCSI channel, 32 MB of ECC cache memory, and an Intel i80303 processor. Up to 14 Ultra160 and Ultra2 SCSI devices are supported.

  Each ServeRAID adapter supports up to 14 drives (and 160 MBps throughput) per channel (for an aggregate of up to 56 drives, for example). Multiple adapters can be installed as needs and available slots dictate.

  Note: ServeRAID 4H is not a supported option with the x450.

- The EXP300 storage expansion unit has a maximum 1 TB of disk storage (14 146.8 GB drives) in a 3U package, allowing up to 14 expansion units to be used in a standard 42U rack (meaning that a full rack of EXP300 units can hold an amazing 14 TB). The EXP300 provides Predictive Failure Analysis (PFA) on key components, including hot-swap fans, hard drives, and redundant power supplies. The EXP300 is optimized for Ultra160 SCSI, with a sustained data transfer rate of 160 MBps.

For more information on IBM SCSI RAID storage solutions, go to:

http://ibm.com/pc/ww/eserver/xseries/scsi_raid.html

Note: ServeRAID will be supported only on external disks on the x450 server. Booting any operating system from the ServeRAID card will not be supported.

IBM Fibre Array Storage Technology

The IBM Fibre Array Storage Technology (FAStT) family of Fibre Channel storage solutions is designed for high-availability, high-capacity requirements. FAStT solutions can support transfers over distances up to 10 km (6.2 miles) at rates of up to 200 MBps.

The FAStT Storage Server is a RAID controller device that contains Fibre Channel (FC) interfaces to connect the host systems and the disk drive enclosures. The Storage Server provides high system availability through the use of hot-swappable and redundant components. We briefly discuss the following three products:

- The IBM TotalStorage™ FAStT200 Storage Server
- The IBM TotalStorage FAStT500 Storage Server
- The IBM TotalStorage FAStT700 Storage Server

Note: The FAStT700 is currently the only certified storage solution for Microsoft clustered configurations. To check the Microsoft Hardware Compatibility List (HCL) for updates to certified solutions, refer to `http://www.microsoft.com/hcl`.

The IBM TotalStorage FAStT200 Storage Server

The FAStT200 Storage Server is a 3U rack-mountable Fibre Channel RAID controller and disk drive enclosure. It targets the entry and midrange segment of the FC storage market. A typical use of the FAStT200 would be in a two-node cluster environment with up to 66 Fibre Channel disk drives attached to the Storage Server.

Two models are available:

- The FAStT200 Storage Server, with a single RAID controller.
- The FAStT200 High Availability (HA) Storage Server, which contains two RAID controllers and can therefore provide higher availability.

Both models feature hot-swap and redundant power supplies and fans and you can install up to 10 slim-line or half-high FC disk drives. If you need to connect more than 10 disks, you can use the EXP500 FC storage expansion enclosures.

Each EXP500 can accommodate 10 additional disk drives, and up to five EXP500s are supported on the FAStT200. This means that the maximum supported number of disk drives is 60.

Hot-swappable and redundant components provide high availability for the FAStT200 Storage Server. A fan or a power supply failure will not cause downtime and such faults can be fixed while the system remains operational. The same is true for a disk drive failure if fault-tolerant RAID levels are used. With two RAID controller units and proper cabling, a RAID controller or path failure will not cause loss of access to data.

Each RAID controller has one host and one drive FC connection. The FAStT200 HA model can use the two host and drive connections to provide redundant connection to the host adapters and to EXP500 enclosures. Each RAID controller unit also contains 128 MB of battery-backup cache.

Tip: The FAStT200 ships with IBM FAStT Storage Manager 7.10. When we wrote this redbook, there was no Storage Manager version for IA-64 architecture. When available, you can download the latest version for IA-64 from http://www.pc.ibm.com/support.

The IBM TotalStorage FAStT500 Storage Server

The FAStT500 Storage Server is a 4U rack-mountable Fibre Channel RAID controller device. It provides the levels of performance, availability, and expandability needed to satisfy high-end storage requirements. You would typically use the FAStT500 Storage Server in advanced cluster environments and possibly with heterogeneous operating systems running on the host systems. Another application would be where multiple servers are being consolidated onto one or more x450 systems and there is a requirement to centralize storage for these systems.

The FAStT500 Storage Server features two RAID controller units, redundant power supplies, and fans. All these components are hot-swappable, which ensures excellent system availability. You use the EXP500 external storage expansion enclosures to install the FC disk drives and you can connect up to 22 EXP500 enclosures to the FAStT500. This means a total of up to 220 disk drives. The enclosures can be connected in a fully redundant manner, which provides a very high level of availability. On the host-side FC connections, you can use up to four mini-hubs.

This allows you to establish up to eight host connections without needing an external hub or a switch. For performance and availability, each RAID controller unit contains 256 MB of battery-backed cache and this amount can be further expanded.

The IBM TotalStorage FAStT700 Storage Server

The FAStT700 Storage Server is the newest addition to the FAStT range of products. As with the FAStT 500 Storage Server, you would typically implement

the FAStT 700 Storage Server in high-end cluster and server consolidation environments and where multiple servers are being consolidated onto a smaller number of x450 systems.

It is the same physical size as the FAStT500 with new higher performance controllers. These new controllers are 2 Gbps and connect via mini-hubs to the new FAStT FC-2 Host Bus Adapter (HBA) and the new 2109 F16 Fibre Channel switch to give full 2 Gbps fabric.

Like the FAStT500, the FAStT700 attaches to up to 220 FC disks via 22 EXP500 expansion units or up to 224 FC disks via 16 EXP700 expansion units to provide scalability for easy growth (18 GB up to 32.6 TB using 146.8 GB drives). To avoid single points of failure, it also features dual hot-swappable RAID controllers, dual redundant FC disk loops, write cache mirroring, redundant hot-swappable power supplies, fans, and dual AC line cords.

FAStT Storage Manager Version 8.21 supports FlashCopy®, Dynamic Volume Expansion, and Remote Mirroring with controller-based support for up to 64 storage partitions. RAID levels 0,1, 3, 5, and 10 are supported and for performance it includes a total of 2 GB battery-backed cache (1 GB per controller).

Restriction: At the time of writing, there was no 64-bit version of FAStT Storage Manager. Consequently, all configuration and management must be performed remotely.

Additional information on the entire range of FAStT storage solutions can be found at:

http://www.storage.ibm.com/hardsoft/disk/fastt/index.html

Enterprise Storage Server (ESS)

ESS provides integrated caching and RAID support for the attached disk devices. ESS can be configured in a variety of ways to provide scalability in capacity and performance. One ESS can support in excess of 28 TB and can utilize 2 Gbps Fibre Channel connectivity.

Redundancy within ESS provides continuous availability. It is packaged in one or more enclosures, each with dual line cords and redundant power. The redundant power system allows ESS to continue normal operation when one of the line cords is deactivated.

ESS provides an image of a set of logical disk devices to attached servers. The logical devices are configured to emulate disk device types that are compatible with the attached servers. The logical devices access a logical volume that is

implemented using multiple disk drives. This allows ESS to connect to all IBM servers, from zSeries to iSeries™, pSeries and xSeries, directly or through a SAN, thus helping the x450 fit into a heterogeneous environment containing a variety of server architectures. ESS offers several choices of host I/O interface attachment methods, including SCSI and Fibre Channel for xSeries.

For more information on the ESS, go to:

http://www.storage.ibm.com/hardsoft/products/ess/index.html

3.3.2 Tape backup

As with your disk subsystem, you need to carefully analyze backup requirements before a tape solution is selected. Considerations when selecting a backup solution should include:

- Currently implemented backup solutions

  If you are consolidating a number of servers onto a single x440 solution, for example, you may want to take the opportunity to move from differing and distributed tape technologies (such as DDS and DLT) and consolidate those into a single, high-performance, automated solution. An example is the IBM Ultrium Autoloader.

- Current and projected capacity requirements

  Select a solution that has the ability to scale as capacity requirements increase.

- Performance requirements

  You need to consider the backup window available, as well as the amount of data being backed up when determining what your backup performance requirements will be. It is also important to consider the need for quick access to data committed to tape when selecting a solution.

- Connection requirements

  Will the tape solution be connected to an existing SAN fabric and if so, will this require additional fabric hardware?

- Hardware and software compatibility

  If you implement a new tape solution, you need to ensure that current backup and management software is still suitable. IBM Tivoli Storage Manager has plans to release 64-bit versions of their Linux and Windows clients, as do other software vendors.

  Disaster recovery procedures may also need to be revised.

IBM offers a full range of high-performance, high-capacity and automated tape solutions for xSeries servers. For detailed information on these products, go to:

http://ibm.com/pc/ww/eserver/xseries/tape.html

Note: The x450 and RXE-100 support 3.3 Volt PCI adapters only. Make sure any SCSI adapters you use to connect your tape subsystem are 3.3 V or dual voltage adapters.

The following Redbooks discuss IBM tape solutions in greater detail:

- *IBM Tape Solutions for Storage Area Networks and FICON™*, SG24-5474
- *Netfinity Tape Solutions*, SG24-5218
- *The IBM LTO Ultrium Tape Libraries Guide*, SG24-5946

3.4 Rack installation

The x450 is 4U high and is intended for use as a rack-drawer server. Due to power distribution considerations, it is recommended that no more than eight 4U x450 chassis be installed in a single 42U rack, leaving 10U available for RXE-100 Remote I/O enclosures, disk or tape storage, or other devices.

The x450 is 27.5 inches deep, and is designed to be installed in a 19-inch rack cabinet designed for 28-inch-deep devices, such as the NetBAY42 ER, NetBAY42 SR, or NetBAY25 SR. Although the x450 system is rack optimized, it may be converted into a tower by installing it in a NetBAY11 SR Standard Rack Cabinet. The NetBAY11 rack supports shipment of fully configured xSeries 450 and other rack-optimized xSeries servers.

Installation considerations include the following:

- The system is not designed to run vertically, and therefore must always be run in a horizontal position.
- For thermal considerations, the x450 must be installed with perforated doors on both front and back. Do not install the x450 in a rack with a glass front door.
- Although installation is supported in non-enterprise racks, it is not recommended, since cable management then becomes an issue.
- The maximum weight of the system, depending on your configuration, is 50 kg (110 lb.). Therefore, this system requires two people to install it in a rack.

If you use a non-IBM rack, the cabinet must meet the EIA-310-D standards with a depth of at least 28 inches. Also, adequate space (approximately two inches for the front bezel and one inch for air flow) must be maintained from the slide

assembly to the front door of the rack cabinet to allow sufficient space for the door to close and provide adequate air flow.

Make sure all the cables attached to the x450 are long enough to permit the server to be slid out of the rack. This would include the normal cables such as power, network, and fiber cables, but also includes the Remote I/O cable for connecting to the RXE-100. See 3.2.2, “Remote Expansion Enclosure” on page 51 for RXE-100 cabling information.

Since the x450 is rack optimized, the IBM xSeries rack configurator should be used to ensure correct placement. The configurator can be downloaded from:

- For EMEA: http://www.pc.ibm.com/europe/configurators/
- For USA: http://www.pc.ibm.com/us/eserver/xseries/library/configtools.html
- For other countries or regions:

  a. Go to http://www.ibm.com
  b. Click **Select a country**
  c. Select your country
  d. Click **Products and Services**
  e. Click **Intel-based servers**
  f. Click **Tools**
  g. Scroll down to find the Rack Configurator section

3.5 Power considerations

The x450 ships with two redundant, hot-swappable power supplies that produce 1050 W each at 220 V, or 550 W each at 110 V. When the x450 is populated with more than two processors, memory, and adapters, the power supplies may not be redundant if they are connected to a 110 V power source.

Therefore, IBM recommends that the x450 be connected to a 220 V power source to ensure power supply redundancy for large configurations.

Tip: If power is not redundant, the non-redundant LED will be lit in the Light Path Diagnostic panel (see 1.10, “Light path diagnostics” on page 25).

Two system power-cord connectors are available on the back of the x450, one for each of the power supplies. Connect each of these power connectors to separate power circuits to ensure availability if one circuit should fail.

The x450 ships with two 2.8 m/9 ft. IEC 320-C13 to IEC 320-C14 power cables for intra-rack power distribution. Models shipped in the US also include two 2.8

m/9 ft. IEC 320-C13 to NEMA 6-15P power cords for attachment to high-voltage power sources.

3.6 Solution Assurance Review

Some level of Solution Assurance Review (SAR) should be performed for all IBM solutions. The level of SAR (self, peer, or expert) should match the complexity of the solution. For example, simpler solutions may need only a self review. However, a combination of the customer environment risk combined with the complexity of the solution may require that an expert level SAR take place, facilitated by a Quality Assurance practitioner and supported by a team of technical experts.

If a solution contains four or more Enterprise X-Architecture servers (currently x360, x440 and x450), then an expert SAR is mandatory.

For further information on what is required, refer to the Solution Assurance Web sites. EMEA and Americas information is available. Procedures for Asia Pacific countries are currently in development.

Trigger Tool

The SAR Trigger Tool provides a recommendation on the level of Quality and Solution Assurance that will be required. It is available from one of the Solution Assurance Web sites.

The three levels are:

- Expert
  - For technically challenging, high-risk solutions
  - Process dictates expert personnel's participation
  - Formal, rigorous
- Peer
  - For low-to-medium-risk solutions
  - Informal, inexpensive
- Self
  - For low-risk solutions
  - Informal, inexpensive

eSAR — Electronic Solution Assurance Review

There is also an eSAR tool available to help you establish whether you require an expert review. This is available from:

- For IBM employees: http://w3.ibm.com/support/assure, then click **eSAR**
- For Business Partners: http://www.ibm.com/support/assure/esar

4

Installation

In this chapter, we describe both the basic Extensible Firmware Interface (EFI) operations that administrators should know prior to installing an operating system and the operating system installation procedures that are specific to the x450.

The following topics are discussed:

- 4.1, “Using the Extensible Firmware Interface” on page 68
- 4.2, “Installing Windows Server 2003” on page 86
- 4.3, “Installing SuSE Linux Enterprise Server” on page 98

Prior to commencing installation, you need to download the latest firmware and drivers. These are all available from the x450 driver matrix at:

http://www.pc.ibm.com/qtechinfo/MIGR-4JTS2T.html

4.1 Using the Extensible Firmware Interface

Intel has introduced the Extensible Firmware Interface (EFI) as a means of giving greater control to the operating system and to the end user as to how the system starts and works. In this chapter we familiarize you with this environment.

The concept of the EFI is described in 1.3, “Extensible Firmware Interface” on page 5.

Table 4-1 shows the common tasks you would perform in either BIOS or via a DOS diskette and how they are now performed on the x450.

Table 4-1 New ways to do familiar tasks in the EFI

If you want to	On 32-bit xSeries system	On the x450
Configure the internal SCSI/RAID controller	Enter adapter’s configuration utility during server boot	Run the LSI Logic SCSI Setup utility from the EFI shell. See “Configuring SCSI controller and internal disks” on page 74.
Configure and set up the server	Press F1 during server boot and enter Configuration/Setup utility	Run the Configuration and Setup utility from the EFI Boot Manager menu. See 4.1.3, “Configuration and Setup utility” on page 83.
Configure the Remote Supervisor Adapter	Boot from Remote Supervisor Adapter CD-ROM or diskette and run the configuration utility	Run the Remote Supervisor utility from the EFI shell. See “Remote Supervisor Adapter firmware update” on page 82.
Run system diagnostics	Press F2 during server boot and enter the Diagnostic utility	Run the Diagnostic utility from the EFI Boot Manager menu. See 4.1.4, “Diagnostic utility” on page 84.
Change the boot sequence	Press F1 during server boot and make the appropriate changes in Configuration/Setup utility	Enter the Boot option maintenance menu from the EFI Boot Manager. See 4.1.5, “Boot maintenance menu” on page 85.
Flash the firmware, BIOS or Service Processor	Boot from a Flash diskette with appropriate updates and run flash utility	Run the Flash Update from the Boot Manager menu. See 4.1.2, “Flash update” on page 80.
Make changes in disk partitioning before the operating system is installed.	Boot from a DOS diskette and run FDISK.EXE	Run DISKPART from the EFI shell. See “Deleting the disk content” on page 78.
Basic file operations on the FAT/FAT32 partitions	Boot from DOS diskette and use **cd**, **dir**, **copy**, **move**, **rm** commands	Enter the EFI shell from Boot Manager menu and use **cd**, **dir**, **cp**, **mv**, **rm** commands. See “File operations in the EFI shell” on page 71.

When you boot the x450, you automatically enter EFI Boot Manager menu (as shown in Figure 4-1). The menu is a starting point from which you initiate any actions. After installing an operating system, the option to start the operating system becomes the first one. By default when booting the server, the first option is loaded if no key is pressed for 30 seconds.

The menu itself may be altered as described in 4.1.5, “Boot maintenance menu” on page 85.

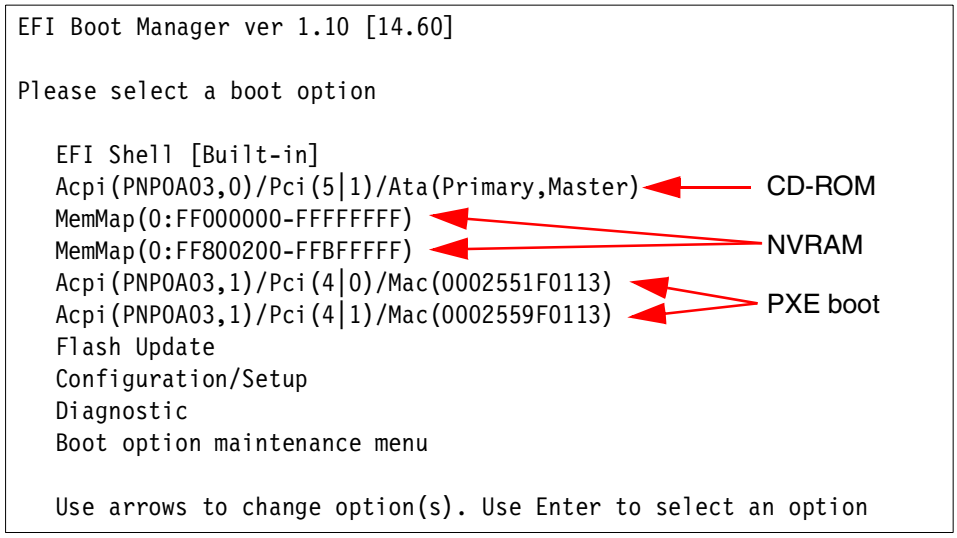

Figure 4-1 EFI Boot Manager menu

In the menu, there are five options marked as CD-ROM, NVRAM and PXE boot as shown in Figure 4-1. These indicate the CD-ROM drive, two default memory mappings (discussed below in 4.1.1, “The EFI shell” on page 70) and two PXE boot devices, referring to the two Broadcom network controllers. If you choose any of these options, the system will try to boot from the selected device.

If it fails (for example, when trying to boot from CD-ROM without bootable media inserted), the following message appears:

```
Loading.: Acpi(PNP0A03,0)/Pci(5|1)/Ata(Primary,Master)
Load of Acpi(PNP0A03,0)/Pci(5|1)/Ata(Primary,Master) failed: Not Found
Paused - press any key to continue
```

Choosing any of the MemMap mappings will always cause a similar failure message because they both contain built-in EFI utilities and are not bootable.

When you select one of the network controllers (marked PXE boot), you initiate the PXE remote boot process. When no PXE boot server is discovered on the network during this process, a failure message appears as well.

Beyond the options mentioned, in this section we cover the following:

- The EFI shell
- Flash update
- Configuration and Setup utility
- Diagnostic utility
- Boot maintenance menu

4.1.1 The EFI shell

The EFI shell allows the loading of EFI applications (including booting installed operating systems) from any EFI-defined file system. The EFI environment replaces DOS to perform maintenance operations, such as upgrading firmware or running system diagnostics.

To start the EFI shell, select **EFI Shell [Built-in]** from the Boot Manager menu. As the shell starts, a device mapping table is displayed, as shown in Figure 4-2.

```
Loading.: EFI Shell [Built-in]
EFI Boot Manager ver 1.10 [14.60]
Device Mapping Table
fs0  : MemMap(0:FF000000-FFFFFFFF)
fs1  : MemMap(0:FF800200-FFBFFFFF)
fs2  : Acpi(PNP0A03,0)/Pci(5|1)/Ata(Primary,Master)/CDROM(Entry0)
fs3  : Acpi(PNP0A03,1)/Pci(3|0)/Scsi(Pun0,Lun0)/HD(Part1,Sig4879D0ED-F
blk0 : MemMap(0:FF000000-FFFFFFFF)
blk1 : MemMap(0:FF800200-FFBFFFFF)
blk2 : Acpi(PNP0A03,0)/Pci(5|1)/Ata(Primary,Master)
blk3 : Acpi(PNP0A03,0)/Pci(5|1)/Ata(Primary,Master)/CDROM(Entry0)
blk4 : Acpi(PNP0A03,1)/Pci(3|0)/Scsi(Pun0,Lun0)
blk5 : Acpi(PNP0A03,1)/Pci(3|0)/Scsi(Pun0,Lun0)/HD(Part1,Sig4879D0ED-F
Shell>
```

Figure 4-2 Entering EFI shell

This mapping table lists both block devices (blk) and file systems (fs). All fixed disks, CD-ROM drives, and USB storage devices will appear as block devices. If the server contains additional controllers and disks, they will appear as additional block devices. The EFI automatically maps known file systems on the block devices to their FSx mappings (FSx is referred to as the *file system identifier*).

Here are some tips for easier identification of devices in the list:

- Even with no storage devices, the server has at least two default mappings — the MemMap entries listed as `fs0:` and `fs1:` in Example 4-2 on page 70. These are read-only RAM disks created during POST so they will always be available. The first MemMap file system is the content of the 4 MB of NVRAM assigned to the EFI, and contains a number of drivers. The second MemMap file system is the content of the NVRAM assigned to the server diagnostics.
- The internal disks contain `Scsi` and `HD` strings. In Figure 4-2 on page 70:

  `Scsi(Pun0,Lun0)/HD(Part1,Sig4879D0ED-FFF4-456D-A15C-DD6C856DF2F5)`

  where `Lun0` is disk's SCSI ID, `Part1` stands for Partition1, and `Sig` is the disk's GUID (discussed in 1.3.1, "GUID Partition Table disk" on page 7).
- The disk partitions are numbered from 1 onwards: the first partition is Part1, the second Part2, and so on.
- EFI System Partition is easily identifiable in Figure 4-2 on page 70 if we realize that EFI is always the first partition on the disk (that is, it is identified by `HD Part1` in the string). In this case it is fs3.
- In Figure 4-2 on page 70, the CD-ROM is represented as fs2.

  Note: Unless CD-ROM media is inserted, there is no FSx mapping created for the CD-ROM.

- The USB memory devices include the string USB.

File operations in the EFI shell

From the EFI shell you may access FAT/FAT32 file systems on existing disk partitions, run EFI executables, and manipulate with disk contents on various media. The EFI executables are text mode utilities similar to DOS or UNIX commands. In general, EFI commands are not case-sensitive.

Tip: All the EFI executable files have an extension of .EFI. The executables may be run only from within EFI, not from within any operating system.

Access to the file system on a device is done through the **`FSn:`** command (where **`n`** is the file system number shown in the device mapping table). This is similar in concept to changing drive letters in DOS. You can use the **`map`** command at any time to remind yourself of the available file systems.

Tip: Adding or removing media may cause arbitrary reassignments of the FSx designations. Always make sure you access the right media by issuing the **map** command.

As an example of the capability of accessing file systems, the following are the steps to access data stored on the USB memory key device:

1. Insert the USB memory key into the USB port (can be hot-added).
2. Start the EFI shell from the Boot Manager menu.
3. Use the **map** command to identify the appropriate mapping for the USB device the EFI automatically created. Sample output of the command is in Figure 4-3.

   If the appropriate mapping does not appear, run **map -r**, where parameter **-r** forces EFI to refresh the device mapping table. You may also have to exit and re-enter the shell before the FSx entry appears.

```
Shell> map
Device Mapping Table
fs0  : MemMap(0:FF000000-FFFFFFFF)
fs1  : MemMap(0:FF800200-FFBFFFFF)
fs2  : Acpi(PNP0A03,0)/Pci(5|2)/Usb(0, 0)/HD(Part1,Sig00000000)
blk0 : MemMap(0:FF000000-FFFFFFFF)
blk1 : MemMap(0:FF800200-FFBFFFFF)
blk2 : Acpi(PNP0A03,0)/Pci(5|1)/Ata(Primary,Master)
blk3 : Acpi(PNP0A03,0)/Pci(5|1)/Ata(Primary,Slave)
blk4 : Acpi(PNP0A03,1)/Pci(3|0)/Scsi(Pun0,Lun0)
blk5 : Acpi(PNP0A03,0)/Pci(5|2)/Usb(0, 0)
blk6 : Acpi(PNP0A03,0)/Pci(5|2)/Usb(0, 0)/HD(Part1,Sig00000000)
```

Figure 4-3 Map command

4. Look for the specific device in the FSx section. In our example, the entry is highlighted in Figure 4-3. The entry will include the string `USB`.

   If the USB device is attached to a USB hub (for example a memory key plugged into the USB port of a USB keyboard), then two USB strings will appear in the entry, as shown in Figure 4-4 on page 73.

```
Shell> map
Device Mapping Table
fs0 : MemMap(0:FF000000-FFFFFFFF)
fs1 : MemMap(0:FF800200-FFBFFFFF)
fs2 : Acpi(PNP0A03,0)/Pci(5|2)/Usb(0,0)/Usb(3,0)/HD(Part1,Sig00000000)
blk0 : MemMap(0:FF000000-FFFFFFFF)
blk1: MemMap(0:FF800200-FFBFFFFF)
blk2: Acpi(PNP0A03,0)/Pci(5|1)/Ata(Primary,Master)
blk3: Acpi(PNP0A03,1)/Pci(3|0)/Scsi(Pun1,Lun0)
blk4: Acpi(PNP0A03,0)/Pci(5|2)/Usb(0, 0)/Usb(3, 0)
blk5: Acpi(PNP0A03,0)/Pci(5|2)/Usb(0,0)/Usb(3,0)/HD(Part1,Sig00000000)
```

Figure 4-4 Map command - USB memory key connected to the USB keyboard

5. Access the device by its name followed by the colon. For example **fs2:**
6. Use the commands **cd**, **dir** (or **ls**), **copy** (or **cp**), **move** (or **mv**), **rm** to manipulate the files on the device.

In the same way, you can access any other devices, for example removable media, CD/DVD-ROM media, and the FAT/FAT32 disk partitions.

Tip: CD-ROM media supported in the x450 must comply with the ISO9660 (El-Torito) format. Joliet-formatted CD-ROMs (which is typically the default) are not readable from within the EFI. When burning a CD-ROM that is to be readable from within EFI shell, specify the following settings:

- File system: ISO9660 or El-Torito
- Physical format: Mode 2, CD-ROM XA

In Easy CD Creator, do this by clicking **File** → **CD Properties**, change the file system to ISO9660, and ensure the physical format of the CD is set to Mode 2: CD-ROM XA.

Note: Although x450 is equipped with a DVD/CD-RW device, writing to writable or rewritable media from the EFI shell is not supported.

There are a variety of built-in shell commands. Of note are:

- **help -b** displays a list of the available commands. The -b parameter pauses the display after each screenful of information and can be used by many EFI commands.
- To view a content of text file, use the **viewer** command.

- To edit any text file, use the **edit** command.
- To restart the server (no confirmation), use **reset**. To power off the server, use **reset -s**.

Configuring SCSI controller and internal disks

The internal LSI SCSI/RAID controller can treat the internal disks in two different ways:

- As two standalone disks
- As two parts of a mirrored volume (that is, RAID-1 capability)

Both ways are supported for operating system installation, although installing on the single non-redundant disk is not recommended.

The embedded LSI controller supports a RAID-1 (mirror) configuration of the two drives in the internal drive bays. The only possible configuration consists of a single logical drive spanning the two disks. There is no way to create multiple logical volumes at the controller level.

The advantage of the embedded LSI controller is that the content of the primary disk (the disk marked as Primary in the Mirroring Properties menu) stays intact and is copied to the secondary disk during the first synchronization. This feature means that the installation of the operating system and the creation of the mirror can be done in any order.

Tip: Even though the content of the primary disk stays intact after mirror creation, we recommend you create a full backup before the procedure.

In the following steps, we describe the process of first creating a RAID-1 array, then installing the operating system once the logical volume has been created.

Boot devices: Booting from any disks other than those attached to the internal LSI controller is not supported.

To run the LSI Logic SCSI Setup utility to create and configure mirror volume, perform the following steps:

1. Boot your server to the Boot Manager menu. Invoke the EFI shell from the menu. Run **drvcfg** to determine the device handle number for the LSI driver. The output will be similar to the following:

   ```
   Shell> drvcfg
   Configurable Components
   Drv[31]  Ctrl[47]  Lang[eng]
   Drv[41]  Ctrl[54]  Lang[eng]
   Drv[41]  Ctrl[55]  Lang[eng]
   Drv[50]  Ctrl[52]  Lang[eng]
   Drv[50]  Ctrl[53]  Lang[eng]
   ```

   Note the driver handle dd, in `Drv[dd]` for each device.

2. For each driver handle, run **dh -d dd** command to display the detailed information. We are looking for the LSI device driver. In our example, we issue the command three times:

   ```
   dh -d 31
   dh -d 41
   dh -d 50
   ```

   The output for driver handle 50 shows the LSI driver:

   ```
   Shell> dh -d 50
    50: Image(PciRom Seg=00000000 Bus=01 Dev=03 Func=00 Image=0000)
   DriverBinding ComponentName Configuration Diagnostics
        Driver Name    : LSI Logic Ultra320 SCSI Driver
        Image Name     : PciRom Seg=00000000 Bus=01 Dev=03 Func=00
   Image=0000
        Driver Version : 01010200
        Driver Type    : DEVICE
        Configuration  : YES
        Diagnostics    : YES
        Managing       :
          Ctrl[52] : LSI Logic Ultra320 SCSI Controller
          Ctrl[53] : LSI Logic Ultra320 SCSI Controller
   ```

   After identifying the LSI driver handle (Driver Name: `LSI Logic Ultra320 SCSI Driver` in the output above), invoke the LSI configuration for the appropriate controller.

   In our example, there are two controllers 52 and 53, displayed in the last two lines of the **dh** command. In the x450, only the first controller (in this case Ctrl[52]) is used to connect the drives in the drive bays. The second controller (Ctrl[53]) is for connecting the external devices.

Tip: The LSI driver handle number may change after each reboot. Each time you configure the LSI controller, perform these steps starting at step 1.

3. To configure the first LSI controller, issue the command **drvcfg dd cc -s**.

   Where:

   – `dd` is the LSI driver handle
   – `cc` is the controller to be configured
   – `-s` indicates that a Setup operation is required

   In our example, the command would be

   drvcfg 50 52 -s

   The command starts the LSI Logic SCSI Setup utility.

4. On the first screen, select the adapter (in yellow text, there will be only one) and press Enter. The Adapter Properties menu appears.

   Low-level format: If you wish to perform a low-level format on the drives before you create a mirrored pair, do the following:

   1. From the Adapter Properties menu, select **Device Properties**.
   2. Use the down arrow key to scroll to the drive you wish to format.
   3. Use the right arrow key to select **Format** (initially this option is off the screen to the right) and press Enter
   4. Press Enter to begin the format.

5. Select **Mirroring Properties**. A window similar to Figure 4-5 on page 77 appears.

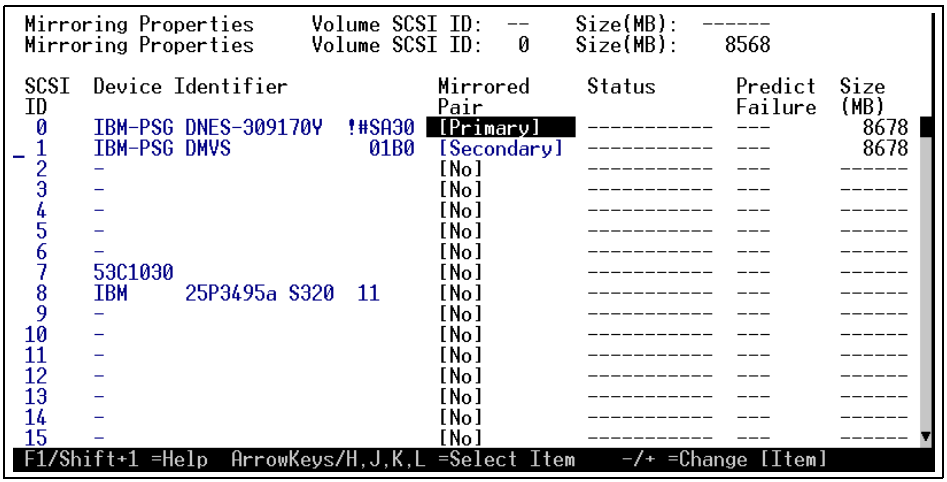

Figure 4-5 Enable mirroring on the LSI adapter

6. For both disks, change disk properties in the Mirrored Pair column.

   The possible values are Primary, Secondary and Hot spare. Pressing the Spacebar key marks one of the disks as Primary and the other as Secondary. The physical disks in the internal drive bays are highlighted (SCSI ID 0 and 1).

7. Exit the window by pressing Esc. No changes are made to the configuration at this point.

   Tip: The option to mark the disk as hot spare is of no use on the x450. Because of the physical limitation of the server (internal maximum of two disks), it is not possible to add a hot spare disk to the mirrored volume configuration anyway.

8. The following message appears:

   ```
   Warning: Data on Secondary or Hot Spare! Mirroring will overwrite!
   ```

   Confirm by pressing Enter and continue.

   Note: After confirming the mirror creation, the data on the secondary disk is lost. The secondary disk will be overwritten by the content of the primary disk whose content stays intact.

9. Save the changes on the next window and exit the menu. The disks will synchronize in the background.

Note: The initial mirror synchronization can take up to several hours. The process itself runs in the background. You can begin the operating system installation at any time, even while the synchronization is running.

During the synchronization, the server may be restarted or turned off as needed. If not finished, the procedure will continue after restarting without user intervention.

10. The following message appears:

    ```
    Drv[dd] Ctrl[cc] Lang[eng] - Options set. Action Required is Stop
    Controller
    ```

    Press Enter to stop the controller.

11. Enter the command **`connect dd`**, where dd is the previously used the LSI driver handle to restart the controller.

    In our example, the command would be **`connect 50`**.

Note: Actions described in steps 10 and 11 do not affect the progress of the mirror synchronization.

Deleting the disk content

There are situations when it is necessary to delete the contents of the disk. For example, you want to install an operating system on a disk that was previously used in a non-EFI system (that is, with a disk with a standard Master Boot Record (MBR), or you want to delete a partition that is reserved and cannot be deleted by other means (for example, a Microsoft Reserved Partition — see “Microsoft Reserved Partition” on page 88).

Tip: If you need to perform a low-level format of the disk, use the LSI Configuration utility, which was described in “Configuring SCSI controller and internal disks” on page 74.

To delete the content of the disk, perform the following steps:

1. Download the DISKPART.EFI drive partitioning utility. Transfer it to the x450 using a device such as a USB memory key or CD-R. This is available from:

   `http://developer.intel.com/technology/efi/diskutil_overview.htm`

2. Invoke the EFI shell from the Boot Manager menu. Change to the device where the diskpart utility is (for example, CD-ROM or USB memory key) and run it. A DiskPart prompt appears, similar to the following:

```
DiskPart Version 1.0
Based on EFI core release Version 1.2.1.0
DiskPart>
```

3. Issue the following commands as shown in Figure 4-6:

   `list` List all the disks in the system that can be deleted.

   `select dd` Select the disk you will delete.

   `clean` Erase the contents of the selected disk. Confirm the action as prompted.

   `exit` Exit the console and return to the EFI shell.

   The overall process is shown in Figure 4-6.

```
DiskPart Version 1.0
Based on EFI core release Version 1.2.1.0
DiskPart> list
  ###  BlkSize          BlkCount
  ---  -------  ----------------
    0      200           10BC001
DiskPart> select 0
Selected Disk =   0
DiskPart> clean
About to CLEAN (DESTROY) disk 0, are you SURE [y/n]?
CLEAN>> y
If you are REALLY SURE, type '$C'
CLEAN>> $C
DiskPart> exit
Exiting....
```

Figure 4-6 Diskpart commands to erase a disk

The described procedure can be used to delete both the MBR and GPT disks.

Tip: The hard drive partitioning utility DISKPART.EFI may be used also to create the disk partitions, but to do this we recommend using tools available during or after operating system installation.

4.1.2 Flash update

As part of your installation procedure for the x450, we recommend that you check the firmware levels on the system components and update to the most current revision:

- System firmware
- Onboard diagnostics
- Remote Supervisor Adapter firmware
- Additional devices if installed, such as ServeRAID adapters and FAStT Fibre Channel host adapters

The latest updates and firmware code can be found at:

http://www.pc.ibm.com/qtechinfo/MIGR-4JTS2T.html

System firmware and diagnostics updates

Unlike an IA-32 platform, the normal boot process does not permit booting from a diskette, so a flash update procedure must start differently. The flash program is embedded in the system. To start a flash update, select **Flash Update** from the Boot Manager menu (Figure 4-1 on page 69).

Once the program loads, it looks exactly like the diskette version but with one change in the EFI version. The diskette version searches for the update files on the diskette, whereas the flash program on x450 reads the root directories of all accessible file systems until it finds firmware update files. The implication is that once the utility successfully locates a flash update file, it will stop searching.

Tip: The routine that searches all devices for suitable flash files stops as soon as it finds the first one in the root directory of a file system. This may not be the newest flash file. The routine does not search subdirectories.

We should realize this behavior when we have, for example, the latest firmware updates on a USB memory key and an older version stored in the root of EFI System Partition. In such situations, the flash update utility will keep offering only the older file for update. To resolve the issue in this specific case, we have to delete the firmware update files from the system partition and repeat the procedure. If you want to delete the files from ESP, enter the EFI shell, identify the EFI System Partition and use the **rm** command (for the exact syntax, use the **help** command).

How to recognize update files used with Flash Update: Diagnostics update file are named MZYT*.US1, where * will be a two-digit number indicating the level of the ROM Diagnostics code followed by optional additional characters. For example, MZYT05A.US1.

Firmware POST/BIOS update file are named MZKT*.FLS. For example, MZKT19AUS.FLS.

To see the current version of system components, run **`ver`** in the EFI shell.

To update the system firmware and the diagnostics (we always recommend you do both at once), follow these steps:

Tip: A full backup should be completed and verified prior to running this, or any other system upgrade.

1. Select **Flash Update** from the Boot Manager menu. The Flash Update utility will display.
2. Press 1 to update the POST/BIOS from the menu.
3. You will be asked if you wish to change the serial number, the machine type and the asset tag. Typically, you would answer no.
4. You will then be asked if you would like to save the current code to a disk.

   If you select **Y,** you need to specify a file name for the backup. The Flash Update utility saves the backup file to the writable file system with lowest file system identifier number (FSx). In most cases this is EFI System Partition, if it exists. If the operation fails, make sure some sort of writable media is attached to the system.

   CD-R/W and MemMap mappings are not considered writable devices in EFI.

   In our system, as shown in Figure 4-2 on page 70, the backup file would be saved on fs3. Do not try to specify full path to the file name, because the utility does not accept full path syntax.
5. At this point, the Flash Update utility should detect the presence of the update files on the disk, and you will be asked to choose which language you wish to use during POST and in Setup.
6. Once the flash update completes, you will be prompted to press Enter to restart your system.
7. After rebooting, return to the Flash Update utility.
8. Enter 2 to select **Update Diagnostics** from the menu.

9. Follow the instructions in the window until you successfully update the diagnostics. Press Enter to reboot the server when requested.
10. After rebooting, verify that you have successfully completed the update by re-issuing the **ver** command at the shell prompt.

Remote Supervisor Adapter firmware update

To update the Service Processor flash, we use the IBM Remote Supervisor utility, which besides updating the flash allows us to configure RSA itself. The utility is not a default part of EFI configuration and may be located on the media shipped with the server or downloaded from http://www.pc.ibm.com/support.

Unlike the Flash Update utility, which requires update files to be in the root directory of any file system, to successfully update Service Processor flash, both IBM Remote Supervisor utility (FLASH2.EFI) and the flash update files must reside in the same directory.

Note: As well as updating RSA firmware, the utility updates firmware on RXE-100 Remote Expansion Enclosure if connected.

To apply Service Processor updates, follow these steps:

1. Enter the EFI shell from the Boot Manager menu.
2. Navigate to the folder that contains FLASH2.EFI and the update files.

Note: The following files are required to be in the same directory as the FLASH2.EFI utility to successfully update RSA and RXE-100 code:

CNETBRUS.PKT	System Management Adapter Boot ROM
CNETMNUS.PKT	System Management Adapter Main Application
DMCFLASH.PKT	Integrated System Management Processor Application
IIOEFLSH.PKT	RXE-100 Remote Expansion Enclosure Application
IIOEFLPB.PKT	RXE-100 Power Backplane Application
CNETRGUS.PKT	Remote Graphics Application

3. Run the program from the prompt by typing `flash2` and pressing Enter. The IBM Remote Supervisor utility will display.
4. Select **Update System Management Firmware** from the main menu and press Enter.
5. Read the warning window and press Enter. You will be presented with the Firmware Update Options menu.

6. Select **Normal Flash - All subsystems** and press Enter. This option will flash all management code in the system. Follow the prompts to complete the flash.
7. Press Esc to return to the main menu and quit the utility.

4.1.3 Configuration and Setup utility

By selecting **Configuration and Setup Utility** from the Boot Manager menu, you enter a boot time user setup interface similar to the others on xSeries servers. This utility allows you to view or customize various configuration parameters.

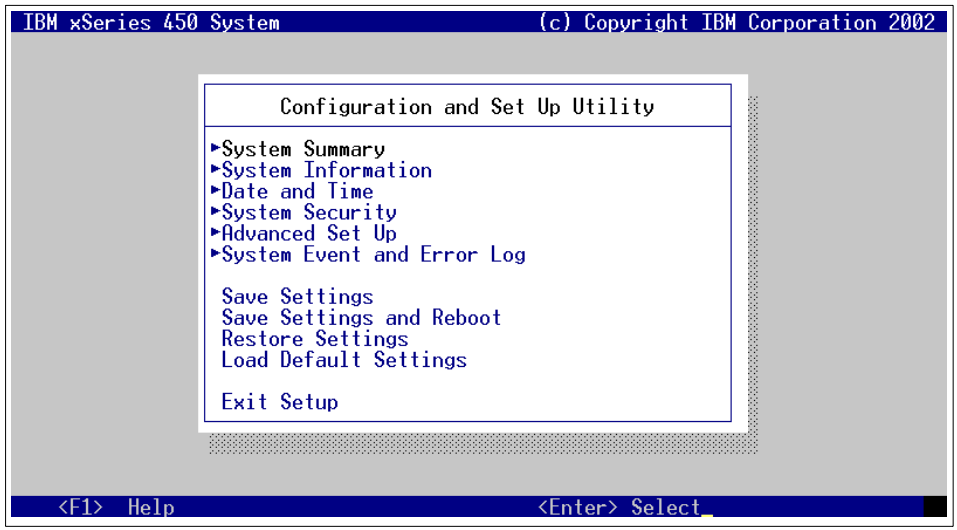

Figure 4-7 Configuration and Set Up utility

Enabling memory mirroring

Memory mirroring (part of IBM's Active Memory technology) provides an additional level of fault tolerance to the memory subsystem. For detailed information and guidelines on memory mirroring, see 1.7, "System memory" on page 17.

To enable memory mirroring on the x450, perform the following steps:

1. Allow the system to boot to the EFI boot menu.
2. Start the Setup program by selecting the **Configuration and Set Up Utility** from the boot menu.
3. On the main window of the Setup program, select **Advanced Set Up**.
4. Select **Memory Configuration** and in the Memory Mirroring section press the right arrow key to change the value to Enabled.

5. Exit the Setup program and save the changes. If you made any changes while in the menus, the system will reboot. If you made no changes, you will be returned to the EFI boot menu.

Note: When adding or replacing memory to the xSeries 450, you may need to run the Setup program for the new memory to work.

Tips:

- When memory is added, it should be detected automatically but may need two reboots to complete the process.
- When replacing faulty memory, you will need to re-enable the faulty DIMM(s) by running Setup, saving the system changes and rebooting again.
- To ensure your memory configuration is correct, enter Configuration and Set Up utility and verify that the settings in the Memory Configuration option match you configuration.

4.1.4 Diagnostic utility

The ROM-based diagnostics package may be started from the EFI Boot Manager by selecting the **Diagnostic Utility**. Alternatively, you can start it using the command **`amidiag`** from the MemMap filesystem (in our example, it was fs1:).

The Diagnostic utility looks similar to the IA-32 ROM version.

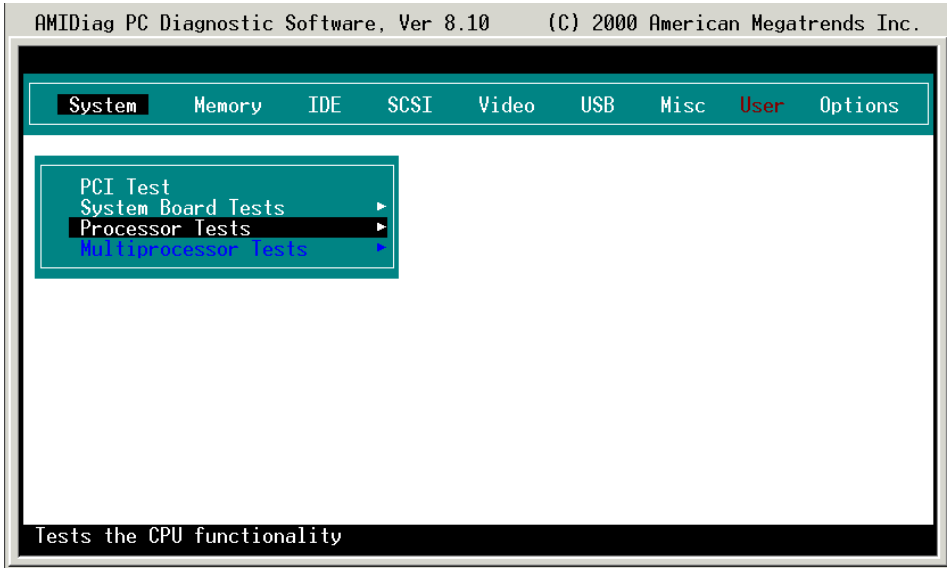

Figure 4-8 Diagnostic utility

In the utility, only data non-destructive tests are available by default. To enable all the tests (including destructive SCSI Disk Write/Format Tests), choose **Options** → **Toggle Hidden Test Display**.

Note: Multiprocessor tests were not available in the Diagnostic utility at the time of writing this book. They should be made available by a diagnostics update.

The processor test only tests the bootstrap processor and does not verify the whole processor system in a multi-processor setup. You can use the Processor Tests option to test a single processor but, if more than one processor is fitted, you will need to remove all the processors and fit each one in turn to correctly test it. Because there is no way to identify what processor is the system bootstrap processor from the EFI environment, this test has limited use.

4.1.5 Boot maintenance menu

The boot maintenance menu allows you to modify the EFI boot menu records, including the entries (add, delete, modify), and to change the order. The order is important because that is the order that the system will try to boot from (similar to the boot sequence option in Setup in IA-32 based xSeries systems).

Important: Do not delete the default boot menu entries, since it is difficult to recreate them manually. If you do delete them, connect an auxiliary console (ANSI terminal connected to the serial port) and press Shift+R when prompted to restore the EFI defaults.

Similarly, operating system boot entries may also be hard to recreate. For Windows-specific information about the Windows boot menu record, see “Exporting and importing the operating system EFI boot record” on page 96.

```
EFI Boot Maintenance Manager ver 1.10 [14.60]

Main Menu. Select an Operation

        Boot from a File
        Add a Boot Option
        Delete Boot Option(s)
        Change Boot Order

        Manage BootNext setting
        Set Auto Boot TimeOut

        Select Active Console Output Devices
        Select Active Console Input Devices
        Select Active Standard Error Devices

        Cold Reset
        Exit

     SystemGuid-->[FFFFFFFF-FFFF-FFFF-FFFF-FFFFFFFFFFFF]
     SerialNumber-->[No Data]
```

Figure 4-9 Boot option maintenance menu

4.2 Installing Windows Server 2003

In this section, we discuss the specifics of installing Windows Server 2003, Enterprise Edition on the xSeries 450.

If you are implementing Microsoft Cluster Service (MSCS) check the Microsoft Cluster Service Hardware Compatibility List (HCL) to confirm that hardware, drivers, and firmware have been Microsoft Certified. The HCL can be found at:

http://www.microsoft.com/hwdq/hcl/

For information about Windows Server 2003 and changes in the 64-bit version, see:

http://www.microsoft.com/windows.netserver/64bit/default.mspx

Note: The dynamic disk concept introduced in Windows 2000 is not affected by introducing GUID Partition Table (GPT) disks in Windows Server 2003. Both Master Boot Record (MBR) and GPT disks can be present in a single dynamic disk group. Volume sets can span both MBR and GPT disks.

The OS loader and the boot partition must reside on a GPT disk. Other hard disks in the system may be either MBR or GPT.

4.2.1 Overall process

The overall process to install Windows Server 2003 on the x450 is as follows:

1. If not already created, configure a mirrored pair of disks using the **drvcfg** command as described in “Configuring SCSI controller and internal disks” on page 74.
2. Optionally, delete all existing partitions (including the system partition and Microsoft Reserved partition) using the **diskpart** command, as described in “Deleting the disk content” on page 78.
3. Ensure the system has the latest of the following installed:
   - System firmware
   - System diagnostics
   - Remote Supervisor Adapter firmware

   Download these from http://www.pc.ibm.com/support
4. Shut down the system and disconnect the RXE-100. Reconnect it after the installation is complete.
5. With the Windows Server 2003 CD-ROM in the drive, boot (or reboot) the system to the EFI Boot Manager menu.
6. Select the CD-ROM entry (will contain the strings “ACPI” and “ATA”) to boot from the disk, as described in “Booting the installation CD automatically” on page 90.

7. If you performed step 2, follow the instructions to create the system partition and Microsoft Reserved partition, as described in 4.2.6, “Text-mode setup” on page 91.
8. Complete the Windows installation as instructed.
9. Install the additional device drivers, as described in 4.2.8, “Post-setup phase” on page 93.
10. Make a backup copy of the Windows entry in the Boot Manager menu, as described in “Exporting and importing the operating system EFI boot record” on page 96.

4.2.2 Microsoft Reserved Partition

The *Microsoft Reserved Partition* (MSR partition) reserves space on each disk drive for subsequent use by the operating system software. For instance, dynamic disk configuration is stored in the MSR partition.

During installation, the MSR partition is automatically created.

Every GPT disk used by Windows must contain an MSR partition. The order of partitions on the disk should be ESP (if the disk is a boot disk), MSR partition, followed by primary data partition(s). It is particularly important that the MSR partition is created before other primary data partitions.

Note: Each Windows *boot drive* must contain an ESP, an MSR partition, and at least one basic data partition that contains the operating system. Each GPT *data drive* must contain at least an MSR partition and one basic data partition.

The order of partitions on the Windows boot disk is shown in Figure 4-10.

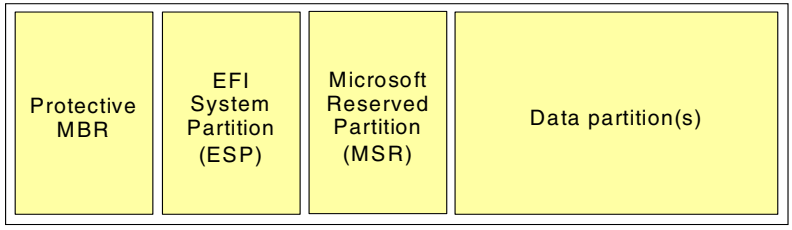

Figure 4-10 Windows boot GPT disk structure

When initially created, the size of the MSR partition depends on the size of the disk drive:

- On drives less than 16 GB in size, the MSR partition is 32 MB.
- On drives greater than or equal to 16 GB, the MSR partition is 128 MB.

Note: The MSR partition is not accessible to the operating system users by any means.

For information on the partitioning required by a Windows 64-bit installation, see

http://www.microsoft.com/technet/prodtechnol/winxppro/reskit/prkb_cnc_helm.asp

4.2.3 Windows installation

The installation of Windows Server 2003 on the server consists of these five phases:

- Pre-installation phase
- Starting the installation
- Text-mode setup
- GUI setup
- Post-setup phase

4.2.4 Pre-installation phase

Perform the following tasks before you begin the installation:

- Before powering on the server, verify that the hardware is configured properly (proper cabling, using only supported disks, adapters and memory modules).
- Before commencing the installation, ensure that firmware levels have been updated to the latest levels. Refer to 4.1.2, “Flash update” on page 80.
- Refer to “Installing additional drivers” on page 93 for information on device drivers required for the installation of Windows Server 2003.
- Disconnect the RXE-100 if you have one. Once the installation is complete, you may reconnect it.

Note: If you do leave the RXE-100 connected during installation, then when you reboot in future, if you ever disconnect the RXE-100, the boot sequence will hang at the black screen immediately after you see the “Starting Windows” text. (This restriction may be corrected by a firmware update by the time you read this.)

4.2.5 Starting the installation

Before you begin, make sure:

- You have created a mirror volume for Windows installation. We recommend you install the operating system on a RAID-1 mirror rather than a single disk. See “Configuring SCSI controller and internal disks” on page 74.
- You use a GPT disk for the installation. If you are about to install Windows on a disk previously used by an IA-32 system (with a MBR disk structure), follow the procedure discussed in “Deleting the disk content” on page 78.

> **Tip:** If the disk you install to is still an MBR disk, the installation procedure will prompt you and will convert the disk to a GPT structure.

There are two ways to start the operating system installation: automatically or manually.

Booting the installation CD automatically

From the Boot Manager menu, select the CD-ROM in the list (as shown in Figure 4-11) and press Enter.

```
EFI Boot Manager ver 1.10 [14.60]

Please select a boot option

EFI Shell [Built-in]
Acpi(PNP0A03,0)/Pci(5|1)/Ata(Primary,Master)
MemMap(0:FF000000-FFFFFFFF)
MemMap(0:FF800200-FFBFFFFF)
Acpi(PNP0A03,1)/Pci(4|0)/Mac(0002551F0113)
Acpi(PNP0A03,1)/Pci(4|1)/Mac(0002559F0113)
Flash Update
Configuration and Set Up Utility
Diagnostic Utility
Boot option maintenance menu
```

Figure 4-11 Boot Manager menu — selecting the CD-ROM

For three seconds, there will be a message on the screen `Press any key to boot from CD-ROM`. If you do not press any key, the installation will not start and you will return to EFI.

Booting the installation CD manually

Starting the installation process manually proceeds as follows:

1. Insert the installation CD-ROM into the drive.
2. Enter the EFI shell from the Boot Manager menu. Access the CD-ROM drive as described in “File operations in the EFI shell” on page 71.
3. Run the SETUPLDR.EFI program. If the CD-ROM drive is mapped as fs2:, the command line looks similar to:

   ```
   fs2:\> setupldr
   ```

For three seconds, there will be a message on the screen `Press any key to boot from CD-ROM`. If you do not press any key, the installation will not start and you will return to EFI.

4.2.6 Text-mode setup

The installation process has now begun.

Tip: When you see the `Setup is starting Windows` message on the video display, the auxiliary console (a terminal session via the serial port) will stop functioning. This is normal.

If you install the operating system on a GPT disk without the EFI System Partition (that is, a first-time installation), you will be asked to create one as shown in Figure 4-12.

```
Windows .NET Server 2003, Enterprise Edition Setup
=================================================

Setup could not locate an existing system partition.

System partitions contain diagnostic or hardware configuration
programs, programs to start operating systems (such as Windows), or
other manufacturer-supplied programs.

Setup will try to create a system partition for you automatically.

To allow Setup to create system partition, press ENTER.
If you want to create a system partition on your own, press ESC.
```

Figure 4-12 ESP creation during Windows installation

Allow Setup to create the EFI System Partition by pressing Enter.

Tip: At this point, in addition to the EFI System Partition, Setup also creates the Microsoft Reserved Partition (see 4.2.2, “Microsoft Reserved Partition” on page 88).

After both partitions are automatically created, the Disk partitioning screen is displayed, as in Figure 4-13.

```
Windows .NET Server 2003, Enterprise Edition
============================================

The following list shows the existing partitions and
unpartitioned space on this computer.

Use the UP and DOWN ARROW keys to select an item in the list.

* To set up Windows on the selected item, press ENTER
* To create a partition in the unpartitioned space, press C.
* To delete the selected partition, press D

8679 MB Disk 0 at Id 0 on bus 0 on symmpi [GPT]

    -:  Partition1 [FAT]                     102 MB ( 101 MB free)
    -:  Partition2 [Reserved]                 31 MB (   0 MB free)
        Unpartitioned space                 8545 MB


ENTER=Install  D=Delete Partition  F3=Quit
```

Figure 4-13 Disk partitioning screen

These partitions are as follows:

- Partition1: EFI System Partition
- Partition2: Microsoft Reserved Partition

Do not attempt to install the operating system into first two partitions. Instead, use the unpartitioned space on the disk to create the data partition.

If the disk you are installing on is still the old Master Boot Record (MBR) partitioning structure, you should delete all partitions, and press S to change to a GUID Partition Table (GPT) partitioning structure. The OS loader and the boot partition must reside on a GPT disk.

Tip: Once created, the MSR partition is a reserved partition and cannot be deleted from the Disk partitioning menu. To delete the MSR partition, the entire disk can be deleted as described in “Deleting the disk content” on page 78.

The sizes of the ESP and the MSR partition are calculated according to the rules described in 1.3.2, “EFI System Partition” on page 8 and 4.2.2, “Microsoft Reserved Partition” on page 88 respectively.

Tip: Do not remove or resize the ESP and MSR partitions after they are created. The ESP partition appears as Partition1 [FAT] and MSR partition as Partition2 [Unknown] in the Disk partitioning menu during text-mode setup.

Follow the instructions to complete the text-mode portion of the installation.

At the end of the text-mode portion, a new Windows entry is added to the EFI Boot Manager menu. After you restart the system at the end of text-mode setup, the entry is launched automatically. (If Windows doesn’t boot automatically, select it from the Boot Manager menu.)

4.2.7 GUI setup

Complete the standard Windows Server 2003 installation per the instructions provided in the product documentation.

Tip: The installation process may take significantly longer than shown by the countdown timer.

4.2.8 Post-setup phase

Once Windows has been installed, you should perform the following steps.

Installing additional drivers

Specific device drivers are available for the x450 that are not included with the base operating system. Drivers that you should obtain separately from the operating system include:

- Broadcom Ethernet controller
- Remote Supervisor Adapter service processor
- Active PCI-X controller

The CD that ships with the server contains these drivers. The latest versions can be downloaded from:

http://www.pc.ibm.com/qtechinfo/MIGR-4JTS2T.html

Important: Drivers that were developed for any 32-bit versions of Windows will not work in 64-bit version of Windows. You must use 64-bit drivers.

After Windows Server 2003 is installed, Device Manager will report a number of unknown devices. These unknown devices correspond to specific components as shown in Figure 4-14.

Note: Some device drivers released for use with Windows Server 2003 may not be digitally signed. Therefore, when you install the drivers, you may have to accept the warning that you are about to install an unsigned driver.

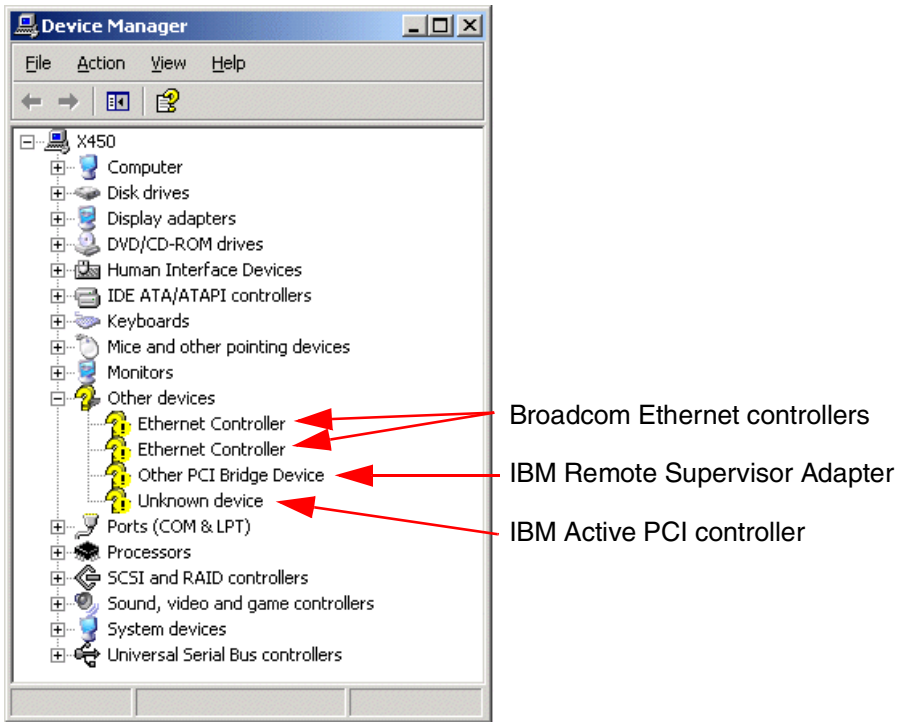

Figure 4-14 Device Manager after Windows installation

When the Windows Server 2003 installation has completed, you should perform install drivers for each of the unknown devices. For each one, right-click the entry in Device Manager and click **Update Driver**. Select the options that let you

specify a location for the drivers and point the installer to the directory where you downloaded the drivers.

If you have other devices installed, such as FAStT Fiber Channel or ServeRAID adapters, update or add the drivers for these. ServeRAID device drivers should be at the same level as the installed ServeRAID firmware and BIOS. The ServeRAID driver that is included in Windows Server 2003 is Version 4.70.

Figure 4-15 shows Device Manager after the Active PCI, RSA, and Broadcom Gigabit Ethernet drivers have been installed.

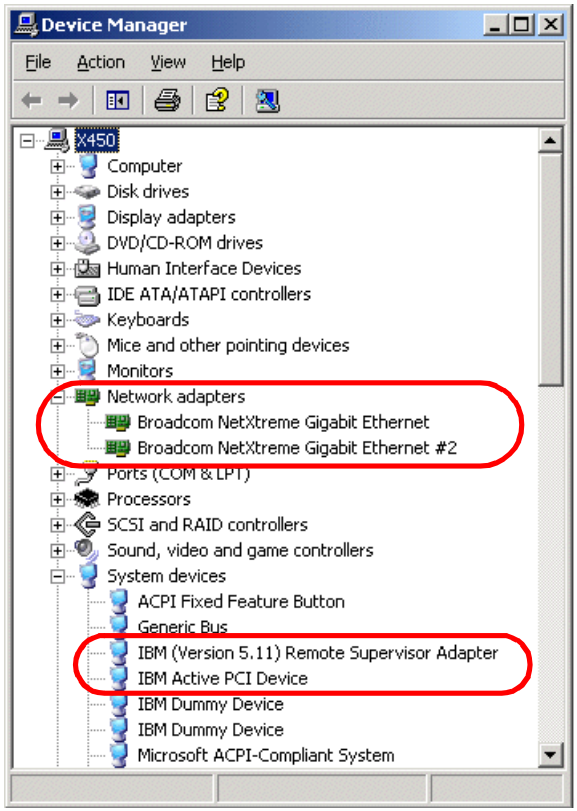

Figure 4-15 Device Manager after device driver updates

Attaching the RXE-100

Once Windows Server 2003 is installed, you can attach the RXE-100. Follow these steps:

1. After the installation of Windows Server 2003 is complete, shut down the operating system and power off the server.

2. Connect the RXE-100 to the x450 as described in “Remote Expansion Enclosure” on page 51.
3. If power is connected to the RXE-100, remove for 10 to 20 seconds.
4. Re-apply power to the RXE-100.
5. Power on the server. The enclosure will power on automatically as the server is started.

Windows Server 2003 supports all 12 slots in the RXE-100.

Accessing EFI System Partition from Windows

To access the EFI System Partition from within operating system, use the following utility.

```
mountvol DriveLetter: /s
```

Where DriveLetter is any spare drive letter in the system.

The ESP will be accessible for both read and write operations, although altering its content is not recommended except, for example ,adding new firmware updates to its root for later use with the Flash Update utility from within EFI.

Important: Do not alter the content of EFI directory on the EFI System Partition. As described in 1.3.2, “EFI System Partition” on page 8, this directory contains the files necessary to boot installed operating systems.

Exporting and importing the operating system EFI boot record

During the installation process, the Setup program adds a new entry to the EFI Boot Manager menu.

It will be added to the top of the menu so that if no keys are pressed once the EFI Boot Manager menu appears, Windows will automatically start. The default delay is 30 seconds and may be changed using the **Set Auto Boot Timeout** option in the Boot option maintenance menu.

The EFI Boot Manager records are stored in NVRAM (non-volatile RAM). If the boot record is deleted from EFI NVRAM, there is no other way to start Windows except for restoring the boot record from backups made using the NVRBOOT.EFI utility.

Tip: After Windows Server 2003 is installed, use NVRBOOT.EFI to export the Windows boot record, then store the file in a safe location. If the system partition gets corrupted, you can easily recreate it and then restore the Windows boot record.

In this section we show you how to save the Boot Manager menu and how to recover it in case of an NVRAM failure.

In 1.3.1, “GUID Partition Table disk” on page 7, we describe how GPT disks are identified by their global unique identifier (GUID). The entries in the Boot Manager menu use the GUID to identify which device to boot from.

There are wo scenarios when the EFI Boot Manager menu will need to be restored:

- Disks rearrangement. If the disks are moved from one server to another, the NVRAM contents and the disk contents are out of sync, meaning that there will be no record in the new server’s Boot Manager menu to point at the recently added disks.
- I/O board replacement. The I/O board contains the system’s NVRAM, so when it is replaced, the operating system boot entry in the EFI Boot Manager menu will need to be manually restored.

To export an EFI boot menu record, perform the following steps:

1. Invoke the EFI shell from the Boot Manager menu. Change to fsx: where x is the system partition (see 4.1.1, “The EFI shell” on page 70).
2. Change the directory to MSUTIL (**`cd msutil`**).
3. Run the NVRBOOT.EFI utility. See Example 4-1.
4. Enter **`X`** (for Export) from the OS Boot Options Maintenance Tool menu.
5. Enter the number of Windows boot record you want to export, or * to export all.
6. Type a file name to store the configuration in and press Enter. A confirmation message as shown in Example 4-1 will display.

   The file will be saved in the root of the EFI System Partition. Back up the file during your regular backup.
7. Press Q and exit the utility.

Note: The utility can export/import all boot menu records, but can only display and modify the Windows boot record.

Example 4-1 Exporting Windows boot record with NVRBOOT.EFI

```
NVRBOOT: OS Boot Options Maintenance Tool [Version 5.2.3683]

    * 1. Windows .NET Server, Enterprise
      2. EFI Shell [Built-in]
      3. Acpi(PNP0A03,0)/Pci(5|1)/Ata(Primary,Master)
```

```
      4. MemMap(0:FF000000-FFFFFFFF)
      5. MemMap(0:FF800200-FFBFFFFF)
      6. Acpi(PNP0A03,1)/Pci(4|0)/Mac(0002551F0113)
      7. Acpi(PNP0A03,1)/Pci(4|1)/Mac(0002559F0113)

  * = Windows OS boot option

(D)isplay (M)odify (C)opy E(x)port (I)mport (E)rase (P)ush (H)elp
(Q)uit

Select> x
Enter OS boot option to export (* = All - Maximum of 30): 1
Enter EXPORT file path: windows.cfg
Saving boot option 1...
Saved Boot Option 1 to file: windows.cfg
Use Import command to retrieve saved boot option

Press enter to continue
```

To recover the boot menu record, use the Import command from the NVRBOOT utility.

To illustrate the export/import process, here's an example. If you were to take a set of disks from one x450 to another, you would do the following:

1. Export the Windows boot entry from the first server and save it on the EFI System Partition (which is a partition on the disk).
2. Insert the disks into the second server.
3. Import the file.
4. Verify the import was successful by booting from the newly created boot entry.

4.3 Installing SuSE Linux Enterprise Server

This section describes the installation of SuSE Linux Enterprise Server 8.0 on the x450.

Tip: To begin the SuSE SLES 8 installation now, go directly to "Installing SuSE Linux Enterprise Server" on page 103.

4.3.1 Background

The goal of the project to port the Linux kernel to the Itanium platform was to have a single optimized port of the Linux kernel for every Itanium-based machine. The kernel has been available at http://www.kernel.org since early 2000.

The development model used to create this new porting of the kernel was the same as used for other Linux development work: open source developers working together around the world on a cooperative effort to bring it available before any other operating system and even before the official release of the processor itself. This is one of the key points of Linux' success, and also of this Itanium 2 port of the kernel.

This effort was also possible thanks to many companies and institutions that contributed on this development effort, such as Intel, IBM, Caldera (now SCO Group), the European Organization for Nuclear Research (CERN), Red Hat, SGI, SuSE, Turbolinux, and VA Linux Systems (now VA Software).

Each company distributed the work among their developers to achieve the best results. The main tasks were:

- Port the kernel itself
- Initial port of GNU Compiler Collection (GCC) to IA-64 architecture
- Port of GAS, emacs, ld, GNUPro toolkit
- Port of performance, measurement, and analysis tools

Intel provided the IA-32 and IA-64 platform specifications and the EFI, and helped on the Apache port. The rest of the work was distributed among the companies involved, and also by private individuals around the world.

Linux for IA-64 provides a new pure 64-bit kernel but maintains compatibility at the API level with Linux for x86 wherever this is possible. The error messages, the system signal, and the ioctl codes will remain as compatible as possible between the two platforms. Also in this version, as started in the Linux kernel from Version 2.1.126 onwards, developers tried to minimize changes to make the code as platform-independent as possible, while optimizing the platform-specific features to improve performance on Itanium processors.

The IA-64 Linux kernel follows the standards defined by Intel and other companies. These include:

- EFI (Extensible Firmware Interface) — see http://developer.intel.com/technology/efi/
- DIG64 (Developers Interface Guide for 64-bit Intel Architecture Servers) — see http://www.dig64.org

- System V UNIX ABI (Application Binary Interface) — see http://www.linuxbase.org/spec/refspecs/elf/gabi4+/contents.html.

4.3.2 Linux IA-64 kernel overview

The following information is helpful if you plan to recompile or port some software to IA-64 kernel. The main objective of this section is to introduce some basic concepts about the Linux kernel and the most important differences between the IA-32 version and the new 64-bit IPF (Itanium Processor Family, or IA-64) architecture. Your preferred Linux distribution may have done this recompilation for you, but it is very common to recompile software under Linux, for example if you get only the source code and not an installable package.

Data types

In 1996, The Open Group decided to extend the *Single UNIX Specification for 64-bit Systems* to create a new standard on 64-bit for any UNIX operating system. Linux kernel developers attempt to follow the Open Group standards as much as possible, so the Linux IA-64 kernel uses the LP64 Data Model, as defined by The Open Group. This data model is LP64 because it allows long and integer types up to 64 bits. 32-bit Intel processors use the LP32 Data Model. This model is an industry standard and it is used in all other 64-bit Linux and UNIX implementations.

Some of the usual data types on C change between the LP64 and the LP32 Data Models, as shown in Table 4-2.

Table 4-2 Comparison table of C types between LP32 and LP64

Data type in C	LP32	LP64
char	8	8
short	16	16
int	32	32
long	32	64
long long	64	64
pointer	32	64
enum	32	32
float	32	32
double	64	64
long double	128	128

Keep these comparisons in mind while programming or porting applications to any 64-bit operating system, not only Linux. It is easy to assume that every data type is the same size, and that's why the LP Data Model standard was created.

Note: The LP64 Data Model example discussed in this redbook is not the only 64-bit Data Model defined by The Open Group. There are other models, such as LP64, ILP64, or LLP64 for 64-bit and ILP32, or LP32 for 32-bit. For more information about the Data Models defined by The Open Group and also about their standardization activities, go to:

http://www.opengroup.org

Byte order

The IA-64 kernel of Linux also uses standard byte order at a byte level. The native byte order is little-endian, while big-endian processes are still possible.

All Intel processors are little-endian based. That means that the bytes at lower addresses have lower significance given a 16-bit or 32-bit word (the word is stored "little-end-first"). In big-endian architectures, the bytes addresses have higher significance (the word is stored "big-end-first"). The IA-64 Linux kernel uses little-endian by default, but allows for the possibility of using big-endian byte order.

4.3.3 Choosing a Linux distribution

Unlike other operating systems, when we talk about Linux, we aren't talking about one company, one version or one distributor. Linux is the kernel, the base of the system, and there are many companies that develop Linux distributions built around the kernel.

In this section, we discuss the major Linux distributors that work with IBM, and their different products available for Itanium 2 platform: Red Hat and UnitedLinux.

Note: The IBM Linux support page is

http://www.pc.ibm.com/qtechinfo/MIGR-48NT8D.html

Here, you can find information about the Linux distributions supported by IBM, and documents, information, and drivers for IBM products.

Note: Debian Linux is also available for Itanium 2 platforms, but we will not discuss this Linux distribution because it is not supported by IBM.

Red Hat Linux 7.2 for Itanium

Red Hat is a major Linux distributor but at the time of writing this redbook, Red Hat did not have any product available for the Itanium 2 platform. The latest released version of Red Hat Linux for Itanium processors is 7.2, but that has now been discontinued.

Red Hat is working on a new release of their Advanced Server that will be available on Itanium 2 platforms (which currently runs on IA-32 architectures only). For the latest information on Red Hat's plans and IBM's support plans, go to:

http://www.redhat.com/software/itanium
http://www.pc.ibm.com/us/compat/nos/redchat.html

The UnitedLinux initiative

In 2002, four of the companies that were developing different Linux distributions - Conectiva, The SCO Group, SuSE and Turbolinux - announced that they will work together to build a unique base for a Linux distribution. The primary idea is to create a base system from which each company will develop their products, solutions, and services.

The official definition of UnitedLinux, as found on their Web site is:

> "UnitedLinux is a standards-based, worldwide Linux solution targeted at the business user and developed by Conectiva, The SCO Group, SuSE, and Turbolinux. Designed to be an enterprise-class, industry-standard Linux operating system, UL provides a single stable, uniform platform for application development, certification, and deployment and allows Linux vendors, Independent Software Vendors (ISVs), and Independent Hardware Vendors (IHVs) to support a single high-value Linux offering rather than many different versions".

UnitedLinux attempts to offer enterprise-level solutions based on Linux, starting from a standard base that each distributor can freely adapt to the needs of its customers. UnitedLinux will also be based on all the Linux standards defined as of this date, such as the Linux Standard Base standard. The main advantage for customers is that they can work with any of UnitedLinux companies, with any of their products, and the base system, the configuration files, and the file locations will remain the same, at the same place and with the same properties. This is a main advantage for customers that use Linux, but do not have a unified distribution across their systems.

For more information about the UnitedLinux initiative, go to:

http://www.unitedlinux.com/en/info/faqs.html

The IPF (IA-64) Linux plans for each of the four UnitedLinux partners are as follows:

- Conectiva Linux

  Conectiva has not officially released any version of their Conectiva Linux product for any Itanium processor. To find out about Conectiva's products and services, go to:

  http://www.conectiva.com

- Caldera SCO Linux 64

  The SCO Group, previously called Caldera, has announced a version of their SCO Linux product for the Itanium family. This product is called Caldera SCO Linux 64.

  To find out about SCO's products on Linux for 64-bit platforms or IBM support for SCO products, go to:

  http://www.sco.com/products/openlinux64
  http://www.pc.ibm.com/us/compat/nos/thecscocgroup.html

- Turbolinux 7.0 for Itanium

  Turbolinux also has a product for Itanium processors, Turbolinux 7.0 for Itanium, but it is not their latest version, which is 8.0.

  To find out about Turbolinux products on Linux for 64-bit platforms or IBM support for SCO products, go to:

  http://www.turbolinux.com/products/tls7i
  http://www.pc.ibm.com/us/compat/nos/turbolinux.html

- SuSE SLES 8 for IA-64

  SuSE Linux Enterprise Server 8 was the only commercial Linux distribution supporting Itanium 2 processors available at the time of writing of this redbook. We had access to a pre-release beta version of this distribution that provided a 2.4.19 kernel optimized for Itanium 2.

  For information on IBM Support of SuSE products, go to:

  http://www.pc.ibm.com/us/compat/nos/suse.html

4.3.4 Installing SuSE Linux Enterprise Server

In this section, we focus on the installation procedure of SuSE SLES 8 and all the issues you may find. We used SuSE SLES 8 Release Candidate 2 to write this redbook.

Requirements and hardware used

According to the *Installation and Administration Addendum* found in the doc directory on the first CD of SuSE SLES 8, the requirements to use this Linux distribution are:

- An Itanium or Itanium 2 processor
- At least 256 MB of main memory
- A hard disk larger than 10 GB
- Firmware version EFI 1.0.2 or 1.1.0 (preferred)

Our testing system was one pre-production x450 unit with:

- One Itanium 2 processor running at 900 MHz
- Two GB of DDR RAM
- EFI Version 1.10
- Firmware Version 0.9.19

We also connected an RXE-100 expansion enclosure after the installation.

Before starting the installation

Before starting the installation, we recommend you do the following:

- Disconnect the RXE-100 expansion enclosure if it is connected to the system.

  Our testing showed that the RXE-100 may not be recognized by the kernel during installation and the installation process may hang. Once the system is installed, you can reconnect it and all your PCI devices will be recognized automatically.

- Disconnect the LS-120 if you have one installed.

  Do not start the installation procedure with an LS-120 drive unit in the spare media bay in the front of the x450. We encountered problems trying to install the preview release of SuSE SLES 8 that disappeared when we removed the unit from the bay (in our machine, the LS-120 drive was on the secondary bus of the primary IDE interface).

SuSE SLES 8 supports LVM volumes during the installation, so you can install it on LVM partitions. We used only internal disks for our tests. Notice that the default file system used on SuSE is the journaling file system ReiserFS. For more information about the ReiserFS file system, go to: http://www.namesys.com/.

Installation procedure

The installation procedure of SuSE Linux Enterprise Server Version 8 is like any other SuSE Linux installation, except that you work on a different architecture, so

the bootable CD-ROM will not boot directly, as it would do on any IA-32 architecture. These are the steps to follow for a SuSE SLES 8 installation:

1. Power on the server and immediately insert the first SuSE CD-ROM in the drive.

   Although it is not necessary to insert the CD at this point, it makes the installation process easier because the CD-ROM file system is automatically detected and mapped.

2. Wait for the EFI Boot Manager menu to appear (Figure 4-16).

3. Select the entry for the CD-ROM drive. In the x450, this is the ACPI entry with the strings “PCI” and “ATA(Primary/Master)”. Press Enter.

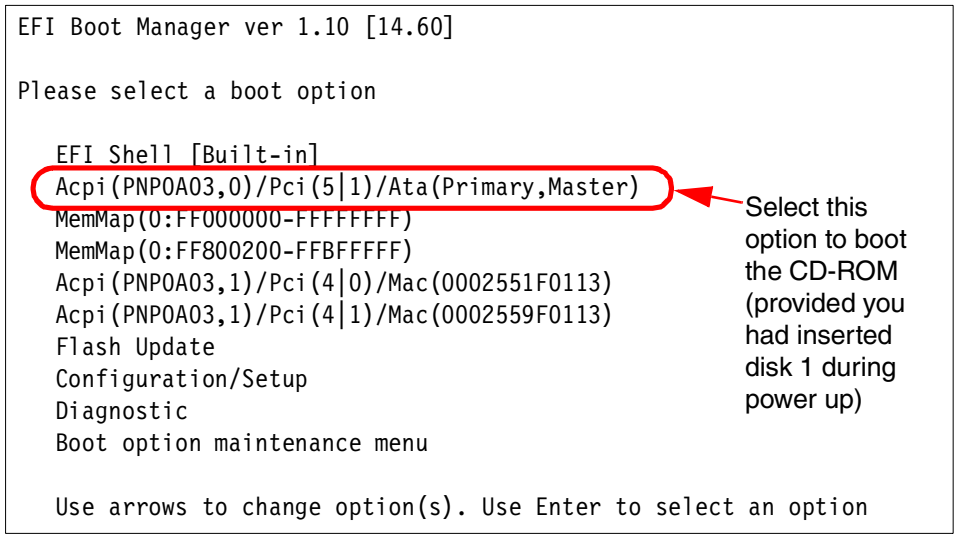

Figure 4-16 EFI Boot Manager — select this option to boot from the CD-ROM

Manually starting the CD-ROM: If you did not have the CD-ROM inserted before the EFI Boot Manager menu loaded, do the following:

1. Insert SuSE Disk 1 into the drive.
2. Select **EFI Shell** and press Enter.
3. Enter the command **`map -r`** to remap all file systems.
4. Enter **`exit`** to go back to the Boot Manager menu.
5. Select **Boot option maintenance menu** and press Enter.
6. Select **Boot from a file** and press Enter.
7. Select the CD-ROM entry **Ata(Primary,Master)/CD**.
8. Browse to the efi/boot folder.
9. Select file **bootia64.efi** and press Enter.

4. The SuSE welcome screen appears.

   The installation procedure is like any other version of SuSE Linux. This will let you select your installation method or boot a rescue environment.

5. Select **Linux** and press Enter to begin the installation. This will boot the YaST2 software that will help you to select the packages that you may want to install. It will also help you with the hardware configuration and then will install all the packages.

Installation

To configure your system and select your packages, here are the steps you will encounter while installing SuSE SLES 8:

1. License Agreement.
2. Language Selection. English is the default.
3. Installation Mode Selection, where you can choose between different installation methods. The default installation should be sufficient. If you need to add or remove packages, you can always do it later.
4. Installation Settings, where you can select what you want to install and where.
5. Package Installation, that is the part of the installation that copies the files from the installation CDs and installs the packages.
6. Prepare for the initial boot, that is the part of the installation that runs the configuration scripts and installs the boot manager.

   Once all your packages are copied to the hard disk, reboot the server.

   Once the system rebooted, on the EFI menu select **SuSE SLES** to start the Linux boot process. This entry will also be configured to be the default entry to boot, so if you do not press any key, the system will automatically boot SuSE SLES.

   You can modify your boot settings EFI boot maintenance menu and the **`elilo.conf`** file as we discuss later.

   The remainder of the installation is completed. The first time SuSE SLES 8 boots, it will finish configuring your system. This will boot the graphical interface and then YaST will be launched. If you encounter problems during this first boot, refer to 4.3.5, “Linux boot process” on page 107.

7. Root password: the installation program will ask you to enter a password for the root user.
8. Add user: the installation program will ask you to create a user. You need to create this user to continue the installation. This is done for security reasons. You should log on as a normal user, and use the root account as little as possible.

9. Desktop settings: you will be prompted to configure your desktop settings and choose which mode, text or graphical, you want to use the next time the machine is booted.
10. Hardware configuration: if you have some additional hardware, such as network interfaces, you can configure them here.

Installation notes: If you install on a mirrored disk configuration while the mirror is being created, the SCSI driver detection on the installation program may take some time. The driver works fine, but it takes longer to load than in a non-mirrored configuration. You may also notice that the I/O throughput slows down.

4.3.5 Linux boot process

The Linux boot process on IA-64 systems uses elilo instead of the classic lilo or grub boot loaders. elilo.conf is a file stored in the root director of the EFI System Partition.

There is no longer a need for boot loaders because the EFI system can manage the boot procedure for the different operating systems. The boot loader is helpful, however, if we need to provide boot parameters to the kernel, for example if we want to recover from some special situations by putting parameters to the boot command line, or we want to boot on initlevel 3 only once and our system is configured to boot on initlevel 5 by default, etc.

If you do not press any key before the delay estimated on the elilo.conf file, Linux will start loading automatically. To reduce the delay or remove it completely, change or remove the corresponding line in elilo.conf.

Example 4-2 shows a typical elilo.conf configuration file. As you can see, the elilo configuration is very similar to the old lilo configuration. One of the main advantages of elilo is that you don't need to execute any command to make your changes active: elilo will automatically take any changes you may make on the elilo.conf file.

Example 4-2 An example of elilo.conf, found in the EFI System Partition root directory

```
timeout=10
read-only

image=vmlinuz
label=linux
root=/dev/sda3
initrd=initdisk
```

```
append="console=ttyS0,115200"
```

Once elilo is loaded, select which Linux kernel image you want to load, and the boot procedure is the same as other Linux startup processes:

1. The Linux kernel image loads on memory.
2. The Setup function is called, enabling basic devices, buses and video.
3. The startup function is called, enabling stacks, decompressing the kernel, and identifying the processor type.
4. The start_kernel function is called, starting all memory procedures, kernel cache and executes /sbin/init.
5. The boot procedure continues.

IA-32 applications

SuSE SLES 8 proved the IA-32 runtime environment installed by default. This environment contains glibc, ncurses, C++ and X11 libraries to guarantee compatibility with IA-32 applications.

If you want to compile any application from code source to IA-32 platform, it is still possible, and you may also compile it to IA-64. This will offer you the possibility to still use older IA-32 applications, if you don't have the source code or you don't want to recompile it to IA-64.

4.3.6 Information about the installed system

Our recently installed Linux system can offer us some information about the x450, such as CPU details, memory, and some other useful information. The following sections outline what the Linux kernel found about the CPU used during our tests. We used the /proc file system to obtain this information.

CPU information

Linux recognizes the x450's processors as an Itanium 2, from the IA-64 architecture. During our tests we only had one single CPU, and it was identified by the system as the CPU number 0. On multiprocessor systems, we will find all the information regarding every CPU.

Example 4-3 /proc/cpuinfo output

```
linux:~ # cat /proc/cpuinfo
processor  : 0
vendor     : GenuineIntel
arch       : IA-64
family     : Itanium 2
```

```
model      : 0
revision   : 6
archrev    : 0
features   : branchlong
cpu number : 0
cpu regs   : 4
cpu MHz    : 900.000000
itc MHz    : 900.000000
BogoMIPS   : 1132.46
```

Memory information

We used 3 GB of DDR memory, and it was fully recognized by Linux. We also set up some swap memory, but it was never used during our tests. Example 4-4 shows the output of the meminfo file on the /proc file system, and more details about the memory.

Example 4-4 /proc/meminfo output

```
linux:~ # cat /proc/meminfo
        total:    used:    free:  shared: buffers:  cached:
Mem:  3112976384 1308147712 1804828672        0 76251136 518733824
Swap: 1044496384        0 1044496384
Total # of HugePages:       32          Available:       32
HugePageSize: 16777216(0x4000KB)
HugePageRegionNumber: 4
MemTotal:      3040016 kB
MemFree:       1762528 kB
MemShared:           0 kB
Buffers:         74464 kB
Cached:         506576 kB
SwapCached:          0 kB
Active:         356800 kB
Inactive:       263072 kB
HighTotal:           0 kB
HighFree:            0 kB
LowTotal:      3040016 kB
LowFree:       1762528 kB
SwapTotal:     1020016 kB
SwapFree:      1020016 kB
```

Modules used by the default kernel

SuSE SLES 8 provides a default kernel that loads almost all the drivers we will need for our x450, and also some drivers we will never use.

As we can see in Example 4-5, it also loads some modules that are not used, identified with `unused`, including:

- parport_pc (the x450 does not provide a parallel port)
- isa_pnp (the x450 does not have any ISA buses)
- keybdev (the x450 does not support a standard keyboard device)
- joydev (for joystick support).

You can unload any unused module with the **rmmod** command. Enter **man rmmod** for more information about rmmod options and parameters.

Example 4-5 /proc/modules output

```
linux:~ # cat /proc/modules
lp                     17112   0 (autoclean)
parport_pc             39960   0 (unused)
parport                67480   0 [lp parport_pc]
videodev               16024   0 (autoclean)
iptable_mangle          4064   0 (autoclean) (unused)
iptable_nat            36320   0 (autoclean) (unused)
ip_conntrack           38816   1 (autoclean) [iptable_nat]
iptable_filter          3272   0 (autoclean) (unused)
ip_tables              29800   5 [iptable_mangle iptable_nat
iptable_filter]
af_packet              34952   1 (autoclean)
bcm5700               175160   1
isa-pnp                83896   0 (unused)
ipv6                  375120  -1 (autoclean)
keybdev                 4416   0 (unused)
hid                    44024   0 (unused)
mousedev               12152   2
st                     63720   0 (autoclean) (unused)
sr_mod                 33232   0 (autoclean)
cdrom                  61504   0 (autoclean) [sr_mod]
sg                     67616   0 (autoclean)
joydev                 13808   0 (unused)
evdev                  10440   0 (unused)
input                   9784   0 [keybdev hid mousedev joydev evdev]
usb-uhci               60088   0 (unused)
usbcore               145016   1 [hid usb-uhci]
nls_iso8859-1           4432   1 (autoclean)
nls_cp437               5936   1 (autoclean)
lvm-mod               123432   0 (autoclean)
ide-scsi               21728   0
reiserfs              495352   1
mptscsih               75456   4
```

```
mptbase                80128   3 [mptscsih]
```

Notice that the Broadcom driver, bcm5700, is used by only one Ethernet card. This is because we were using only one port during our tests and also that SuSE SLES 8 provides LVM support by default.

Partitions on IA-64 Linux

The EFI system implements a new way to work with partitions. That means that the IA-32 tools to manage partitions on Linux may not work properly.

For example, if you run **fdisk** to see the content of a partition or disk, you don't see the real content of the disks. Modifying this partition can destroy the Linux installation.

Example 4-6 shows the fdisk output for one of our test partitions. We had two disks configured on the system: one to run SuSE SLES 8, and the other to run Windows 2003 Server.

Example 4-6 fdisk output of two EFI partitions

```
Disk /dev/sda: 255 heads, 63 sectors, 1106 cylinders
Units = cylinders of 16065 * 512 bytes

   Device Boot    Start       End    Blocks   Id  System
/dev/sda1             1      1107   8887079+  ee  EFI GPT

Disk /dev/sdb: 255 heads, 63 sectors, 1106 cylinders
Units = cylinders of 16065 * 512 bytes

   Device Boot    Start       End    Blocks   Id  System
/dev/sdb1             1      1107   8887079+  ee  EFI GPT
```

The two disks included more than one partition. In fact, there were at least three partitions on each: one for EFI, one for the operating system boot images, and another for the data. As you can see in the example, fdisk doesn't see any of these partitions, but only a big partition of type EFI GP.

Instead of using **fdisk**, use **parted**.

The parted utility is included on the SuSE SLES base system. As shown in Example 4-7 on page 112, the parted output does not show the same information in the partition table, but recognizes the real partitions on the system. The parted utility can identify every partition on the partition table.

Example 4-7 GNU parted output (some lines removed)

```
linux:~ # parted
GNU Parted 1.6.3
...
Using /dev/sda
(parted) print
Disk geometry for /dev/sda: 0.000-8568.000 megabytes
Disk label type: gpt
Minor    Start       End     file system  Name                  Flags
1          0.017     54.906  fat32                              boot
2         54.906   1051.062  linux-swap
3       1051.062   8565.375  reiserfs
(parted)
```

The FAT32 partition is the EFI System Partition where the kernel images and other files needed to boot Linux are stored. The EFI System Partition is also available as the /boot partition once you log into Linux.

For more information about the parted utility, use **man parted** or go to:

http://www.gnu.org/software/parted

4.3.7 Using the serial port for the Linux console

You can redirect all kernel console messages to the serial port. This can be done by entering the following command at the elilo prompt or adding it to the elilo.conf file.

```
append="console=ttyS0,115200"
```

You can also redirect not only kernel messages, but any other system message from the default messages console (which is virtual terminal number 10 on SuSE Linux, and /var/log/messages on other Linux distributions) to the serial port. To do so, you must use the syslog daemon. This daemon controls the flow of system messages and redirects them wherever we want them to be redirected. To redirect all syslog messages to the serial port, edit the /etc/syslog.conf file and add the following string to the line that you want to be redirected to the serial port:

```
/dev/ttyS0
```

Example 4-8 shows an example of this.

Example 4-8 Redirection of kernel messages to the serial port with syslogd daemon

```
# This is only a part of the /etc/syslog.conf file
# Your file may be different. Here we will send all the kernel
```

```
# warning messages to the serial port

kern.warn; /dev/ttyS0
```

To activate the changes you made in the syslog daemon, enter the command:

```
/etc/init.d/syslog restart
```

Refer to the syslog daemon man page for more information about the configuration.

4.3.8 RXE-100 Expansion Enclosure

The RXE-100 Expansion Enclosure is now fully supported under Linux. The x450 will recognize any device attached to any of the 12 RXE-100 slots.

When installing SuSE SLES 8, it is recommended that you disconnect the RXE-100 before the operating system installation, and then reconnect it again once the system is installed. The installation program uses a limited version of the kernel that may not recognize your devices on the RXE enclosure or may crash the installation.

Active PCI support (Hot Plug) is not supported on Linux.

Tip: To easily view all your recognized PCI devices, you can use the YaST tool. To invoke YaST, simply enter **`yast`** at the command prompt, then go to the Hardware and PCI information.

As you can see in Figure 4-17 on page 114, all the PCI slots are recognized, including the 12 external PCI slots.

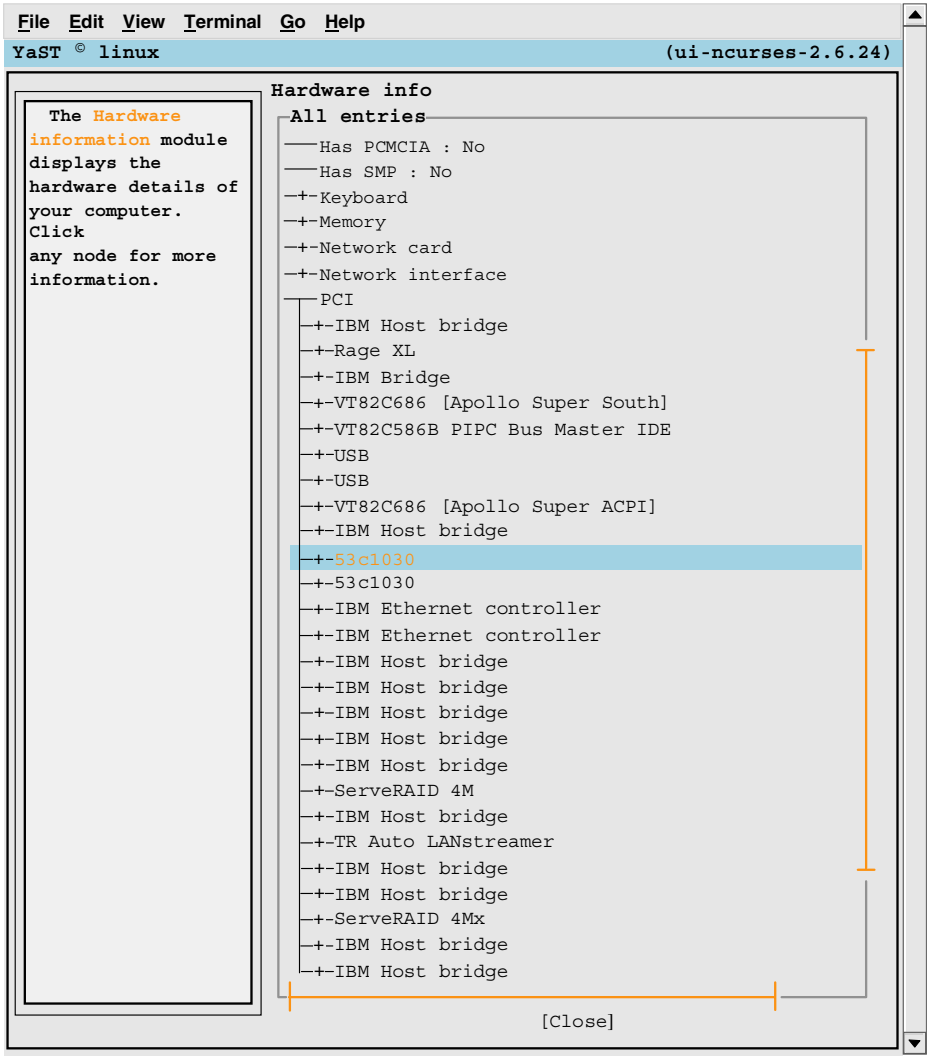

Figure 4-17 PCI information on YaST

4.3.9 Upgrading drivers

Any update to the Linux installation should be done via the official support service provided from IBM and SuSE. To visit the xSeries support page, go to:

http://www.pc.ibm.com/support

The most significant drivers and modules that you should maintain on the x450 are:

- Broadcom Gigabit Ethernet

  The Broadcom Gigabit Ethernet adapter is supported on Linux with the “bcm5700” module. This module supports all the 5700 series of Gigabit adapters from Broadcom. This module comes with SuSE SLES 8 and both Broadcom ports are detected during the installation.

- LSI SCSI chipset

  The default kernel that comes with SuSE SLES 8 has the LSI chipset support precompiled. This will automatically support any SCSI devices connected to your LSI chipset.

- ATI Rage XL video controller

  The video controller used in the x450 is an ATI Rage XL device. This device is supported by XFree86, with the “ati” driver. This driver supports almost all the ATI graphics controllers. ATI controllers are supported on both major versions of XFree86, the 3.3.6 and the 4.2.1. To follow the status of any supported ATI graphic device contsupport by the XFree86 project, see:

  `http://www.xfree86.org/current/Status6.html#6`

  For generic information about the status of any driver, see:

  `http://www.xfree86.org/current/Status.html`

  ATI has also a support page for Linux, where you may find all the information you may need for finding any new drivers or simply know more about ATI’s support to Linux:

  `http://mirror.ati.com/support/faq/linux.html`

- XA-64 chipset

  At the time of the writing, there was no official kernel update to support the new 64-bit chipset. The x450 will work in 64-bit native mode without any kernel patches. If IBM or SuSE release any specific patch for the XA-64 chipset, you will find more information about how to upgrade your current kernel on the xSeries support page:

  `http://www.pc.ibm.com/support`

5

Management

As systems grow in power and complexity, systems management becomes a critical component of a complete system solution. This chapter discusses several techniques and tools that can be used to manage the x450. The topics covered are:

- Configuring the integrated Remote Supervisor Adapter
- Out of band management with IBM Director
- Managing the x450 using the Remote Supervisor Adapter

Important: At the time of writing, a 64-bit version of IBM Director Agent was not available for either Windows or Linux. Consequently, this chapter concentrates on out-of-band alerting and management through the Remote Supervisor Adapter.

5.1 The Remote Supervisor Adapter

The x450 is equipped with a Remote Supervisor Adapter integrated into the system. See Figure 5-1 for the location of the Remote Supervisor Adapter.

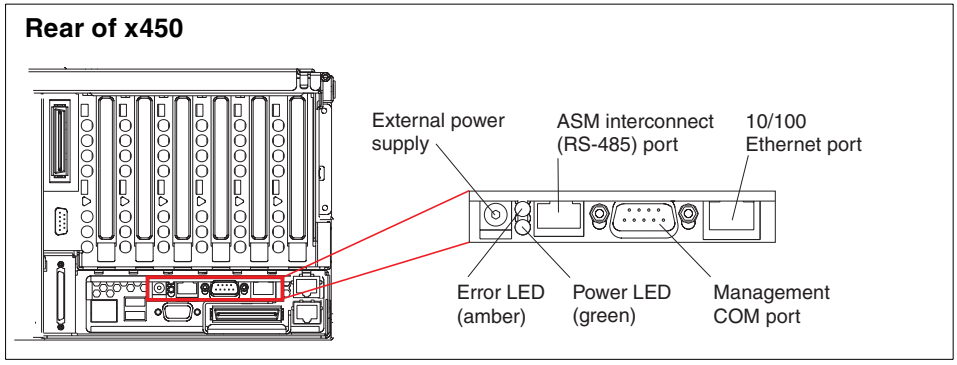

Figure 5-1 Location of Remote Supervisor Adapter in the x450

The features of the Remote Supervisor Adapter include the following:

- Continuous health monitoring and control
- Advanced Predictive Failure Analysis (PFA)
- Configurable notification and alerts
- Event logs that are time stamped, saved in nonvolatile memory, and can be attached to e-mail alerts
- Remote graphical console redirection
- Remote access either via Ethernet, serial, or Advanced System Management (ASM) interconnect connectivity
- Point-to-point protocol (PPP) support
- Simple Network Management Protocol (SNMP) support
- Domain Name System (DNS) and Dynamic Host Configuration Protocol (DHCP) support
- Remote power control
- Remote firmware update and access to critical server settings

5.1.1 The Remote Supervisor Adapter

The functions and information made available by the Remote Supervisor Adapter can typically be accessed either in-band via IBM Director (via IBM Director Agent installed on the server), or via one of the out-of-band interfaces.

In-band means by means of the operating system and device drivers installed on the server. Out-of-band means bypassing the operating system and connecting to the adapter directly.

Important: At the time of writing, a 64-bit version of IBM Director Agent was not available for either Windows or Linux. Consequently, this chapter concentrates on out-of-band alerting and management through the Remote Supervisor Adapter.

There are three physical connections to the adapter for out-of-band management (that is, management without going through the operating system on the server), as shown in Figure 5-1 on page 118:

- Ethernet port, supporting TCP/IP connections
- Serial (COM) port, supporting null modem or PPP connections
- RS-485 ASM interconnect port, for connection to other service processors

The adapter supports three user interfaces for out-of-band management:

- Web browser
- Telnet
- ANSI terminal

The recommended and usually the easiest ways to perform out-of-band management are:

- Using a Web browser to connect to the adapter via its Ethernet port
- Connecting to the adapter from a Remote Supervisor Adapter installed in another server, using the RS-485 ASM interconnect port

5.1.2 Connecting via a Web browser

To use a Web browser to access the Remote Supervisor Adapter remotely, ensure that your browser meets the following requirements:

- Microsoft Internet Explorer 4.0 (with Service Pack 1) or later
- Netscape Navigator 4.72 or later (Version 6.x is not supported)
- Java-enabled Web browser
- Support for JavaScript 1.2 or later
- Minimum display resolution of 800 x 600 pixels and 256 colors

Note: Neither the Web interface nor the text-based interface support double-byte character set (DBCS) languages.

By default, TCP/IP on the Remote Supervisor Adapter is as follows:

- DHCP enabled
- Host name set to `ASMA` followed by the adapter's burned-in MAC address

To initially configure the adapter, you will need to know either the host name or the IP address of the adapter. Since both of these items are initially unknown, follow these steps to determine their values from the System Log:

Tip: The following assumes you have a DHCP server configured. If you do not, see the instructions in 5.1.3, "Configuring a static IP address" on page 121.

1. Connect an active Ethernet connection to the Ethernet port of the adapter.
2. Power on the server.
3. From the EFI Boot Manager menu (Figure 4-1 on page 69), select **Configuration/Setup**.
4. Select **System Event and Error Log** and press Enter.
5. Scroll through the event logs until you see an entry similar to the following:

   ```
   ENET[1] DHCP-HstNme=450rsa,DN=itso.ral.ibm.com,IP@=9.24.104.109,
   GW@=9.24.104.1, NMsk=255.255.254.0,DNS1@=9.24.106.15,DNS2@=9.37.0.5
   ```

6. From this information, you can connect to the adapter via a Web browser. In our example, the URL would be:

   ```
   http://450rsa.itso.ral.ibm.com
   ```

7. When prompted, enter a valid user ID and password. By default, only one user ID is configured:

   User ID: USERID (all caps)
   Password: PASSW0RD (all caps and zero, not the letter O)

   Tip: We strongly recommend you disable this default user ID and create new more secure user IDs before you put your server into production. Access to the Remote Supervisor Adapter should be carefully restricted.

8. Once you are connected, the main menu should appear.

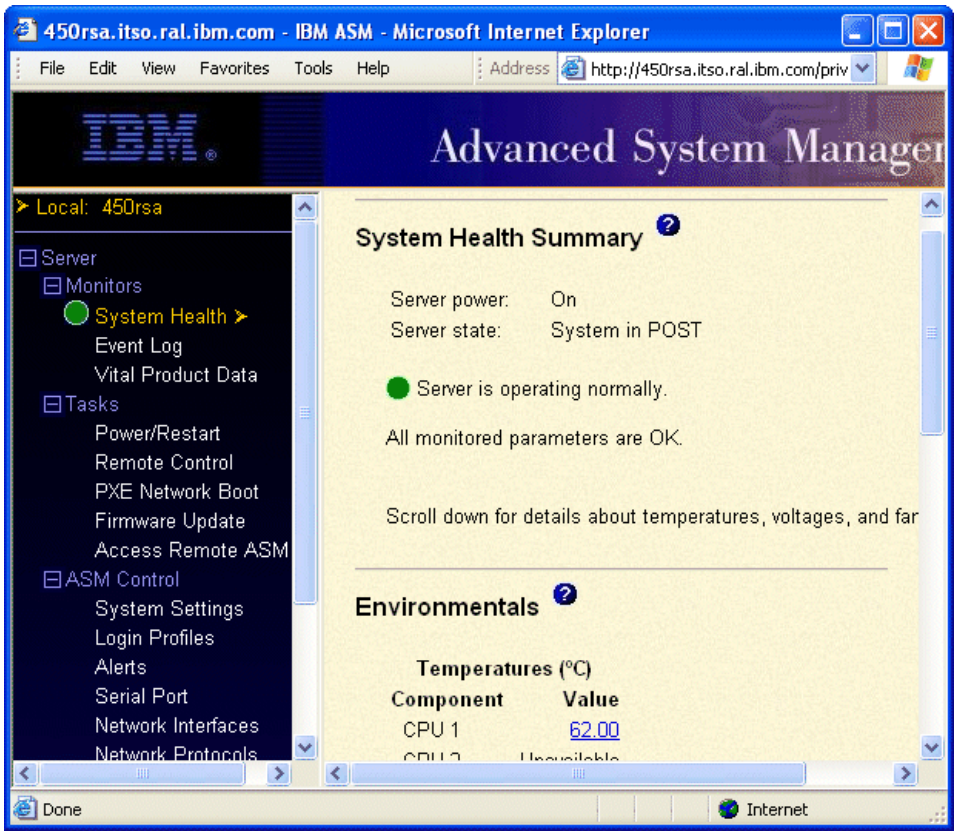

Figure 5-2 Remote Supervisor Adapter Web interface

9. You should change the adapter name (via System Settings) and host name (via Network Interfaces). We recommend you make these two names the same to eliminate confusion.
10. We also recommend that you configure a static IP address for your Remote Supervisor Adapter. This will ensure the adapter is still accessible from the network even when the DNS and DHCP servers are unavailable. This is described in the next section.

5.1.3 Configuring a static IP address

If you do not have a DHCP server on the network, you will need to configure the adapter to use a static IP address.

The easiest way is to connect to the x450's adapter using the ASM interconnect bus. You can then change the TCP/IP setting via the Network Interfaces menu

using a Web browser. See 5.1.4, “Connecting via the ASM interconnect” on page 124.

However, if another Remote Supervisor Adapter is not available, or close enough to be connected via the ASM interconnect, then the next easiest way is to use the EFI-based flash/configuration utility FLASH2.EFI.

Download this utility as part of the latest x450 Remote Supervisor Adapter firmware update from:

`http://www.pc.ibm.com/support`

Once you have downloaded it, you will need to copy it to a USB memory key or a CD-R, or copy it directly to the EFI system partition. (For Windows, see “Accessing EFI System Partition from Windows” on page 96. For Linux, see “Partitions on IA-64 Linux” on page 111.) Only the FLASH2.EFI file is needed to configure the adapter.

To configure a static IP address using FLASH2, perform the following steps:

1. Boot the x450 into the EFI shell.
2. Place the flash2 application on media that can be read by the x450. This could be a CD, memory key, or the application can be copied to a partition on a hard disk drive (assuming a partition has been created).
3. From the EFI shell prompt, navigate to where you have placed the flash2 application (enter `fs2:` or `fs3:`, for example, to select the file system that has been mapped to the system partition).
4. Enter the command **`flash2`** and press Enter.
5. Select **Configuration Settings**.
6. Select **Ethernet Settings**. The Ethernet Settings window opens, as shown in Figure 5-3 on page 123.

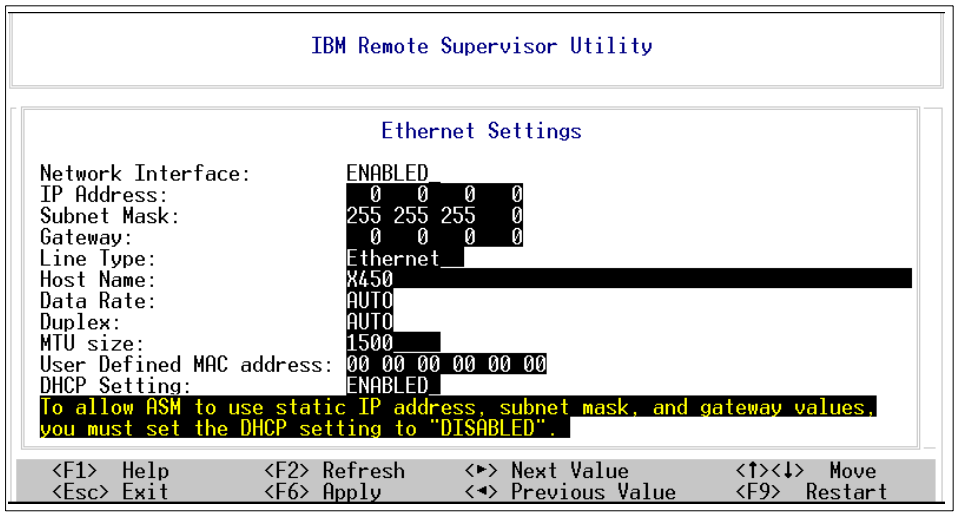

Figure 5-3 RSA Ethernet settings

7. Enter the values specific to your organization for the IP address, subnet mask, and gateway.
8. Enter your host name. To eliminate confusion, consider making this field the same value as the system name under **General**.

   For detailed information about the Ethernet settings, see the *Remote Supervisor Adapter User's Guide*.
9. Change the DHCP Setting field to Disabled. The default value is Enabled.
10. Press F6 to apply the changes.
11. Press Esc. After the Warning window opens, press Enter to restart the Remote Supervisor Adapter.
12. Exit the utility and return to the EFI Boot Manager menu.

You should now be able to connect to the adapter using the static IP address you supplied.

Tip: For information on how to configure other connectivity options using the serial port (null modem or PPP), see the *Remote Supervisor Adapter User's Guide*.

5.1.4 Connecting via the ASM interconnect

In addition to the Ethernet interface, the ASM interconnect is often used to manage the Remote Supervisor Adapter when you have other service processors in other xSeries servers.

To connect the Remote Supervisor Adapter in the x450 to an ASM interconnect bus, you will need:

- Advanced System Management Interconnect Cable Kit, part 03L9309
- Standard Cat5 Ethernet cable (not crossover)

Follow the instructions that ship with the cable kit, or review the redbook *Implementing Systems Management Solutions using IBM Director*, SG24-6188.

Once you have connected the x450 Remote Supervisor Adapter to other service processors including another Remote Supervisor Adapter, do the following:

1. From your Web browser, connect to the other Remote Supervisor Adapter (for example, one in an x440).
2. Select **Access Remote ASM** from the main menu (Figure 5-2 on page 121).
3. A list of the other service processors connected on the ASM interconnect bus will be displayed.

> **Tip:** It can take up to about a minute (the online help says 45 seconds) before a new service processor connected to the interconnect bus appears on this page.

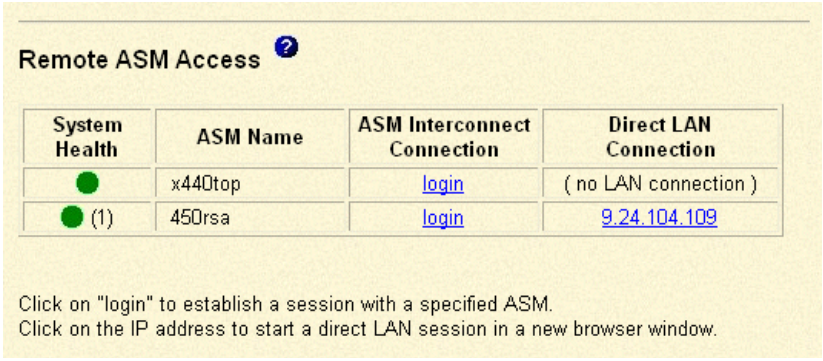

Figure 5-4 Connecting remotely to the Remote Supervisor Adapter in the x450

4. Next to the x450, (host name 450rsa in our example), click **login**. If you click the IP address in that row, it connects to the adapter via its Ethernet port.

5. When prompted, enter a valid user ID and password. By default, only one user ID is configured:

   User ID: USERID (all caps)
   Password: PASSW0RD (all caps and zero, not the letter O)

   Tip: We strongly recommend you disable this default user ID and create new more secure user IDs before you put your server into production. Access to the Remote Supervisor Adapter should be carefully restricted.

6. Once you are connected, the main menu of the x450 Remote Supervisor Adapter should appear.

5.1.5 Installing the device driver

You should install the ASM device driver once the operating system is installed. The device driver enables communication between the adapter and the installed operating system (Windows or Linux). In particular, the driver provides:

- The ability to gracefully shut down the server
- Indications as to whether the operating system is booted
- The ability to determine if the operating system is hung and automatically restart the system

The device driver is on the Resource CD-ROM that is shipped with the server. The latest version of the driver is also available from:

`http://www.pc.ibm.com/support`

Installation procedures are documented in the following sections:

- Windows: “Installing additional drivers” on page 93
- Linux: 4.3.9, “Upgrading drivers” on page 114

5.1.6 Configuring the remote control password

If you want to use the remote control features of the Remote Supervisor Adapter Web interface, you must configure the remote control password using the Configuration/Setup utility program.

To set up a remote control password, complete the following steps:

1. Start the server.
2. From the EFI Boot Manager startup options, select **Configuration/Setup** and press Enter.
3. Select **System Security**.

4. Select **Remote Control Security**.
5. In the Enter New Remote Control Password field, enter the Remote Control password that you wish to use each time you log into the remote control features of the Remote Supervisor Adapter Web Interface.
6. Select **Change Control Password** and then press Enter.

5.2 Management using the Remote Supervisor Adapter

Once you have configured the Remote Supervisor Adapter as described in 5.1, "The Remote Supervisor Adapter" on page 118, you can then use the adapters to manage x450 by forwarding events to IBM Director. IBM Director can then be configured to act upon those events in a variety of ways (for example, by sending e-mail or a pager alert).

Tip: This redbook assumes you are familiar with IBM Director. For information on how to configure and use IBM Director, see the redbook *Implementing Systems Management Solutions using IBM Director*, SG24-6188.

The Remote Supervisor Adapter supports the following methods of forwarding alerts to IBM Director:

- Simple Network Management Protocol (SNMP), as described in 5.2.2, "Configuring SNMP" on page 127
- Communication directly with IBM Director, either via Ethernet or via modem, as described in 5.2.3, "Sending alerts directly to IBM Director" on page 128

Tip: While all methods are supported, it is recommended that only one method be utilized. Otherwise, you may have the situation where a single event on the x450 triggers multiple events in IBM Director.

5.2.1 Configuring which alerts to monitor

Before configuring the method of sending the alert to IBM Director, you should first define which events the Remote Supervisor Adapter is to send.

To enable alerting on the Remote Supervisor Adapter, do the following:

1. From the Web browser main menu, click **Alerts** in the navigation bar.

2. Scroll to the Monitored Alerts section. There are three types of events listed, as follows:
   - Critical Alerts (these are hard disk drive, multiple fan failure, power failure, tamper, temperature, voltage, and VRM failure)
   - Warning Alerts (these are single fan failure, temperature, voltage, and redundant power supply)
   - System Alerts (these are boot failure, loader timeout, O/S timeout, PFA, POST timeout, power off, and power on)
3. Select the events you wish to monitor. We recommend you enable all events.
4. Click **Save**.

Tip: Below the Monitored Alerts section is the Monitored Local Alerts section. These options are applicable only when the IBM Director Agent is installed on the x450. When the agent is installed and those options are selected, the events are sent to IBM Director agent and appear as sourced from the x450.

5.2.2 Configuring SNMP

The Remote Supervisor Adapter supports the sending of alerts to IBM Director using the Simple Network Management Protocol (SNMP).

Before alerts can be sent, you need to configure the adapter for SNMP so that IBM Director Server can discover and manage it.

Note: We assume you have already set up a system running IBM Director Server. If you need help in configuring IBM Director, please see the documentation that comes with IBM Director or the redbook *Implementing Systems Management Solutions using IBM Director*, SG24-6188.

Configure SNMP as follows:

1. If you have not already done so, configure the adapter so that it can be accessed from a Web browser, either through a direct Ethernet connection or via an ASM interconnect. See 5.1.2, “Connecting via a Web browser” on page 119 or 5.1.4, “Connecting via the ASM interconnect” on page 124.
2. Connect to the adapter using the Web browser. From the home page, click **Network Protocols**.
3. In the SNMP section (see Figure 5-5 on page 128), enter the following information:

   SNMP agent Enabled (Default is Disabled)

SNMP traps Enabled (Default)

Community name Set you your SNMP community name (Default is public for IBM Director)

Host Name or IP address Host name or IP address of the system running IBM Director Server

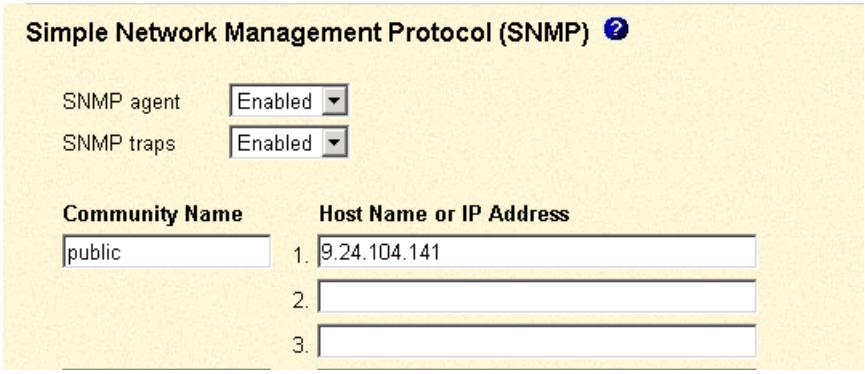

Figure 5-5 Remote Supervisor Adapter — SNMP configuration

4. Once all values are entered, click **Save**. You will be prompted that you need to restart the Remote Supervisor Adapter to bring the changes into effect.
5. In the left-hand navigation pane, click **Restart ASM**.
6. Perform a discovery at the IBM Director Console. IBM Director should detect the Remote Supervisor Adapter in the x450 as an SNMP device.

   By default, the name listed in IBM Director is “Remote Supervisor Adapter <IP address>” where IP Address is the address of the Remote Supervisor Adapter you configured. You can choose to make this name more descriptive for your environment by right-clicking and choosing **Rename**.
7. Proceed to 5.2.4, “Creating a test event action plan in IBM Director” on page 130 for instructions on how to verify the functionality of the Remote Supervisor Adapter.

5.2.3 Sending alerts directly to IBM Director

The Remote Supervisor Adapter can be configured to send alerts directly to IBM Director Server via one of these ports on the adapter:

- Ethernet port
- COM port and an attached modem
- ASM interconnect connection to another service processor that has an Ethernet connection or modem attached

The ASM interconnect bus transparently handles the transfer of all outgoing alerts to Ethernet connections or modems on any of the service processors connected together. In this way, for example, as long as one of the connected service processors has an Ethernet connection or modem, then all alerts from the Remote Supervisor Adapter in the x450 will be sent without any extra configuration required.

Important: Sending alerts directly to IBM Director, while convenient, may cause difficulties processing incoming events at the IBM Director Server. This is because the events that are sent directly from the x450 Remote Supervisor Adapter are reported as being sourced from the IBM Director Server itself, instead of being sourced from the x450.

Once the Director Agent is available, alerts can be sent through it and event processing will be easier.

To configure the Remote Supervisor Adapter to send alerts to IBM Director Server via Ethernet, perform the following steps:

1. If you have not already done so, configure the Ethernet adapter settings on the adapter. See 5.1.2, “Connecting via a Web browser” on page 119 for details on how to complete this configuration.
2. Connect to the Remote Supervisor Adapter using a Web browser. From the home page, click **Alerts**.
3. The upper portion of this Web page lists those recipients configured for remote alerts. By default, this list is empty. Click on an empty link (labeled ~not used~) to configure a remote recipient.
4. To configure IBM Director as a remote recipient, you will need to enter the following values. All fields not listed can be left in their default (blank) state:

Status	Enabled (default is Disabled)
Name	Enter a descriptive name
Notification Method	IBM Director over LAN or IBM Director over modem
Number	IP Address of the system running IBM Director Server

5. Click **Save** to save your configuration. You should see your recipient listed in the list.

Tip: At time of writing, clicking the **Generate Test Alert** button did not properly transmit an alert to IBM Director when the **IBM Director over LAN** method was selected.

5.2.4 Creating a test event action plan in IBM Director

Once you have configured the Remote Supervisor Adapter to communicate with IBM Director, we recommend you test the functionality to ensure proper communication.

For more information on alerts and event management in IBM Director, see the redbook *Implementing Systems Management Solutions using IBM Director*, SG24-6188.

To verify the functionality of the Remote Supervisor Adapter in the x450, perform the following steps:

1. Launch the IBM Director Console, either directly from the IBM Director Server system or from a workstation with the IBM Director Console installed.
2. Click the **Event Action Plan Builder** icon located in the IBM Director Console.
3. The Event Action Plan Builder window will appear. Right-click **Event Action Plan** in the left-hand pane and select **New** to create a new plan.
4. Provide a descriptive name for your plan and click **OK**. For our testing, we entered **`RSA Test`**.

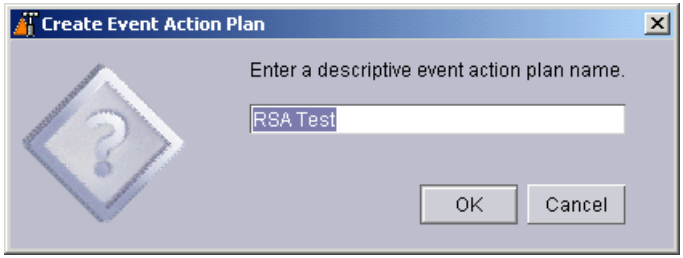

Figure 5-6 New event action plan

5. Once you click **OK**, your new event action plan will appear in the left-hand pane under Event Action Plans.
6. In the middle pane, click the blue bullet next to **Simple Event Filter**. A menu tree will drop down showing all pre-defined event filters created during installation of IBM Director. Click and drag **All Events** to the event action plan you just created. This will add the filter to the action plan. The All Events filter will cause any event generated by the system to be reported to IBM Director, and is good for testing. For production, you may wish to limit the events sent to IBM Director.

7. In the right-hand pane, right-click **Send an Event Message to a Console User** and choose **Customize**. This action will pop up the Message Console on the console currently being used when the alert is triggered.
8. A new window will appear. For the message, type `&Text`. This will show the error message in the pop-up window. For User(s), type `*`. This will cause the message to appear on all current Console users. See Figure 5-7 for more details.

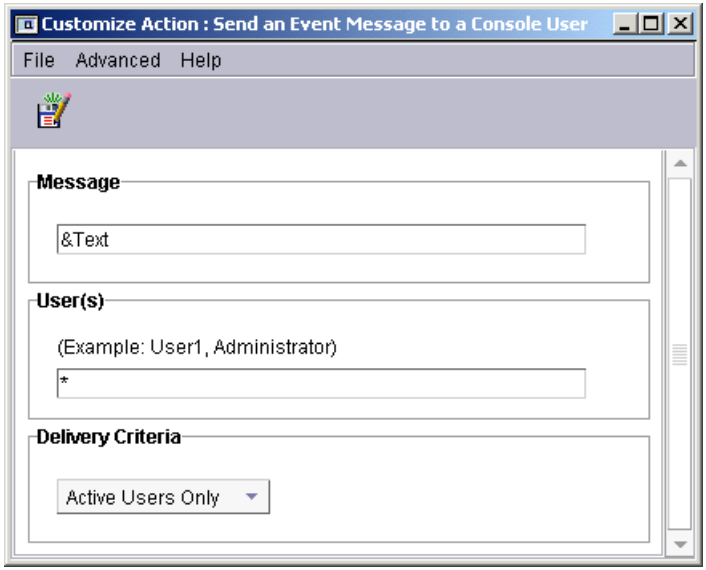

Figure 5-7 Customize action configuration window

9. Click the Save icon in the upper left-hand corner of the window (looks like a diskette) to save the changes.
10. Your new customized action should appear in the right-hand pane. Click and hold the action and drag the icon to the filter you created.
11. Close the Event Action Plan Builder.
12. If you chose to configure the Remote Supervisor Adapter using Alert Forwarding, continue with step 13. If you chose to configure your adapter using SNMP, jump to step 14.
13. In the IBM Director Console, click and drag your test event to your IBM Director Server system in the middle pane. See Figure 5-8 on page 132 for how your configuration should look.

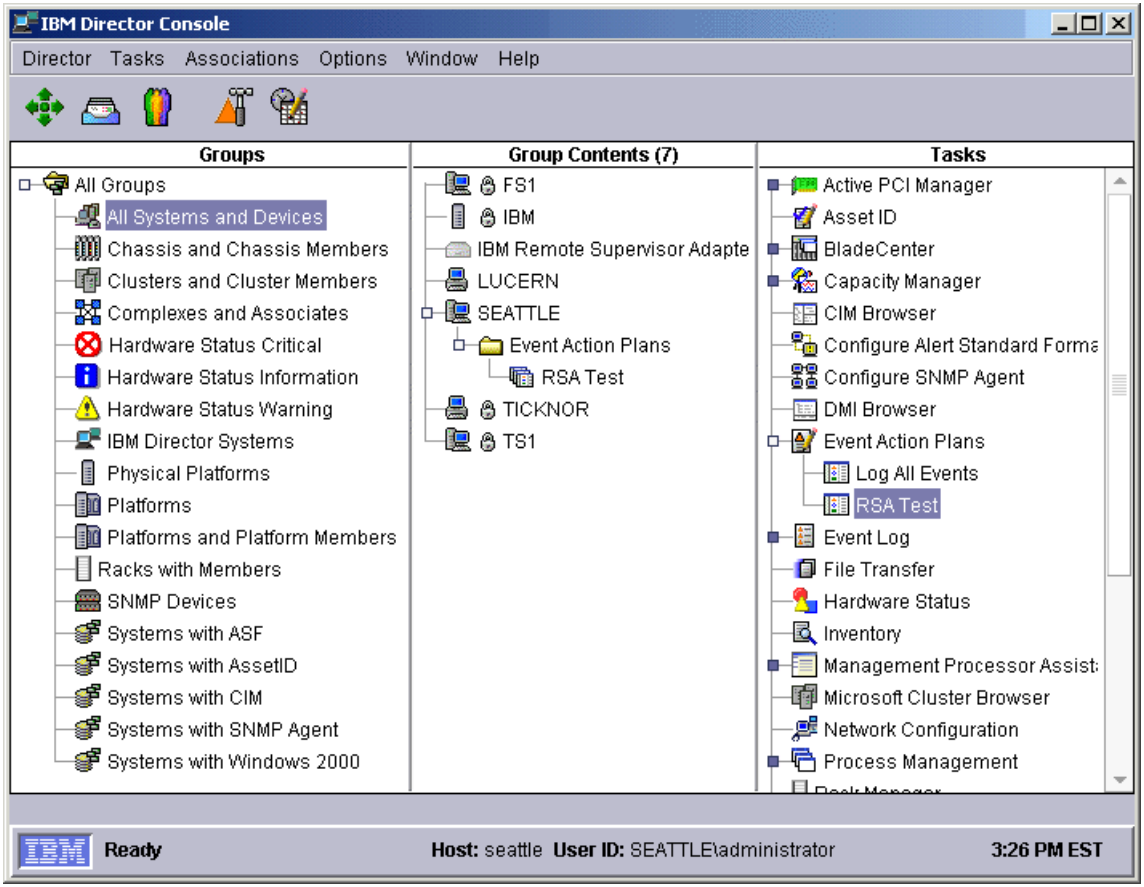

Figure 5-8 Completed configuration

14. In the IBM Director Console, click and drag your completed test configuration to the Remote Supervisor Adapter for the x450 located in the middle pane.
15. In either case, a pop-up window will appear notifying you that the event action plan has been added (Figure 5-9 on page 133).

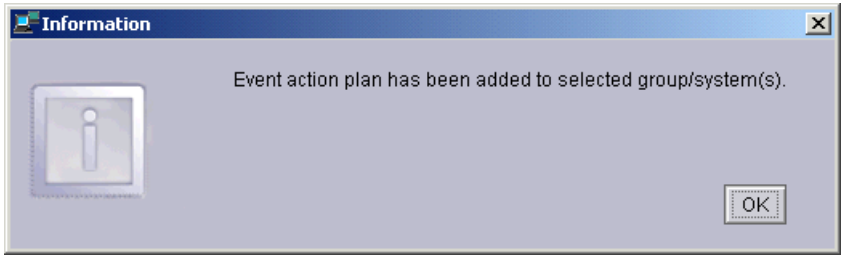

Figure 5-9 Event action plan added

16. To complete testing, we recommend removing one of the fans from the top of the x450. Within approximately 10 seconds, an alert should appear on the Director Console similar to the one in Figure 5-10.

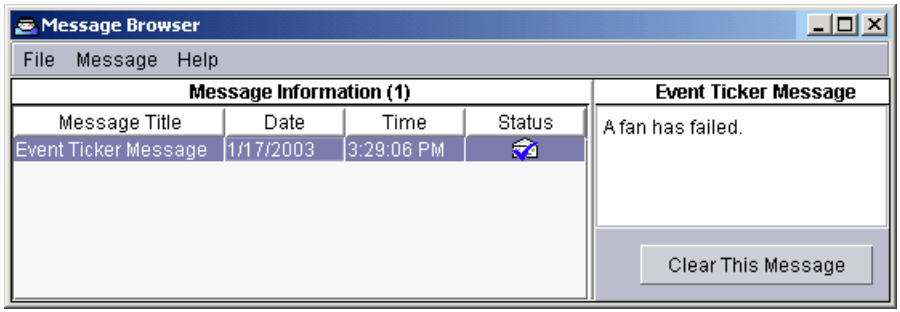

Figure 5-10 Message alert

Once functionality has been confirmed, you can now build event action plans to suit your environment.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

IBM Redbooks

For information on ordering these publications, see “How to get IBM Redbooks” on page 137.

- *Tuning IBM @server xSeries Servers for Performance*, SG24-5287
- *Implementing Systems Management Solutions using IBM Director*, SG24-6188
- *IBM @server xSeries 440 Planning and Installation Guide*, SG24-6196

Referenced Web sites

These Web sites are also relevant as further information sources:

- http://www.pc.ibm.com/us/compat/nos/matrix.shtml
- http://www.intel.com/design/itanium2
- http://developer.intel.com/technology/efi/download.htm
- http://www.intel.com/technology/efi/index.htm
- http://www.microsoft.com/hwdev/platform/firmware/EFI/default.asp
- http://www.intel.com/products/server/processors/server/itanium2/index.htm
- http://www.sap.com/benchmark
- http://www-3.ibm.com/software/data/db2/itanium
- http://www.spec.org
- http://www.ibmlink.ibm.com/usalets&parms=H_202-214
- http://www-3.ibm.com/software/data/Intel/ia64.html
- http://www.pc.ibm.com/us/compat/lan/matrix.html
- http://www.storage.ibm.com/linux
- http://www.ibm.com/pc/us/compat

- http://www.pc.ibm.com/us/compat/x450/ibm_22P6801.html
- http://www.pc.ibm.com/europe/configurators
- http://www.pc.ibm.com/us/eserver/xseries/library/configtools.html
- http://w3.ibm.com/support/assure
- http://www.ibm.com/support/assure/esar
- http://www.pc.ibm.com/support
- http://ibm.com/pc/ww/eserver/xseries/tape.html
- http://ibm.com/pc/ww/eserver/xseries/scsi_raid.html
- http://www.storage.ibm.com/hardsoft/disk/fastt/index.html
- http://www.microsoft.com/hwdq/hcl
- http://www.microsoft.com/windows.netserver/64bit/default.mspx
- http://www.pc.ibm.com/qtechinfo/MIGR-4JTS2T.html
- http://developer.intel.com/technology/efi/diskutil_overview.htm
- http://www.microsoft.com/technet/prodtechnol/winxppro/reskit/prkb_cnc_helm.asp
- http://www.kernel.org
- http://www.dig64.org
- http://www.linuxbase.org/spec/refspecs/elf/gabi4+/contents.html
- http://www.opengroup.org
- http://www.pc.ibm.com/qtechinfo/MIGR-48NT8D.html
- http://www.redhat.com/software/itanium
- http://www.pc.ibm.com/us/compat/nos/redchat.html
- http://www.unitedlinux.com/en/info/faqs.html
- http://www.conectiva.com
- http://www.sco.com/products/openlinux64
- http://www.pc.ibm.com/us/compat/nos/thecscocgroup.html
- http://www.turbolinux.com/products/tls7i
- http://www.pc.ibm.com/us/compat/nos/turbolinux.html
- http://www.pc.ibm.com/us/compat/nos/suse.html
- http://www.namesys.com
- http://www.gnu.org/software/parted
- http://www.xfree86.org/current/Status6.html#6

- http://www.xfree86.org/current/Status.html
- http://mirror.ati.com/support/faq/linux.html

How to get IBM Redbooks

You can order hardcopy Redbooks, as well as view, download, or search for Redbooks at the following Web site:

ibm.com/redbooks

You can also download additional materials (code samples or diskette/CD-ROM images) from that site.

IBM Redbooks collections

Redbooks are also available on CD-ROMs. Click the CD-ROMs button on the Redbooks Web site for information about all the CD-ROMs offered, as well as updates and formats.

Index

Numerics

03K9309, ASM Interconnect Cable Kit 51
22P6801, NetXtreme 1000 T Ethernet adapter 47
22P7801, NetXtreme 1000 SX Fiber Ethernet 47
31P5998, RXE-100 6-slot Expansion Kit 51
31P6087, 3.5 m management cable kit 53
31P6088, 8 m management cable kit 53
31P6102, 3.5 m Remote I/O cable kit 53
31P6103, 8 m Remote I/O cable kit 53
31P6301, NetXtreme 1000T Ethernet adapter 47
32-bit PCI adapters 44
73P2027, 512 MB DDR DIMM 41
73P2028, 1 GB DDR DIMM 41
73P2030, 2 GB DDR DIMM 41
73P5832, USB Conversion Option 49
73P7076, 1.3 GHz processor option 40
73P7077, 1.4 GHz processor option 40
73P7078, 1.5 GHz processor option 40
8687-1RX 2, 40
8687-2RX 2, 40
8687-3RX 2, 40

A

Active Memory 17
  mirroring 19
Active PCI driver
  Windows 2000 94
adapter speeds 45
adapter teaming 47

B

BIOS 6
boot maintenance menu 85
bootstrap processor 16
Broadcom 22, 47
  adapter teaming 47
Broadcom driver
  Windows 2000 94
Broadcom Gigabit Ethernet controller 22

C

cabling
  Ethernet 22, 47
  RXE-100 53
cache 11, 17
Chipkill memory 20
cyclone chip 3

D

DDR memory 17
diagnostics 84
DIMM sockets 14
DIMMs 17
  installation order 41
  maximum 15, 41
  standard memory 5
diskpart command 78
drive bays 1, 24
drvcfg command 75

E

EFI 5, 68
  boot maintenance menu 85
  file systems 73
  flash updates 80
  map command 72
  RSA firmware 82
  shell 70
  system partition 96
  tasks 68
EFI System Partition 8
Enterprise Storage Server 60
Enterprise X-Architecture 1
Ethernet, Gigabit 2, 22
  adapter teaming 47
Explicitly Parallel Instruction Computing 12
Extensible Firmware Interface 5

F

fans 23–24
features 1
Fibre Channel storage options 58
flash updates 80
flash2 command 82

front panel 24
frontside bus 4, 11

G
Gigabit Ethernet 22
GPT 7
GUID Partition Table 7

H
hardware configuration 14–15
hot-swap fans 23

I
IBM XA-64 chipset 3
installing
  in a rack 62
  SuSE Linux 98
  Windows Server 2003 86
Intel PRO/1000 XT 47
Intel Xeon Processor MP 10
interleaving 5

L
LEDs, light path diagnostics 25
Level 3 cache 11
Level 4 cache 17
Light Path Diagnostics
  SMP Expansion Module 15–16
light path diagnostics 25
Linux
  distributions 101
  installing 98
  support 27
low-level format 76
LSI
  configuring 75

M
Madison 2
map command 72
MBR 7
McKinley 2
memory 17
  Active Memory 17
  Chipkill memory 20
  disabling during POST 20
  memory mirroring 19
  Memory ProteXion 18
  memory scrubbing 18
  standard DIMMs 2
memory mirroring 42, 83
memory-board assembly 14, 41
Microsoft Reserved Partition 88
models 2
mountvol command 96

N
NetBurst architecture 4
NetXtreme 47

O
operating systems 27
OS loader 6

P
parallel ports 23
PCI bridges 5
PCI scan order 46
PCI slots 21
  RXE-100 51
PCI-X
  See also RXE-100
PCI-X and PCI coexistence 45
PCI-X slots in the x450 44
performance 43
planning 39–65
power pod 16, 40
power supplies 24
processor-board assembly 15
protective MBR 7

R
rack installation 62
RAID
  configuring 75
RAID-1 76
rear panel 5
Red Hat Linux 102
Redbooks Web site 137
  Contact us xii
redundancy 23
Remote Expansion I/O port 5
Remote Supervisor Adapter
  connectivity 50

driver
Windows 2000 94
features 26
firmware updates 82
RS-485 23
rules
PCI slots 44
RXE-100 52
RXE-100 51, 113
cabling 53
slots 52

S

scan order, PCI 46
scrubbing, memory 18
SCSI
configuring 75
ports 22
SDRAM 5, 17
serial ports 23
shark 60
SMP Expansion Module 13
LEDs 15–16
standard memory 2
Summit chipset 3
support
operating systems 27
SuSE Linux
installing 98
support 27

T

teaming 47
tornado chip 3
TotalStorage solutions 58

U

Ultra 160 SCSI ports 22
UnitedLinux 102
USB ports 23

V

video controller 22
video PCI adapters 46

W

Windows Server 2003
Device Manager 94
installing 86
RXE-100 95
support 27
winnipeg chip 3

X

XA-64 chipset 3
X-Architecture 1
XceL4 Server Accelerator Cache 17
Xeon Processor MP 10

Z

ZIF socket 10

IBM Redbooks

IBM @server xSeries 450 Planning and Installation Guide

(0.2”spine)
0.17”<->0.473”
90<->249 pages

IBM @server xSeries 450 Planning and Installation Guide

Describes the technical details of the new 64-bit server

Covers Windows Server 2003 and SuSE Linux Enterprise Server

Helps you prepare for and perform an installation

The IBM @server xSeries 450 is IBM's new 64-bit Itanium Processor Family (IPF) Architecture server and is the first implementation of the 64-bit IBM XA-64 chipset, as part of the Enterprise X-Architecture strategy. This IBM Redbook is a comprehensive resource on the technical aspects of the server, and is divided into five key subject areas:

- Chapter 1, "Technical description" introduces the server and its subsystems and describes the key features and how they work.
- Chapter 2, "Positioning" examines the types of applications that would be used on a server such as the x450.
- Chapter 3, "Planning" describes the considerations when planning to purchase and planning to install the x450.
- Chapter 4, "Installation" covers the process of installing Windows Server 2003, Enterprise Edition and SuSE Linux Enterprise Server on the x450.

Chapter 5, "Management" describes how to use the Remote Supervisor Adapter.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-6927-00 ISBN 0738429465