

The Lenovo logo is displayed in white text on a black rectangular background.

Lenovo System x3750 M4 Implementation Guide

Last Update: September 2016

**Detailed technical information
about the server and options**

**Guidance on how to maximize
performance in all subsystems**

**IMM2 and UEFI key functions and
options explained**

**Suitable for IT specialists looking to
implement this 4-socket 2U rack
server**

David Watts
Shelly Faith
Miroslav Peic



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

Last update on September 2016

This edition applies to:

- ▶ Lenovo System x3750 M4, machine type 8732 with Intel Xeon Processor E5-4600 v2
- ▶ System x3750 M4, machine type 8752 with Intel Xeon Processor E5-4600 v2
- ▶ System x3750 M4, machine type 8722 with Intel Xeon Processor E5-4600

© Copyright Lenovo 2016. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract

Contents

Notices	iv
Trademarks	v
Preface	vi
The team who wrote this paper	vi
Comments welcome	vii
Do you have the latest version?	viii
Chapter 1. Product overview	1
1.1 Introduction	2
1.2 Key features	2
1.2.1 Scalability and performance	3
1.2.2 Availability and serviceability	3
1.2.3 Manageability and security	4
1.2.4 Energy efficiency	5
1.3 Locations of key components and connectors	6
1.4 Standard specifications	8
1.5 Standard models	11
1.6 Architecture	12
1.7 Processors	13
1.7.1 Processor options	15
1.7.2 Processor failure and recovery	16
1.8 Memory	17
1.8.1 Memory types	18
1.8.2 Memory options	18
1.8.3 Memory performance	20
1.8.4 Memory operational modes	21
1.8.5 Population sequence	22
1.9 Internal disk storage options	27
1.9.1 Backplanes	27
1.9.2 Backplane combinations	30
1.9.3 RAID controllers	40
1.9.4 RAID controller upgrades	42
1.9.5 Battery mounting	45
1.9.6 Battery versus flash	46
1.9.7 Drive selection	46
1.9.8 Drive options	47
1.10 Optical drive	48
1.11 PCI Express slots	48
1.12 Integrated networking: Type 8753 and 8752	50
1.13 Integrated networking: Type 8722	51
1.14 Network adapters	53
1.15 Power supplies	53
1.16 Integrated virtualization	55
1.17 System management	56
1.17.1 Light path diagnostic panels	56
1.17.2 Integrated Management Module II	57
1.17.3 Remote management	58
1.18 Physical and electrical specifications	58

Chapter 2. Initial setup of the IMM2	61
2.1 IMM2 overview	62
2.2 Setting up IMM2 for remote access	63
2.2.1 Network interface	63
2.2.2 Methods to access: Web and command-line interface	65
2.2.3 IMM2 web interface quick overview	68
2.2.4 Remote presence	73
2.3 Initial configuration of IMM2	76
Chapter 3. Initial setup of UEFI and hardware components	77
3.1 System Settings in UEFI	78
3.1.1 Boot Manager menu	78
3.1.2 System Event Log	81
3.1.3 User Security	81
3.1.4 Default UEFI settings and Advanced Settings Utility values	82
3.1.5 Operating modes	83
3.1.6 Settings for performance of individual subsystems	86
3.1.7 Optimizing boot time	89
3.2 Configuration to enable memory mirroring or sparing	90
3.2.1 Memory mirroring	91
3.2.2 Memory sparing	92
3.3 Configuring the RAID adapter by using UEFI	93
3.3.1 Controller Management	95
3.3.2 Virtual Drive Management	100
3.3.3 Drive Management	105
3.3.4 Enclosure Management	106
3.4 Configuring embedded Ethernet using UEFI	106
Chapter 4. Installation and management	115
4.1 Firmware considerations	116
4.1.1 Firmware update tools and preferred practices	116
4.1.2 Fix Central	116
4.1.3 Bootable Media Creator	118
4.1.4 UpdateXpress System Pack Installer	122
4.1.5 Updating firmware through the IMM2	125
4.2 Supported operating systems	126
4.3 Installing an operating system	127
4.3.1 Use of ServerGuide	129
4.3.2 Installation instructions	129
4.3.3 Installing Windows	129
4.3.4 Installing VMware	130
4.3.5 Installing Linux	130
4.4 Post-installation management tools	131
4.4.1 IMM2 email notifications	131
Abbreviations and acronyms	135
Related publications	137
Lenovo Press publications	137
Other publications and online resources	137

Notices

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

Lenovo 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:

Lenovo (United States), Inc.
1009 Think Place - Building One
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO 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 jurisdictions 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. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

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

Any references in this publication to non-Lenovo 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 Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Trademarks

Lenovo, the Lenovo logo, and For Those Who Do are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. These and other Lenovo trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by Lenovo at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of Lenovo trademarks is available on the Web at <http://www.lenovo.com/legal/copytrade.html>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Advanced Settings Utility™	Lenovo®	System x®
BladeCenter®	Lenovo XClarity™	ThinkServer®
Bootable Media Creator™	Lenovo (logo)®	ToolsCenter™
Dynamic System Analysis™	ServeRAID™	UpdateXpress System Packs™
eXFlash™	ServerGuide™	vNIC™
Flex System™	ServerProven®	

The following terms are trademarks of other companies:

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

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

ActiveX, Internet Explorer, Microsoft, Nina, Windows, Windows Server, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

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

Preface

The Lenovo® System x3750 M4 is a 4-socket server that features a streamlined design that is optimized for price and performance with best-in-class flexibility and expandability. Models of the x3750 M4 are powered with the Intel Xeon E5-4600 v2 processors for an entry-level, 4-socket solution. The x3750 M4 provides maximum storage density, with flexible PCIe and 10 Gb Ethernet networking options in a 2U form factor.

The x3750 M4 is ideal for technical computing workloads with floating-point computations, small to medium databases that require fast I/O, and applications that require 4-socket performance. This system is for workloads that do not need the extra scalability that the X6 systems provide.

This paper introduces the server and provides technical details about all hardware subsystems. It explains how to set up the Integrated Management Module II (IMM2) for remote access. It then describes important setup tasks in the Unified Extensible Firmware Interface (UEFI), including disk and network configuration. This paper also describes ways to update firmware and includes pointers on how to install the supported operating systems.

This paper is intended for clients, Lenovo Business Partners, and Lenovo employees that want to learn more about the System x3750 M4 and how to prepare it to run high-performance applications.

The team who wrote this paper

This paper was produced by a team of specialists from around the world:

David Watts is a Senior IT Consultant and the program lead for Lenovo Press. He manages residencies and produces pre-sale and post-sale technical publications for hardware and software topics that are related to System x®, ThinkServer®, Flex System™, and BladeCenter® servers. He has authored over 300 books and papers. David has worked in the IT industry, both in the U.S. and Australia, since 1989, and is currently based in Morrisville, North Carolina. David holds a Bachelor of Engineering degree from the University of Queensland (Australia).

Shelly Faith is a System x and BladeCenter Technical Support Specialist in Atlanta, GA, providing remote technical support for IBM clients worldwide. She has five years of experience with System x support and worked for IBM for four years. She worked exclusively for over three years with high-performance System x servers, BladeCenter servers and chassis, and high-performance computing (HPC) clusters. She received a degree in Computer and Electronics Technology from Ivy Tech Community College and is a Microsoft Certified Professional.

Miroslav Peic is an IT specialist at IBM Global Technology Services in Austria. He provides technical support and services on System x, BladeCenter, and System Storage entry-level disk systems products. He has 10 years experience in IT and co-authored the Lenovo Press publication *eX5 Implementation Guide*, SG24-7909. He has a Graduate Degree in Applied Computer Science and has many industry certifications, including the Microsoft Certified Systems Administrator 2003 and VMware VCP 4.1. He trains and provides technical support to other IBM professionals, IBM Business Partners, and clients. He started his career at IBM in 2008.



The team (l-r): David, Shelly, and Miro

Thanks to the following people for their contributions to this project:

System x marketing:

- ▶ Randy Lundin
- ▶ Iliyas Pathan
- ▶ Steve Simmons
- ▶ Dan Sprissler

System x development:

- ▶ Mike Anderson
- ▶ Milton Cobo
- ▶ James Day
- ▶ Rodrigo De la Parra
- ▶ Chris LeBlanc
- ▶ Nina Newton
- ▶ Ron Rubano
- ▶ Greg Sellman

Comments welcome

Your comments are important to us!

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

- ▶ Use the online feedback form found at the web page for this document:

<http://lenovopress.com/redp4874>

- ▶ Send your comments in an email to:
comments@lenovopress.com

Do you have the latest version?

We update our books and papers from time to time, so check whether you have the latest version of this document by clicking the **Check for Updates** button on the front page of the PDF. Pressing this button will take you to a web page that will tell you if you are reading the latest version of the document and give you a link to the latest if needed. While you're there, you can also sign up to get notified via email whenever we make an update.

Product overview

The System x3750 M4 is a four-socket server that features a streamlined design, which is optimized for price and performance with best-in-class flexibility and expandability.

Note: There are three machine types of the x3750 M4:

- ▶ Machine type 8753 with E5-4600 v2 processors and Lenovo-signed firmware
- ▶ Machine type 8752 with E5-4600 v2 processors and IBM-signed firmware
- ▶ Machine type 8722 with E5-4600 processors and IBM-signed firmware

Only 8753 is currently available. The other two machine types are now withdrawn from marketing.

The x3750 M4 provides maximum storage density, with flexible Peripheral Component Interconnect Express (PCIe) and 10 Gb Ethernet networking options in a 2U form factor.

The following topics are included in this chapter:

- ▶ 1.1, “Introduction” on page 2
- ▶ 1.2, “Key features” on page 2
- ▶ 1.3, “Locations of key components and connectors” on page 6
- ▶ 1.4, “Standard specifications” on page 8
- ▶ 1.5, “Standard models” on page 11
- ▶ 1.6, “Architecture” on page 12
- ▶ 1.7, “Processors” on page 13
- ▶ 1.8, “Memory” on page 17
- ▶ 1.9, “Internal disk storage options” on page 27
- ▶ 1.10, “Optical drive” on page 48
- ▶ 1.11, “PCI Express slots” on page 48
- ▶ 1.12, “Integrated networking: Type 8753 and 8752” on page 50
- ▶ 1.13, “Integrated networking: Type 8722” on page 51
- ▶ 1.14, “Network adapters” on page 53
- ▶ 1.15, “Power supplies” on page 53
- ▶ 1.16, “Integrated virtualization” on page 55
- ▶ 1.17, “System management” on page 56
- ▶ 1.18, “Physical and electrical specifications” on page 58

1.1 Introduction

Figure 1-1 shows the System x3750 M4.



Figure 1-1 The System x3750 M4

The x3750 M4 has outstanding memory performance that is achieved by supporting three registered DIMM modules (RDIMMs) per channel. This performance is at speeds of up to 25% faster than the Intel specification. The enhanced speed still provides world-class Lenovo reliability for your configuration. Load-reduced DIMM (LRDIMM) speeds are also 25% beyond the Intel specification for 1.35 V RDIMM. This speed not only improves performance, but reduces overall system power.

The x3750 M4 offers a flexible, scalable design, and a simple upgrade path to 16 hard disk drives (HDDs) or 32 solid-state drives (SSDs). The x3750 M4 solution provides the HDDs or SSDs with up to eight PCIe 3.0 slots and up to 1.5 TB of memory. The flexible embedded Ethernet solution provides two standard Gigabit Ethernet (GbE) ports onboard. This solution also provides a dedicated 10 GbE slot that allows for a choice of two copper or two fiber optic connections. Comprehensive systems management tools with the next-generation Integrated Management Module II (IMM2) make it easy to deploy, integrate, service, and manage.

1.2 Key features

The System x3750 M4 blends outstanding flexibility and expandability. The x3750 M4 2+2 socket design enables pay-as-you-grow processing with the new Intel Xeon E5-4600 v2 series processors and memory scalability to help lower cost and manage growth. The 5+3 PCIe socket design enables you to pay for PCIe capabilities as needed.

The x3750 M4 provides unmatched features and capabilities in a dense 2U design. This system can support the following features:

- ▶ Four processor sockets
- ▶ 48 DIMM sockets
- ▶ Mix and match internal storage with up to 16 HDDs or 32 SSDs
- ▶ Six hot-swap dual rotor fans
- ▶ Two power supplies
- ▶ Integrated 10 GbE networking with options for fiber or copper

1.2.1 Scalability and performance

The x3750 M4 offers numerous features to boost performance, improve scalability, and reduce costs:

- ▶ The Intel Xeon processor E5-4600 v2 family improves productivity by offering superior system performance. It provides increased productivity by offering processors up to 12 cores each and processors with core speeds up to 3.3 GHz, L3 cache up to 30 MB, and two QuickPath Interconnect (QPI) interconnect links operating at up to 8 GTps.
- ▶ The x3750 M4 2+2 processor socket design enables pay-as-you-grow processing with the Intel Xeon E5-4600 v2 series processors and memory scalability to help lower cost and manage growth.
- ▶ Up to four processors, 48 cores, and 96 threads maximize the concurrent execution of multithreaded applications.
- ▶ Intelligent and adaptive system performance with Intel Turbo Boost Technology 2.0 allows processor cores to run at maximum speeds during peak workloads by temporarily going beyond processor TDP.
- ▶ Intel Hyper-Threading Technology boosts performance for multithreaded applications by enabling simultaneous multithreading within each processor core, up to two threads per core.
- ▶ Intel Virtualization Technology integrates hardware-level virtualization hooks that allow operating system vendors to better use the hardware for virtualization workloads.
- ▶ 48 Load Reduced DIMMs (LRDIMMs) of 1866 MHz DDR3 ECC memory provide speed, high availability, and a memory capacity of up to 1.5 TB.
- ▶ The use of eXFlash™ solid-state drives (SSDs) instead of, or along with, traditional spinning drives (HDDs), can improve I/O performance. An SSD can support up to 100 times more I/O operations per second (IOPS) than a typical HDD.
- ▶ Up to 16 HDDs or 32 eXFlash SSDs, together with an optical drive at the same time, provide a flexible and scalable all-in-one platform to meet your increasing demands.
- ▶ The server offers a SAS switch backplane option to allow up to 16x 2.5-inch drives to attach to a single controller.
- ▶ The 5+3 PCI Express socket design of the server allows you to pay for PCIe capabilities as needed.
- ▶ The server has a dedicated flexible mezzanine LOM slot offering a variety of 10 GbE or Gigabit Ethernet adapters that do not occupy one of the 8 standard PCIe slots .
- ▶ The server offers PCI Express 3.0 I/O expansion capabilities that improve the theoretical maximum bandwidth by almost 100% (8 GTps per link using 128b/130b encoding) compared to the previous generation of PCI Express 2.0 (5 GTps per link using 8b/10b encoding).
- ▶ With Intel Integrated I/O Technology, the PCI Express 3.0 controller is integrated into the Intel Xeon processor E5 family. This integration reduces I/O latency and increases overall system performance.

1.2.2 Availability and serviceability

The x3750 M4 provides many features to simplify serviceability and increase system uptime:

- ▶ The server offers Chipkill, memory mirroring, and memory rank sparing for redundancy in the event of a memory failure.

- ▶ The server provides restart recovery for any failed processor. If a failure of processor 1 occurs, the server connects the Southbridge to processor 2 for reboot.
- ▶ Tool-less cover removal provides easy access to upgrades and serviceable parts, such as the processor, memory, and adapter cards.
- ▶ The server offers hot-swap drives, supporting RAID redundancy for data protection and greater system uptime.
- ▶ The server has up to two redundant hot-swap power supplies and six hot-swap dual-rotor N+N redundant fans to provide availability for business-critical applications.
- ▶ The power source independent light path diagnostics panel and individual light path LEDs lead the technician to failed (or failing) components, which simplifies servicing, speeds up problem resolution, and helps improve system availability.
- ▶ The x3750 M4 provides error checking, Predictive Failure Analysis (PFA), or both on the following components. Alerts are generated in advance of a possible failure, therefore increasing uptime.
 - Memory ECC correction (Chipkill technology)
 - Microprocessor built-in self test (BIST) and internal error checking
 - PCIe Bus parity checking
 - SAS Bus Parity
 - HDD and SSD drive CRC checking
 - Diskette CRC checking
 - USB CRC checking
 - Over temperature detection
 - Over/under voltage detection
 - Power supply error checking
 - Fan failure detection
 - Serial interface parity, overrun, and frame checking
- ▶ Solid-state drives (SSDs) offer more reliability than traditional mechanical HDDs for greater uptime.
- ▶ The built-in Integrated Management Module Version II (IMM2) continuously monitors system parameters, triggers alerts, and performs recovery actions in case of failures to minimize downtime.
- ▶ Built-in diagnostics, using Dynamic Systems Analysis (DSA) Preboot, speed up troubleshooting tasks to reduce service time.
- ▶ Three-year customer-replaceable unit and onsite limited warranty, 9 x 5 next business day. Optional service upgrades are available.

1.2.3 Manageability and security

Powerful systems management features simplify local and remote management of the x3750 M4:

- ▶ The server includes an Integrated Management Module II (IMM2) to monitor server availability and perform remote management. Remote presence support is standard.
- ▶ The integrated industry-standard Unified Extensible Firmware Interface (UEFI) enables improved setup, configuration, and updates, and simplifies error handling.
- ▶ Integrated Trusted Platform Module (TPM) 1.2 support enables advanced cryptographic functionality, such as digital signatures and remote attestation.
- ▶ Industry-standard Advanced Encryption Standard (AES) NI support for faster, stronger encryption.

- ▶ Systems Director is included for proactive systems management. It offers comprehensive systems management tools that increases uptime, reduces costs, and improves productivity through advanced server management capabilities.
- ▶ Intel Execute Disable Bit functionality can prevent certain classes of malicious buffer overflow attacks when combined with a supported operating system.
- ▶ Intel Trusted Execution Technology provides enhanced security through hardware-based resistance to malicious software attacks, allowing an application to run in its own isolated space, protected from all other software running on a system.

1.2.4 Energy efficiency

The x3750 M4 offers the following energy-efficiency features to save energy, reduce operational costs, and increase energy availability:

- ▶ Energy-efficient planar components help lower operational costs.
- ▶ Highly efficient 750 W DC, 900 W AC, and 1400 W AC power supplies. 80 PLUS Platinum certification at high-voltage AC.
- ▶ The Intel Xeon processor E5-4600 v2 product family offers better performance over the previous generation E5-4600, although it fits into the same TDP limits.
- ▶ Intel Intelligent Power Capability powers individual processor elements on and off as needed to reduce power draw.
- ▶ Low-voltage Intel Xeon processors draw less energy to satisfy the demands of power and thermally constrained data centers and telecommunication environments.
- ▶ Low-voltage 1.35 V DDR3 memory RDIMMs consume 19% less energy compared to 1.5 V DDR3 RDIMMs.
- ▶ Solid-state drives (SSDs) consume as much as 80% less power than traditional spinning 2.5-inch HDDs.
- ▶ The server uses hexagonal ventilation holes, which are a part of Calibrated Vected Cooling technology. Hexagonal holes can be grouped more densely than round holes, providing more efficient airflow through the system.
- ▶ Systems Director Active Energy Manager provides advanced data center power notification and management to help achieve lower heat output and reduced cooling needs.

1.3 Locations of key components and connectors

Figure 1-2 shows the front of the server.

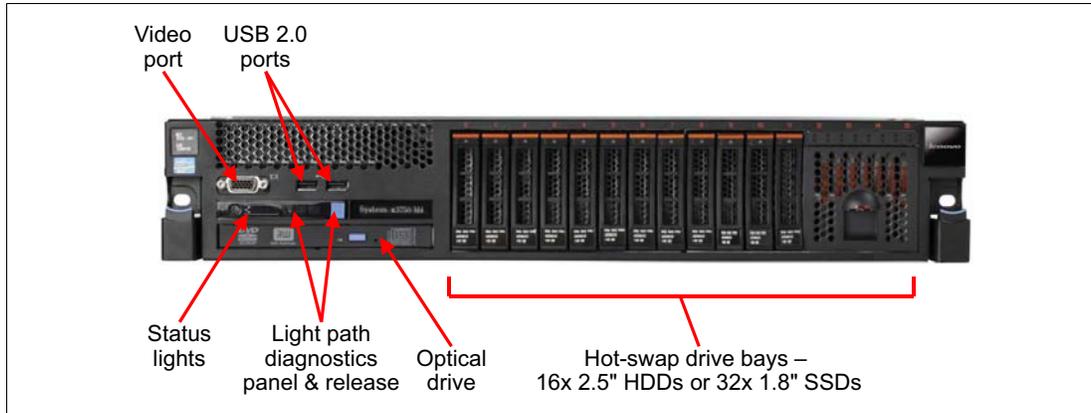


Figure 1-2 Front view of the System x3750 M4

Figure 1-3 shows the rear of the server for machine type 8753 and 8752 (with Intel Xeon Processor E5-4600 v2 family). The arrangement of PCIe slots is different compared to machine type 8722.

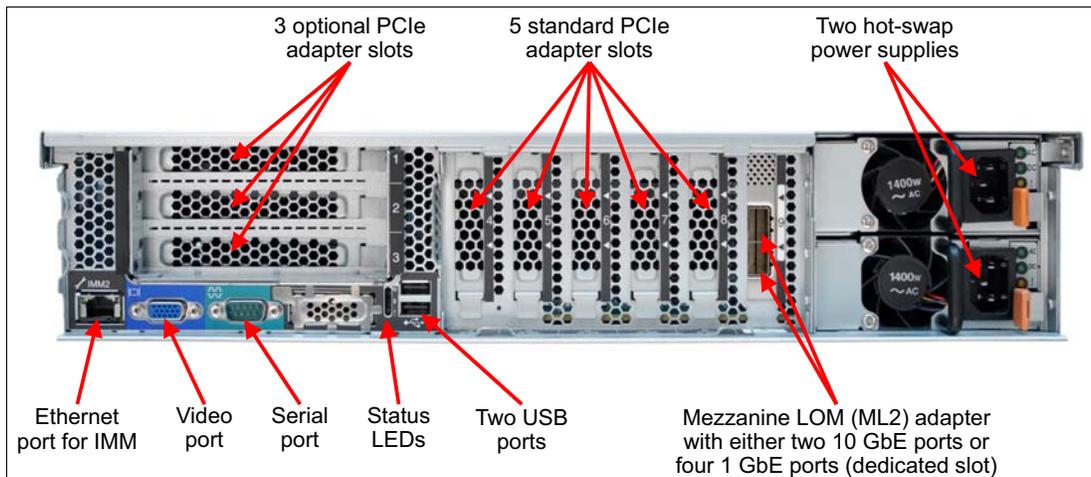


Figure 1-3 Rear view of the System x3750 M4 (8753 and 8752)

Figure 1-4 shows the rear of the server for machine type 8722 (with Intel Xeon Processor E5-4600 family).

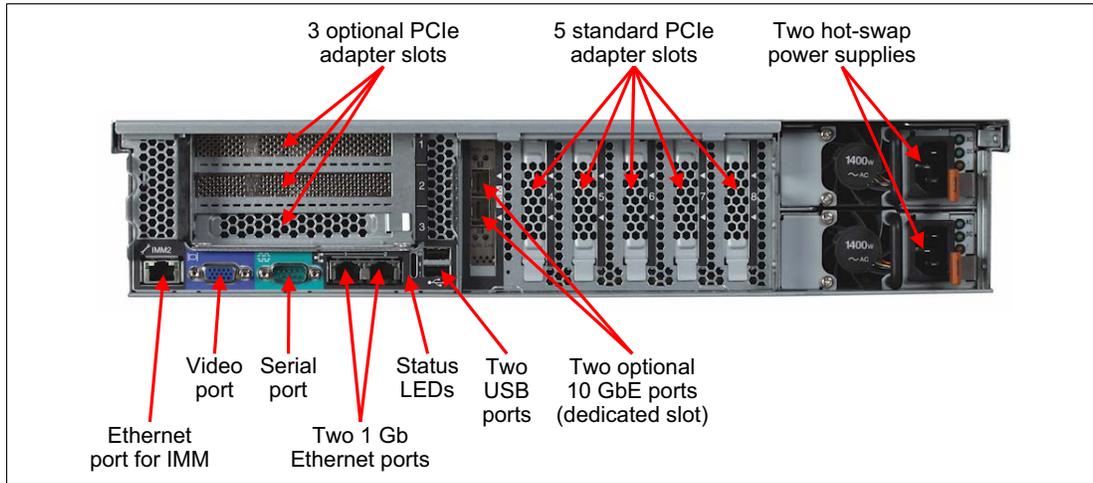


Figure 1-4 Rear view of the System x3750 M4 (8722)

Figure 1-5 shows the locations of key components inside the server for machine type 8753 and 8752 with (Intel Xeon Processor E5-4600 v2 family).

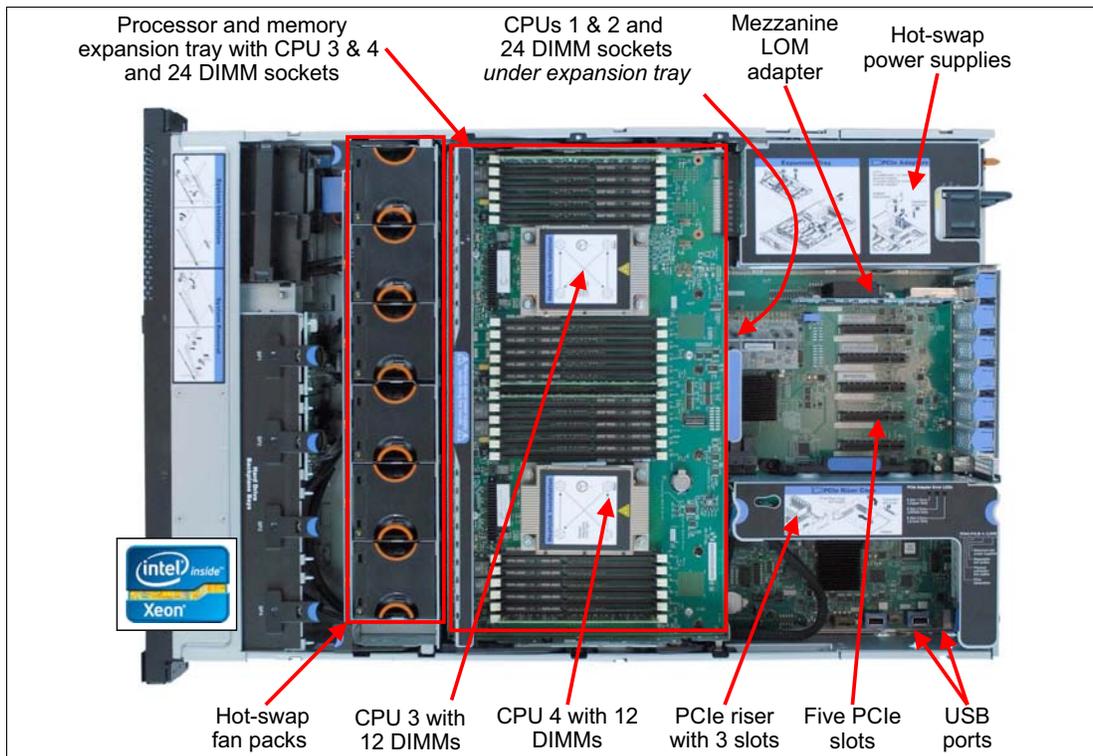


Figure 1-5 Inside view of the System x3750 M4 (8753 and 8752)

Figure 1-6 shows the locations of key components inside the server for machine type 8722 (with Intel Xeon Processor E5-4600 family).

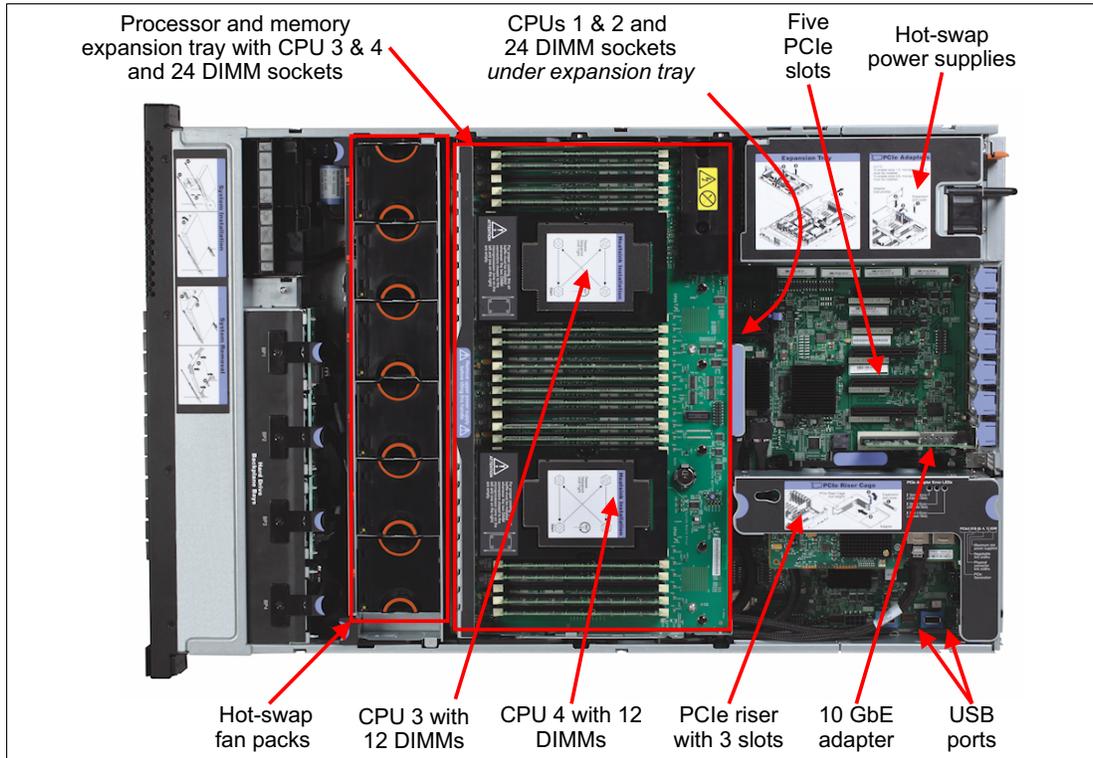


Figure 1-6 Inside view of the System x3750 M4 (8722)

1.4 Standard specifications

Table 1-1 lists the standard specifications of the System x3750 M4.

Table 1-1 Standard specifications

Components	Specification
Form factor	2U rack
Machine type	Machine type 8753: Intel Xeon Processor E5-4600 v2 family with Lenovo firmware Machine type 8752: Intel Xeon Processor E5-4600 v2 family with IBM firmware Machine type 8722: Intel Xeon Processor E5-4600 family with IBM firmware
Processor	Machine type 8753 and 8752: Up to four Intel Xeon E5-4600 v2 processors, each with 12 cores (2.4 GHz), 10 cores (up to 2.4 GHz), eight cores (up to 3.3 GHz), six cores (2.6 GHz), or four cores (2.2 GHz). Up to 1866 MHz memory speed. Up to 30 MB L3 cache per processor. Machine type 8722: Up to four Intel Xeon E5-4600 processors, each with eight cores (up to 2.7 GHz), six cores (up to 2.9 GHz), or four cores (up to 2.0 GHz). Up to 1600 MHz memory speed. Up to 20 MB L3 cache per processor. Two processor sockets on the system board and two processors on the processor and memory expansion tray (standard on most models). Two QPI links up to 8.0 GTps.
Chipset	Intel C600 series

Components	Specification
Memory	<p>Up to 48 DIMM sockets (12 DIMMs per processor). RDIMMs and LRDIMMs are supported, but memory types cannot be intermixed. There are 24 DIMM sockets on the system board. There are another 24 DIMM sockets on the processor and memory expansion tray (standard on most models).</p> <p>Machine type 8753 and 8752: Memory speed is up to 1866 MHz. Machine type 8722: Memory speed is up to 1600 MHz.</p>
Memory maximums	<p>Machine type 8753 and 8752:</p> <ul style="list-style-type: none"> ▶ With RDIMMs: Up to 768 GB with 48x 16 GB RDIMMs and four processors ▶ With LRDIMMs: Up to 1.5 TB with 48x 32 GB LRDIMMs and four processors <p>Machine type 8722:</p> <ul style="list-style-type: none"> ▶ With RDIMMs: Up to 768 GB with 48x 16 GB RDIMMs and four processors ▶ With LRDIMMs: Up to 1.5 TB with 48x 32 GB LRDIMMs and four processors
Memory protection	ECC, Chipkill (for x4-based memory DIMMs), memory mirroring, and memory sparing.
Disk drive bays	Up to 16 2.5-inch hot-swap SAS/SATA bays or up to 32 1.8-inch hot-swap SSD eXFlash bays. Drive bays can be in any combination of four 2.5-inch drives or eight 1.8-inch eXFlash SSDs.
Maximum internal storage	Up to 25.6 TB with 1.6 TB 2.5-inch SSDs, up to 19.2 TB with 1.2 TB 2.5-inch SAS HDDs, up to 16 TB with 1 TB 2.5-inch NL SAS/SATA HDDs. Intermix of SAS/SATA supported.
RAID support	<p>Machine type 8753 and 8752: RAID 0, 1, 10 with integrated ServeRAID™ M5210e with LSI SAS3108 RAID on Chip (ROC) controller. Optional upgrades to RAID 5 and 50 are available (1 GB cache no battery, or 1 or 2 GB flash-backed cache). Optional upgrades to RAID 6 and 60.</p> <p>Machine type 8722: RAID 0, 1, 10 with integrated ServeRAID M5110e with LSI SAS2208 6 Gbps RAID on Chip (ROC) controller. Optional upgrades to RAID 5 and 50 are available (512 B cache no battery, or 512 MB or 1 GB flash-backed cache). Optional upgrade to RAID 6 and 60.</p>
Optical drive bays	There is one bay for an optional Multi-burner drive.
Tape drive bays	None internal. Use a supported external tape drive.
Network interfaces	<p>Machine type 8753 and 8752: Dedicated mezzanine LOM adapter slot for a choice of 2-port 10 GbE (RJ45 or SFP+) or 4-port 1 GbE controller. One port can optionally be shared with the IMM2 management processor. Standard models include Intel I350-T4 ML2 Quad Port GbE Adapter (I350 based) except model C2x, which includes Broadcom NetXtreme II ML2 Dual Port 10GbaseT (BCM57712 based)</p> <p>Machine type 8722: Emulex BE3 controller with two standard integrated Gigabit Ethernet 1000BASE-T ports (RJ-45) and two optional 10 Gb ports through an adapter in a dedicated slot. The 10 GbE options are 10Base-T dual port (copper) or SFP+ dual port (fiber). One port can optionally be shared with the IMM2 management processor.</p>

Components	Specification
PCI Expansion slots	<p>Up to eight slots, five on the system board, up to three on an optional riser card. The following slots are available:</p> <ul style="list-style-type: none"> ▶ Slot 1: PCIe 3.0 x8; full-height, half-length (optional with riser card, requires CPU 2) ▶ Slot 2: PCIe 3.0 x8; full-height, half-length (optional with riser card, requires CPU 2) ▶ Slot 3: PCIe 3.0 x8; full-height, half-length (optional with riser card, requires CPU 2) ▶ Slot 4: PCIe 3.0 x8; low profile (requires CPU 2) ▶ Slot 5: PCIe 3.0 x8; low profile (requires CPU 2) ▶ Slot 6: PCIe 3.0 x8; low profile ▶ Slot 7: PCIe 3.0 x8; low profile ▶ Slot 8: PCIe 3.0 x8; low profile <p>Slots 1, 2, and 3 are physically x16 slots. Alternative two-slot riser with one x16 and one x8 slot also is available (machine type 8752 only).</p>
Ports	<ul style="list-style-type: none"> ▶ Front: Two USB 2.0 and one DB-15 video on front. ▶ Rear: Two USB 2.0, one DB-15 video, one DB-9 serial, one RJ-45 systems management ports, mezzanine LOM adapter with four 1 GbE network ports or two 10 GbE (RJ-45 or SFP+) network ports in dedicated slot. ▶ Internal: Two internal USB ports (for the embedded hypervisor).
Cooling	Calibrated Vectored Cooling with up to six N+N redundant hot swap fans (all six are standard); each fan has two rotors.
Power supply	Up to two hot-swap redundant AC power supplies (80 PLUS Platinum certification). Standard models use 900W supplies; available 1400W AC and 750W DC options. Second power supply requires processor expansion tray (88Y7365) or the power interposer card (88Y7367) installed.
Video	Matrox G200eR2 with 16 MB memory that is integrated into the IMM2. Maximum resolution is 1600 x 1200 at 75 Hz with 16 M colors.
Hot-swap parts	Hard disk drives, power supplies, and fans.
Systems management	UEFI, Integrated Management Module II (IMM2), Predictive Failure Analysis, Light Path Diagnostics, Automatic Server Restart, Lenovo XClarity™ Administrator, IBM Systems Director and Active Energy Manager, and ServerGuide™. IMM Advanced Upgrade software features for remote presence are standard with the x3750 M4.
Security features	Power-on password, administrator's password, and Trusted Platform Module (TPM).
Supported operating systems	Microsoft Windows Server 2008 R2, 2012, 2012 R2; RHEL 6 x 64; SLES 11 x 64; VMware vSphere 5.1, 5.5. For more information, see 4.2, "Supported operating systems" on page 126.
Limited warranty	Three-year customer-replaceable unit and onsite limited warranty with 9x5 NBD.
Service and support	<p>The following optional service upgrades are available through Lenovo warranty upgrade offerings:</p> <ul style="list-style-type: none"> ▶ Four-hour or two-hour response time ▶ Eight-hour fix time ▶ One-year or two-year warranty extension ▶ Remote technical support for Lenovo hardware and some Lenovo and third-party applications
Dimensions	Height: 86 mm (3.4 in.), width: 445 mm (17.5 in.), depth: 746 mm (29.4 in.)
Weight	Minimum configuration: 25 kg (55 lb.), maximum: 30 kg (65 lb.)

The x3750 M4 servers are shipped with the following items:

- ▶ Statement of Limited Warranty
- ▶ Important Notices
- ▶ Rack Installation Instructions
- ▶ Documentation CD that contains the Installation and User's Guide
- ▶ System x Gen-III Slides Kit
- ▶ System x Gen-III Cable Management Arm (CMA)
- ▶ 2.8 m (9.2 ft.) C13-C14 power cord (one for models with one power supply, and two for models with two power supplies)

1.5 Standard models

For a list of models see the Lenovo Press product guides:

- ▶ Machine type 8753: <https://lenovopress.com/tips1252>
- ▶ Machine type 8752: <http://lenovopress.com/tips1159>
- ▶ Machine type 8722: <http://lenovopress.com/tips0881>

1.6 Architecture

Figure 1-7 shows the block diagram of the system. Machine type 8753 and 8752 have an ML2 adapter slot that supports various ML2-form factor adapters, including 4-port 1 Gb Ethernet and 2-port 10 Gb Ethernet. Machine type 8722 has an integrated Emulex BE3 Ethernet controller that supplies two 10 Gb ports and two 1 Gb ports.

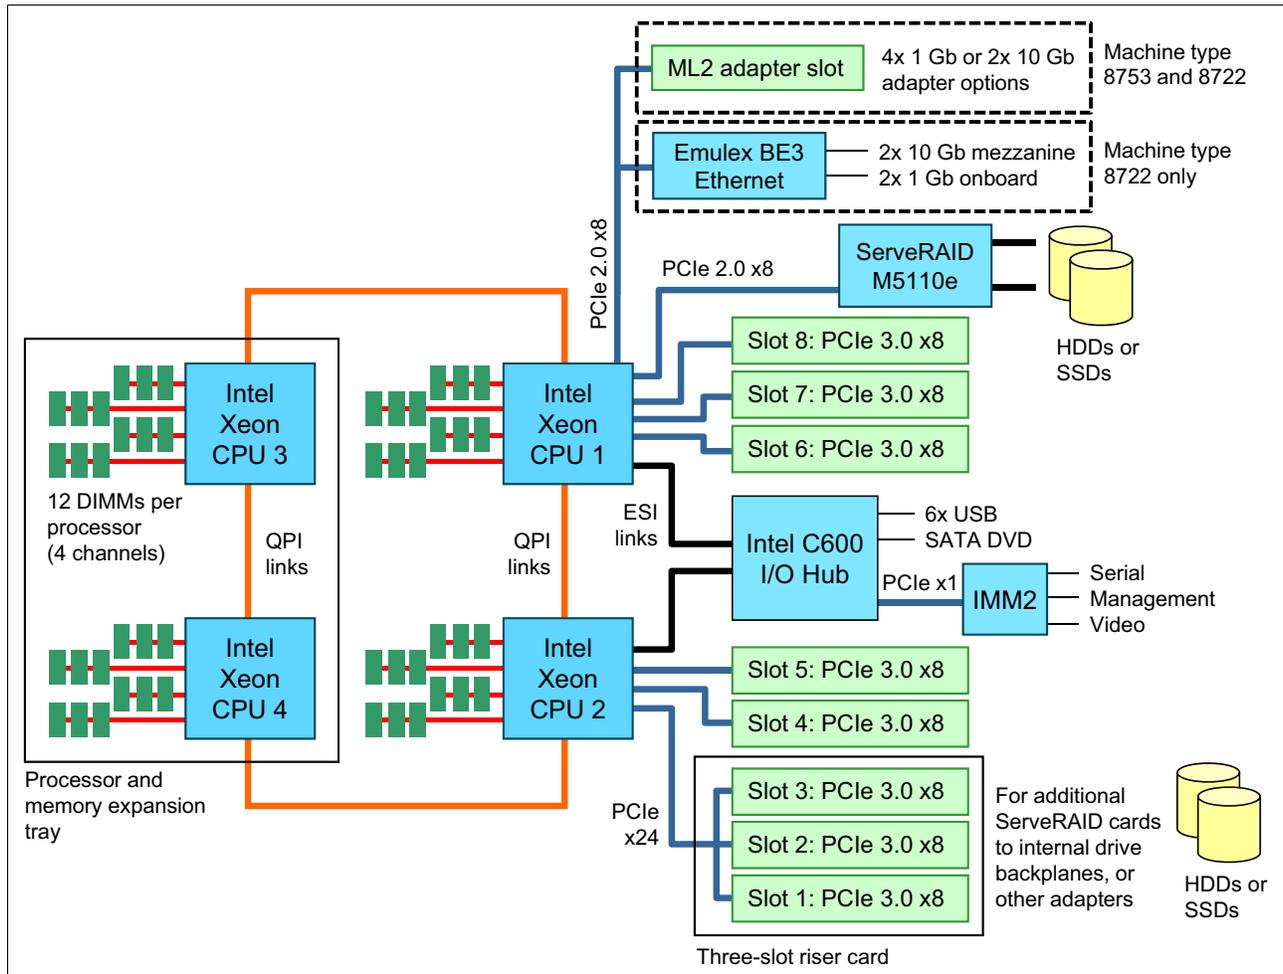


Figure 1-7 The System x3750 M4

Figure 1-8 shows the processors and memory in the server. The processor and memory expansion tray is removed and placed on top of the front of the server to show the lower two processors. The figure also shows where the expansion tray plugs into the QPI ports.

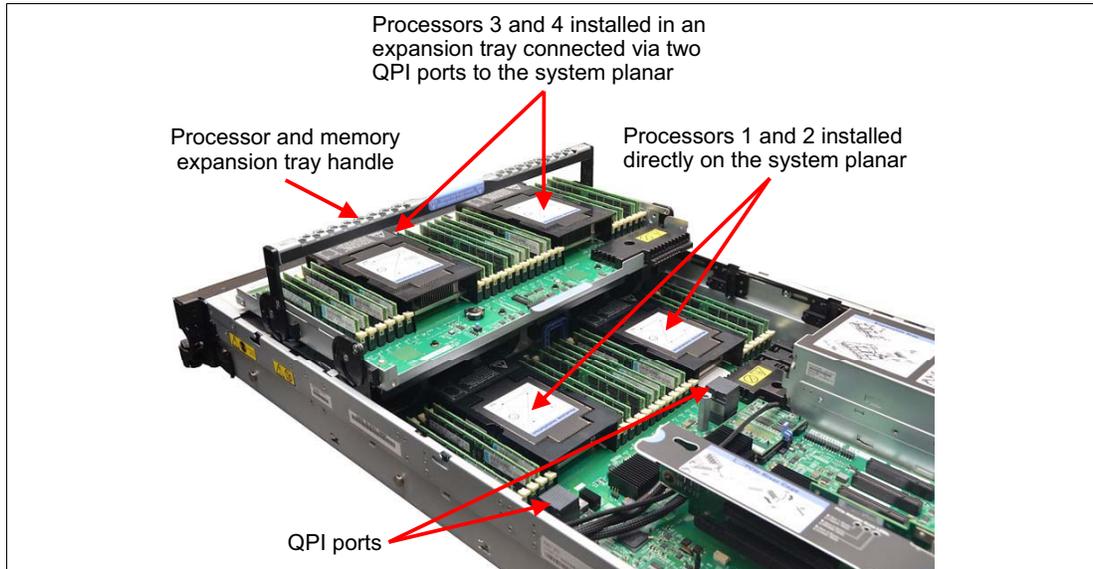


Figure 1-8 Processor layout

1.7 Processors

The server supports one, two, or four processors. Processors 1 and 2 are installed in sockets on the system board. Processors 3 and 4 are installed on the processor and memory expansion tray. Most models (except for A2x) have the expansion tray that is installed as standard. For model A2x or for configure-to-order (CTO), use the ordering information that is listed in Table 1-2.

Table 1-2 Processor and memory expansion tray

Part number	Feature code	Description
88Y7365	A29Z	System x3750 M4 processor and memory expansion tray

The following processor quantities are supported:

- ▶ One processor, which is installed in socket 1
- ▶ Two processors, which are installed in sockets 1 and 2
- ▶ Four processors, which are installed in all four sockets

For all processor options, see 1.7.1, “Processor options” on page 15.

The Intel Xeon E5-4600 and E5-4600 v2 series processors are enhanced with various features. The following enhancements are some of the highlights for both families:

▶ Intel Turbo Boost 2.0

Intel enhanced the Turbo Boost Technology. In previous generations of Xeon processors, Turbo Boost increased the CPU clock rate temporarily. This adjustment was done to handle the load when the core count, power consumption, current, and processor temperature were all below their limits. With Turbo Boost 2.0, Intel integrated more intelligence. This increase helps Turbo Boost manage the power and thermal headroom, and it is smart enough to not engage when memory and I/O are bottlenecks.

With these improvements, the Turbo Boost is active for longer periods and is engaged more often.

For more information about Turbo Boost, see this website:

<http://www.intel.com/technology/turboboost/>

▶ Intel Data Direct I/O

For I/O, Intel no longer has the separate I/O hub. It now integrates PCI Express 3.0 I/O into the processor. Data Direct I/O helps to optimize data transfer between local CPU and PCIe devices. The combination of Data Direct I/O and PCI 3.0 provides a higher I/O performance with lower latencies and reduced power consumption.

Intel Data Direct I/O offers the following advantages:

- Higher bandwidth
- Faster access from CPU
- Reduced latency

For more information about Data Direct I/O, see this website:

<http://www.intel.com/content/www/us/en/io/direct-data-i-o.html>

Support: The use of Data Direct I/O requires support from the adapter, device driver, and operating system.

In addition, the Intel Xeon E5-4600 v2 processors include the following new features:

▶ New 22nm manufacturing process

Latest micro-architecture on leading 22nm manufacturing process. This results in up to 50% more cores and cache in the same physical package, which means improved price, performance, and energy efficiency.

▶ Intel Platform Protection Technology with Secure Key

This new security feature provides faster and more secure encryption. It works in part by generating a random number: a digital circuit generates 3 Gb stream of random bits that is based on thermal noise, then a conditioner generates 256-bit random seeds, then a pseudorandom-number generator creates 128-bit keys to be used by the RDRAND instruction.

▶ Intel Data Protection Technology with OS Guard

This feature offers improved protection against malware by preventing the execution of calls to the operating system from compromised applications in the user mode or code pages. For supported operating systems, it protects system from kernel and Ring 0 level attacks from user mode applications. This new feature provides more protection over existing Intel XD (eXecute Disable) bit technology.

- ▶ Advanced Programmable Interrupt Controller virtualization (APICv)

This new feature improves virtualization performance by reducing the overhead that is required to service every APIC interrupt.

1.7.1 Processor options

The x3750 M4 supports the processor options that are listed in Table 1-3 and Table 1-4.

The tables also show which server models include what processor as a standard feature. If there is no corresponding where-used model for a particular processor, this processor is only available through CTO.

Table 1-3 shows the Intel Xeon Processor E5-4600 v2 family options (machine type 8753 and 8722)

Table 1-3 Processor options: Intel Xeon Processor E5-4600 v2 family (machine type 8753 and 8752)

Part number	Feature codes ^a	Description	Models where used
00D1948	A43K / A43U	Intel Xeon E5-4603 v2 4C 2.2 GHz 10 MB 1333MHz 95W	8752-A2x
00D1953	A43L / A43V	Intel Xeon E5-4607 v2 6C 2.6 GHz 15 MB 1333MHz 95W	8752-A3x
00D1958	A43M / A43W	Intel Xeon E5-4610 v2 8C 2.3 GHz 16 MB 1600MHz 95W	8752-B1x
00D1963	A43N / A43X	Intel Xeon E5-4620 v2 8C 2.6 GHz 20 MB 1600MHz 95W	8752-B2x
00D1978	A43R / A440	Intel Xeon E5-4627 v2 8C 3.3 GHz 16 MB 1866MHz 130W	8752-A1x
00D0574	A452 / A453	Intel Xeon E5-4624L v2 10C 1.9 GHz 25 MB 1866MHz 70W	N/A
00D1968	A43P / A43Y	Intel Xeon E5-4640 v2 10C 2.2 GHz 20 MB 1866MHz 95W	8752-C1x
00D1973	A43Q / A43Z	Intel Xeon E5-4650 v2 10C 2.4 GHz 25 MB 1866MHz 95W	8752-C2x
00D1988	A43T / A442	Intel Xeon E5-4657L v2 12C 2.4 GHz 30 MB 1866MHz 115W	N/A

a. The first feature code corresponds to the first processor installed. The second feature is for all other processors.

Table 1-4 shows the Intel Xeon Processor E5-4600 family options.

Table 1-4 Processor options: Intel Xeon Processor E5-4600 family (machine type 8722)

Part number	Feature code ^a	Description	Models where used
88Y7446	A2AL / A2AM	Intel Xeon E5-4603 4C 2.0 GHz 10 MB 1066 MHz 95 W	8722-A2x
88Y7342	A2AE / A2AF	Intel Xeon E5-4607 6C 2.2 GHz 12 MB 1066 MHz 95 W	8722-A3x
88Y7336	A2AC / A2AD	Intel Xeon E5-4610 6C 2.4 GHz 15 MB 1333 MHz 95 W	8722-B1x, D1x
88Y7354	A2AJ / A2AK	Intel Xeon E5-4617 6C 2.9 GHz 15 MB 1600 MHz 130 W	8722-A1x
88Y7330	A2AA / A2AB	Intel Xeon E5-4620 8C 2.2 GHz 16 MB 1333 MHz 95 W	8722-B2x
88Y7348	A2AG / A2AH	Intel Xeon E5-4640 8C 2.4 GHz 20 MB 1600 MHz 95 W	8722-C1x
88Y7324	A2A8 / A2A9	Intel Xeon E5-4650 8C 2.7 GHz 20 MB 1600 MHz 130 W	8722-C2x, D2x
88Y7458	A2AQ / A2AR	Intel Xeon E5-4650L 8C 2.6 GHz 20 MB 1600 MHz 115 W	N/A

a. The first feature code corresponds to the first processor installed. The second feature is for all other processors.

1.7.2 Processor failure and recovery

All System x servers can tolerate the failure of a processor by disabling that processor and rebooting the system. However, for most systems, this recovery is not possible if processor 1 fails. However, the System x3750 M4 server is designed so that it can tolerate the failure of processor 1. This feature decreases downtime if a CPU fails, such as a CPU BIST failure or CPU iERR. Downtime is decreased by allowing the machine to run on another processor until time can be taken to engage Lenovo service.

When the server detects a failure of processor 1, the Intel C600 I/O hub detects the failure and reroutes the booting instructions to processor 2 instead. The system boots to processor 2, but with reduced functionality. This failover to a surviving processor is automatic and is handled by the system at boot time.

If you want to use this recovery, it is important to carefully plan your network and storage connectivity. Table 1-5 lists the devices that are inaccessible if there is a processor failure and which processor is the boot processor. For more information, see Figure 1-7 on page 12.

Unsupported processors: If there is a processor failure, a second processor is also automatically taken offline. This second processor is disabled because the server cannot operate with three processors. Note the following points:

- ▶ If processor 1 fails, processor 2 is also taken offline
- ▶ If processor 2 fails, processor 1 is also taken offline
- ▶ If processor 3 fails, processor 4 is also taken offline
- ▶ If processor 4 fails, processor 3 is also taken offline

Table 1-5 shows the effect of a disabled processor.

Table 1-5 Effect of a disabled processor

Disabled processor	Processor 1	Processor 2	Processor 3	Processor 4
Boot processor	Processor 2	Processor 1	Processor 1	Processor 1
System components to which you lose access	<ul style="list-style-type: none"> ▶ All memory that is connected to processor 1 ▶ Processor 3 ▶ All memory that is connected to processor 3 ▶ Onboard ServeRAID M5110e SAS controller ▶ Any drives that are connected to the onboard SAS controller ▶ Onboard 1 Gb Ethernet ▶ 10 Gb Ethernet that uses the mezzanine slot ▶ PCIe slots 6, 7 and 8 	<ul style="list-style-type: none"> ▶ All memory that is connected to processor 2 ▶ Processor 4 ▶ All memory that is connected to processor 4 ▶ PCIe slots 1-5 ▶ Any drives that are connected to ServeRAID adapters that are installed in these PCIe slots 	<ul style="list-style-type: none"> ▶ All memory that is connected to processor 3 ▶ Processor 4 ▶ All memory that is connected to processor 4 	<ul style="list-style-type: none"> ▶ All memory that is connected to processor 4 ▶ Processor 3 ▶ All memory that is connected to processor 3

The embedded Ethernet and SAS are directly connected to CPU 1, along with PCIe slots 6, 7, and 8; therefore, a CPU 1 failure essentially leaves those devices disabled. More or “failover” network adapters and storage adapters are needed to regain network connectivity and to boot to an operating system. For more information about this solution, see Figure 1-7 on page 12.

As shown in the block diagram in Figure 1-7 on page 12, one possible boot device (the USB ports boot device) is connected directly to the Southbridge chipset. In this configuration, the path is not lost if there is a CPU failure. Therefore, one option is to use ESXi that is installed on an embedded USB memory key.

Consider the following configuration choices to enable the server to continue operating if any processor fails:

- ▶ Boot from a USB hypervisor.
- ▶ If you are connecting to external storage, configure dual paths to your SAS fabric via a host bus adapter (HBA) that is installed in one of slots 1 - 5 and one of slots 6 - 8.
- ▶ If you are booting from local drives, ensure that a boot volume is accessible from the drives that are connected to the onboard ServeRAID controller. Also, ensure that a boot volume is accessible from the drives that are connected to a ServeRAID adapter that is installed in one of slots 1 - 3.
- ▶ Configure a network adapter in one of slots 1 - 5 in addition to the onboard network controllers.

For information about boot order settings in the UEFI, see 3.1, “System Settings in UEFI” on page 78.

1.8 Memory

The System x3750 M4 supports DDR3 memory. The server supports up to 48 DIMMs when four processors are installed, with 12 DIMMs for each processor. A total of 24 DIMM sockets (for processors 1 and 2) are on the system board. The remaining DIMM sockets (for processors 3 and 4) are on the processor and memory expansion tray. Each processor has four memory channels, and there are three DIMMs per channel.

The x3750 M4 memory system is carefully tuned so that the server supports higher memory frequencies than the Intel processor specification. For example, you can use low-voltage DIMMs but still operate them at the rated speed.

The following rules apply when the memory configuration is selected:

- ▶ The server supports RDIMMs and LRDIMMs. Unbuffered DIMMs (UDIMMs) are not supported.
- ▶ Mixing different types of memory (RDIMMs and LRDIMMs) is not supported.
- ▶ Mixing 1.5 V and 1.35 V DIMMs in the same server is supported; in such a case, all DIMMs operate at 1.5 V.
- ▶ The maximum number of ranks per one channel is eight (except for LRDIMMs, where more than eight ranks are supported). This exception occurs because one quad-rank LRDIMM provides the same electrical load on a memory bus as a single-rank RDIMM.

- ▶ The maximum quantity of DIMMs that can be installed in the server depends on the number of processors that are installed. Table 1-9 on page 21 shows the maximum quantity of DIMMs when all four processors are installed. When two processors are installed, the maximum quantity that is supported is a half of the quantity that is shown.
- ▶ All DIMMs in the server operate at the same speed, which is determined as the lowest value of the following solutions:
 - Memory speed that is supported by the specific processor.
 - Maximum operating speeds for the selected memory configuration, which depends on the rated speed, operating voltage, and the quantity of DIMMs per channel. This configuration is shown under the section “Max Operating Speed” in Table 1-9 on page 21.

1.8.1 Memory types

The System x3750 M4 supports the following DIMMs:

- ▶ RDIMM: Registered DIMMs
- ▶ LRDIMM: Load reduced DIMMs

These two types cannot be mixed. UDIMMs are not supported.

LRDIMMs are a new type of DDR3 DIMM. Each memory channel can have only eight memory ranks total (this means that only two quad-rank DIMMs can be installed on a channel, for example). LRDIMMs solve this limitation by adding a buffer and control logic to the data lines of the DIMM. This addition reduces the electrical loading on the channel and enables the DIMM to operate as though it were only a single or dual-rank DIMM¹.

LRDIMMs are the most useful for high-capacity memory requirements. With the added buffer, a small latency is introduced to memory operations. However, this latency is offset by the higher capacity and potential higher frequency in which the DIMM can operate. This outcome is a result of the decreased electrical load that it places on the channel.

For more information about LRDIMMs, see this website:

<http://edn.com/design/systems-design/4368420/Basics-of-LRDIMM>

1.8.2 Memory options

Table 1-6 lists the memory options that are available for the x3750 M4 server of machine type 8753 and 8752 (with Intel Xeon E5-4600 v2 processors).

Table 1-6 Memory options: Machine type 8753 and 8752 (with Intel Xeon E5-4600 v2 processors)

Part number	Feature code	Description	Maximum quantity supported	Where used
RDIMMs - 1866 MHz				
00D5048	A3QL	16 GB (1x16 GB, 2Rx4, 1.5V) PC3-14900 CL13 ECC DDR3 1866 MHz LP RDIMM	48 (12 per processor)	N/A
00D5032	A3QG	8 GB (1x8 GB, 1Rx4, 1.5V) PC3-14900 CL13 ECC DDR3 1866 MHz LP RDIMM	48 (12 per processor)	8752-A1x

¹ Whether the LRDIMM appears on the memory channel as a single or dual-rank DIMM (a feature called *rank multiplication*) depends on the specific DIMM that is used.

Part number	Feature code	Description	Maximum quantity supported	Where used
00D5020	A3QD	4 GB (1x4 GB, 1Rx4, 1.5V) PC3-14900 CL13 ECC DDR3 1866 MHz LP RDIMM	48 (12 per processor)	N/A
RDIMMs - 1600 MHz				
46W0672	A3QM	16 GB (1x16GB, 2Rx4, 1.35V) PC3L-12800 CL11 ECC DDR3 1600 MHz LP RDIMM	48 (12 per processor)	N/A
00D5036	A3QH	8 GB (1x8GB, 1Rx4, 1.35V) PC3L-12800 CL11 ECC DDR3 1600 MHz LP RDIMM	48 (12 per processor)	All other 8752 models
00D5024	A3QE	4 GB (1x4GB, 1Rx4, 1.35V) PC3L-12800 CL11 ECC DDR3 1600 MHz LP RDIMM	48 (12 per processor)	N/A
LRDIMMs				
46W0761	A47K	32 GB (1x32 GB, 4Rx4, 1.5V) PC3-14900 CL13 ECC DDR3 1866 MHz LP LRDIMM	48 (12 per processor)	N/A

Table 1-7 lists the memory options that are available for the x3750 M4 server of machine type 8722 (with Intel Xeon E5-4600 processors).

Table 1-7 Memory options: Machine type 8722 (with Intel Xeon E5-4600 processors)

Part number	Feature code	Description	Maximum quantity supported	Models where used
RDIMMs - 1600 MHz				
00D4968	A2U5	16 GB (1x 16 GB, 2Rx4, 1.5 V) PC3-12800 CL11 ECC DDR3 1600 MHz LP RDIMM	48 (12 per processor)	N/A
90Y3109	A292	8 GB (1x 8 GB, 2Rx4, 1.5 V) PC3-12800 CL11 ECC DDR3 1600 MHz LP RDIMM	48 (12 per processor)	8722-A1x
49Y1559	A28Z	4 GB (1x 4 GB, 1Rx4, 1.5 V) PC3-12800 CL11 ECC DDR3 1600 MHz LP RDIMM	48 (12 per processor)	N/A
RDIMMs - 1333 MHz				
49Y1563	A1QT	16 GB (1x 16 GB, 2Rx4, 1.35 V) PC3L-10600 CL9 ECC DDR3 1333 MHz LP RDIMM	48 (12 per processor)	N/A
49Y1397	8923	8 GB (1x 8 GB, 2Rx4, 1.35 V) PC3L-10600 CL9 ECC DDR3 1333 MHz LP RDIMM	48 (12 per processor)	All other 8722 models
49Y1406	8941	4 GB (1x 4 GB, 1Rx4, 1.35 V) PC3L-10600 CL9 ECC DDR3 1333 MHz LP RDIMM	48 (12 per processor)	N/A
LRDIMMs				
90Y3105	A291	32 GB (1x 32 GB, 4Rx4, 1.35 V) PC3L-10600 CL9 ECC DDR3 1333 MHz LP LRDIMM	48 (12 per processor)	N/A
49Y1567	A290	16 GB (1x 16 GB, 4Rx4, 1.35 V) PC3L-10600 CL9 ECC DDR3 1333 MHz LP LRDIMM	48 (12 per processor)	8722-D2x

1.8.3 Memory performance

The tables in this section show the maximum speed that the memory channels operate at based on the DIMMs that are installed. Installing three DIMMs per channel can slow the memory, depending on the parts that are installed, as the tables show. All memory channels operate at the same speed.

Maximize memory: Use 1600 MHz RDIMMs and limit the configuration to two DIMMs per channel to maximize memory speed. Use a full complement of LRDIMMs to maximize memory capacity.

The following key refers to the highlighted references in the tables:

- ▶ The entries that are highlighted in red text indicate that the System x3750 M4 supports higher memory frequencies than the Intel processor specification. In some instances, this configuration also results in larger memory capacity than the specification recommends.
- ▶ Table cells that are highlighted with a gray background indicate when the specific combination of DIMM voltage and the number of DIMMs per channel still allows the DIMMs to operate at the rated speed.

Table 1-8 shows the maximum memory speeds that are based on DIMMs per channel for the x3750 M4 server of machine type 8752 (with Intel Xeon E5-4600 v2 processors).

Table 1-8 Maximum memory speeds: Machine type 8752 (with Intel Xeon E5-4600 v2 processors)

DIMM specification	RDIMM					LRDIMM	
	Single-rank DIMMs			Dual-rank DIMMs		Quad-rank LRDIMMs	
Part numbers	00D5024 (4 GB) 00D5036 (8 GB)		00D5020 (4 GB) 00D5032 (8 GB)		46W0672(16 GB)	00D5048 (16 GB) 46W0761 (32 GB)	
Rated speed	1600 MHz		1866 MHz		1600 MHz	1866 MHz	
Rated voltage	1.35 V		1.5 V		1.35 V	1.5 V	
Operating voltage	1.35 V	1.5 V	1.5 V	1.35 V	1.5 V	1.5 V	
Max. qty supported ^a	48	48	48	48	48	48	
Max. DIMM capacity	8 GB	8 GB	8 GB	16 GB	16 GB	16 GB	
Max. memory capacity	384 GB	384 GB	384 GB	768 GB	768 GB	768 GB	
Max. memory at maximum speed	256 GB	256 GB	256 GB	512 GB	512 GB	512 GB	
Max operating speed (MHz)							
1 DIMM per channel	1333	1600	1866	1333	1600	1866	1866
2 DIMMs per channel	1600	1600	1866	1600	1600	1866	1866
3 DIMMs per channel	1066	1333	1333	1066	1333	1333	1333

a. The maximum quantity that is supported is shown for four processors installed. When two processors are installed, the maximum quantity that is supported is a half of the quantity that is shown. When one processor is installed, the quantity is one quarter of what is shown.

Table 1-9 on page 21 shows the maximum memory speeds that are based on DIMMs per channel for the x3750 M4 server of machine type 8722 (with Intel Xeon E5-4600 processors).

Table 1-9 Maximum memory speeds: Machine type 8722 (with Intel Xeon E5-4600 processors)

DIMM specification	RDIMM						LRDIMM	
	Single-rank DIMMs			Dual-rank DIMMs			Quad-rank LRDIMMs	
Ranks	Single-rank DIMMs			Dual-rank DIMMs			Quad-rank LRDIMMs	
Part numbers	49Y1406 (4 GB)	49Y1559 (4 GB)	49Y1397 (8 GB) 49Y1563 (16 GB)	90Y3109 (8 GB) 00D4968 (16 GB)	49Y1567 (16 GB) 90Y3105 (32 GB)			
Rated speed	1333 MHz	1600 MHz	1333 MHz	1600 MHz	1333 MHz			
Rated voltage	1.35 V		1.5 V	1.35 V		1.5 V	1.35 V	
Operating voltage	1.35 V	1.5 V	1.5 V	1.35 V	1.5 V	1.5 V	1.35 V	1.5 V
Max. qty supported ^a	48	48	48	48	48	48	48	48
Max. DIMM capacity	4 GB	4 GB	4 GB	16 GB	16 GB	16 GB	32 GB	32 GB
Max. memory capacity	192 GB	192 GB	192 GB	768 GB	768 GB	768 GB	1.5 TB	1.5 TB
Max memory at maximum speed	128 GB	128 GB	128 GB	512 GB	768 GB	512 GB	1.5 TB	1.5 TB
Max operating speed (MHz)								
1 DIMM per channel	1333	1333	1600	1333	1333	1600	1333	1333
2 DIMMs per channel	1333	1333	1600	1333	1333	1600	1333	1333
3 DIMMs per channel	1066	1066	1066	1066	1333	1333	1333	1333

a. The maximum quantity that is supported is shown for four processors installed. When two processors are installed, the maximum quantity that is supported is a half of the quantity that is shown. When one processor is installed, the quantity is one quarter of what is shown.

The following memory protection technologies are supported:

- ▶ ECC
- ▶ Chipkill (for x4-based memory DIMMs; look for “x4” in the DIMM description)
- ▶ Memory mirroring
- ▶ Memory rank sparing

If memory mirroring is used, DIMMs must be installed in pairs (minimum of one pair per CPU). Both DIMMs in a pair must be identical in type and size.

If memory rank sparing is used, a minimum of one quad-rank DIMM or two single-rank or dual-rank DIMMs must be installed per populated channel. The DIMMs do not need to be identical. In rank sparing mode, one rank of a DIMM in each populated channel is reserved as spare memory. The size of a rank varies, depending on the DIMMs that are installed.

1.8.4 Memory operational modes

As clients put more critical data in memory and memory workload demands increase, it is critical for platforms to provide a level of protection for data that is in memory.

The System x3750 M4 offers enhanced memory availability by using memory mirroring or memory rank sparing. Both of these memory operation modes reserve a portion of memory if a rank or an entire DIMM fails.

You select which memory mode you want the server to operate in within the UEFI. The default is Independent Mode where mirroring nor rank sparing is enabled. For more information about the necessary UEFI settings, see “Memory” on page 87.

Tip: The mode that you select affects all processors. You cannot operate the server in a mix of memory operation modes.

Memory mirroring

In this mode, the data that is stored in DIMMs on channel 1 of each processor is mirrored to the DIMMs on channel 2 of the same processor. Similarly, channel 3 is mirrored to channel 4. If a failure occurs, the memory controller switches from the DIMMs on the primary channel to the DIMMs on the backup channel.

DIMMs must be installed in pairs when you use mirroring. Install at least one pair per CPU. The DIMMs in a pair must have the same type and size across channels 1 and 2, and channels 3 and 4. DIMM slots within a channel do not have to be identical.

Memory mirroring reduces the total memory capacity by 50%.

Memory rank sparing

In the Memory Rank Sparing mode, one memory rank of a DIMM acts as a spare for the other ranks on the same channel. The spare rank is held in reserve and is not available as system memory. In this mode, if the number of correctable errors in an active rank reaches a certain threshold, that rank is disabled and the spare rank is activated. The spare rank must be identical or a larger capacity than all of the other ranks on the same channel.

The amount of memory that is reserved by sparing depends on the total number of the ranks in a channel and the size of the DIMMs. For example, if three dual-rank DIMMs are installed in a channel, one rank is reserved for sparing, which leaves five ranks for data. If the DIMMs are 4 GB each, that means that channel has 10 GB of available memory.

Tip: With memory rank sparing, one rank in every memory channel is reserved.

1.8.5 Population sequence

Each Memory DIMM is connected to a processor. The system board provides four memory channels and each memory channel has three DIMM slots.

Table 1-10 shows the DIMM connectors that are associated with each processor.

Table 1-10 DIMM connectors that are associated with each microprocessor

Processor	Location	DIMM connectors that are associated with the processor
Processor 1	System board	1 - 12
Processor 2	System board	13 - 24
Processor 3	Processor and memory expansion tray	25 - 36
Processor 4	Processor and memory expansion tray	37 - 48

Table 1-11 shows the memory channels that are distributed on the memory slot and processor.

Table 1-11 DIMMs per CPU for each Memory Channel

	Memory Channel 1	Memory Channel 2	Memory Channel 3	Memory Channel 4	Location
DIMMs for CPU 1	7, 8, 9	10, 11, 12	4, 5, 6	1, 2, 3	System board
DIMMs for CPU 2	19, 20, 21	22, 23, 24	16, 17, 18	13, 14, 15	System board
DIMMs for CPU 3	31, 32, 33	34, 35, 36	28, 29, 30	25, 26, 27	Processor and memory expansion tray
DIMMS for CPU 4	43, 44, 45	46, 47, 48	40, 41, 42	37, 38, 39	Processor and memory expansion tray

For one processor that is operating in Independent mode, install DIMMs in the order that is shown in Table 1-12.

Table 1-12 DIMM installation order (one processor)

Number of DIMMs	Processor 1											
	Channel 1			Channel 2			Channel 3			Channel 4		
	DIMM 7	DIMM 8	DIMM 9	DIMM 10	DIMM 11	DIMM 12	DIMM 4	DIMM 5	DIMM 6	DIMM 1	DIMM 2	DIMM 3
1										X		
2						X				X		
3						X	X			X		
4			X			X	X			X		
5			X			X	X			X	X	
6			X		X	X	X			X	X	
7			X		X	X	X	X		X	X	
8		X	X		X	X	X	X		X	X	
9		X	X		X	X	X	X		X	X	X
10		X	X	X	X	X	X	X		X	X	X
11		X	X	X	X	X	X	X	X	X	X	X
12	X	X	X	X	X	X	X	X	X	X	X	X

For two processors that operate in Independent mode, install DIMMs in the order that is shown in Table 1-13.

Table 1-13 DIMM installation order (two processors)

Number of DIMMs	Processor 1												Processor 2											
	Channel 1			Channel 2			Channel 3			Channel 4			Channel 1			Channel 2			Channel 3			Channel 4		
	DIMM 7	DIMM 8	DIMM 9	DIMM 10	DIMM 11	DIMM 12	DIMM 4	DIMM 5	DIMM 6	DIMM 1	DIMM 2	DIMM 3	DIMM 19	DIMM 20	DIMM 21	DIMM 22	DIMM 23	DIMM 24	DIMM 16	DIMM 17	DIMM 18	DIMM 13	DIMM 14	DIMM 15
2									X													X		
3					X				X													X		
4					X				X								X					X		
5					X	X			X								X					X		
6					X	X			X								X	X				X		
7		X			X	X			X								X	X				X		
8		X			X	X			X					X			X	X				X		
9		X			X	X			X	X				X			X	X				X		
10		X			X	X			X	X				X			X	X				X	X	
11		X		X	X	X			X	X				X			X	X				X	X	
12		X		X	X	X			X	X				X		X	X	X				X	X	
13		X		X	X	X	X		X	X				X		X	X	X				X	X	
14		X		X	X	X	X		X	X				X		X	X	X	X			X	X	
15	X	X		X	X	X	X		X	X				X		X	X	X	X			X	X	
16	X	X		X	X	X	X		X	X			X	X		X	X	X	X			X	X	
17	X	X		X	X	X	X		X	X	X		X	X		X	X	X	X			X	X	
18	X	X		X	X	X	X		X	X	X		X	X		X	X	X	X			X	X	X
19	X	X	X	X	X	X	X		X	X	X		X	X		X	X	X	X			X	X	X
20	X	X	X	X	X	X	X		X	X	X		X	X	X	X	X	X	X			X	X	X
21	X	X	X	X	X	X	X	X	X	X	X		X	X	X	X	X	X	X			X	X	X
22	X	X	X	X	X	X	X	X	X	X	X		X	X	X	X	X	X	X	X	X	X	X	X
23	X	X	X	X	X	X	X	X	X	X	X		X	X	X	X	X	X	X	X	X	X	X	X
24	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X

Number of DIMMs	Processor 1				Processor 2				Processor 3				Processor 4																																		
	Ch1		Ch2		Ch3		Ch4		Ch1		Ch2		Ch3		Ch4																																
	DIMM 7	DIMM 8	DIMM 9	DIMM 10	DIMM 11	DIMM 12	DIMM 4	DIMM 5	DIMM 6	DIMM 1	DIMM 2	DIMM 3	DIMM 19	DIMM 20	DIMM 21	DIMM 22	DIMM 23	DIMM 24	DIMM 16	DIMM 17	DIMM 18	DIMM 13	DIMM 14	DIMM 15	DIMM 31	DIMM 32	DIMM 33	DIMM 34	DIMM 35	DIMM 36	DIMM 28	DIMM 29	DIMM 30	DIMM 25	DIMM 26	DIMM 27	DIMM 43	DIMM 44	DIMM 45	DIMM 46	DIMM 47	DIMM 48	DIMM 40	DIMM 41	DIMM 42	DIMM 37	DIMM 38
33	X	X		X	X	X	X		X	X	X		X	X		X	X		X	X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		
34	X	X		X	X	X	X		X	X	X		X	X		X	X		X	X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		
35	X	X		X	X	X	X		X	X	X		X	X		X	X		X	X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		
36	X	X		X	X	X	X		X	X	X		X	X		X	X		X	X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		
37	X	X	X	X	X	X	X		X	X	X		X	X		X	X		X	X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		
38	X	X	X	X	X	X	X		X	X	X		X	X		X	X		X	X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		
39	X	X	X	X	X	X	X		X	X	X		X	X		X	X		X	X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		
40	X	X	X	X	X	X	X		X	X	X		X	X		X	X		X	X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		
41	X	X	X	X	X	X	X		X	X	X		X	X		X	X		X	X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		
42	X	X	X	X	X	X	X		X	X	X		X	X		X	X		X	X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		
43	X	X	X	X	X	X	X		X	X	X		X	X		X	X		X	X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		
44	X	X	X	X	X	X	X		X	X	X		X	X		X	X		X	X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		
45	X	X	X	X	X	X	X		X	X	X		X	X		X	X		X	X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		
46	X	X	X	X	X	X	X		X	X	X		X	X		X	X		X	X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		
47	X	X	X	X	X	X	X		X	X	X		X	X		X	X		X	X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		
48	X	X	X	X	X	X	X		X	X	X		X	X		X	X		X	X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		

For more information about the population sequence for memory mirroring and rank sparing, see *System x3750 M4 Installation and Service Guide*, which is available at this website:

<http://ibm.com/support/entry/portal/docdisplay?lnocid=MIGR-5090828>

Minimum DIMMs: A minimum of one DIMM must be installed for each processor.

1.9 Internal disk storage options

The server can support up to 16 x 2.5-inch drives, up to 32 x 1.8-inch drives, or a combination of both 2.5-inch and 1.8-inch hot-swap drives. This configuration is possible by using the supported SAS/SATA backplane configurations.

The server supports the following options:

- ▶ 2.5-inch hot-swap SAS or SATA HDDs
- ▶ 2.5-inch hot-swap SATA SSDs
- ▶ 1.8-inch hot-swap SATA SSDs

You can mix drives in the same server if you do not mix drives on the same array.

The following topics are included in this section:

- ▶ 1.9.1, “Backplanes” on page 27
- ▶ 1.9.2, “Backplane combinations” on page 30
- ▶ 1.9.3, “RAID controllers” on page 40
- ▶ 1.9.4, “RAID controller upgrades” on page 42
- ▶ 1.9.5, “Battery mounting” on page 45
- ▶ 1.9.6, “Battery versus flash” on page 46
- ▶ 1.9.7, “Drive selection” on page 46
- ▶ 1.9.8, “Drive options” on page 47

1.9.1 Backplanes

The System x3750 M4 server supports various internal storage configurations that are based on four different backplanes. Table 1-15 shows the backplane options that are available for an x3750 M4 server.

Table 1-15 Internal storage expansion options

Part number	Feature code	Name	Maximum supported
00D2020	A2A3	4x 2.5-in. HS SAS HDD Backplane	4 ^a
00D2022	A2A4	8x 2.5-in. HS SAS HDD Backplane	2
00D2011	A2A2	8x 2.5-in. HS SAS HDD Backplane with RAID expansion (SAS expander)	1 ^b
00D2024	9281	eXFlash 8x 1.8-in. HS SAS SSD Backplane	4

a. When a server is ordered via CTO, only one 4-drive backplane can be selected. If the order has more than 4-drive bays, 8-drive backplanes are selected instead.

b. Only one backplane with SAS expander (00D2011) can be installed in a server and can be used connected to a single 8-drive backplane (00D2022), a single 4-drive backplane (00D2020), or two 4-drive backplanes (00D2020). The SAS expander backplane cannot be used with the eXFlash backplane.

Standard models (all models except A2x) ship with at least one backplane. All backplane options include the necessary cables to connect to the integrated RAID controller or any installed RAID adapters.

The following backplane options are available:

- ▶ Four-drive backplane for 2.5-inch drives

Up to four backplanes can be installed, each requiring one SAS connection to a supported controller. One RAID controller can connect to two of these backplanes.

Figure 1-9 shows the 4x 2.5-inch backplane and the connectors.

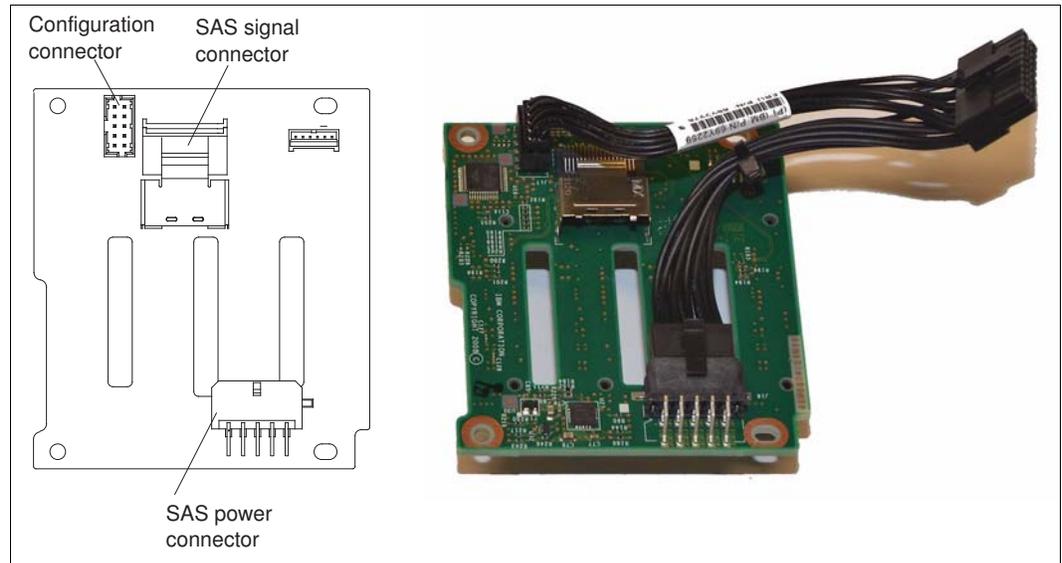


Figure 1-9 4x 2.5-inch drive backplane (drive connectors are on the other side of the board)

- ▶ Eight-drive backplane for 2.5-inch drives

Up to two backplanes can be installed, each requiring two SAS connections. One RAID controller is needed per backplane.

Figure 1-10 shows the 8x 2.5-inch backplane and the connectors.

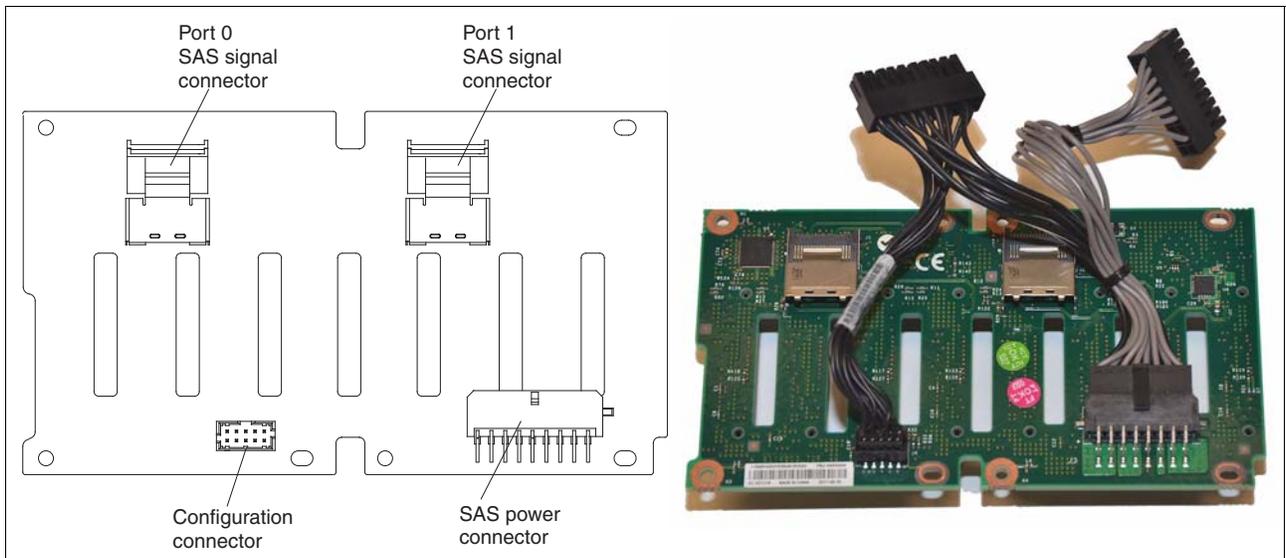


Figure 1-10 8x 2.5-inch drive backplane

- ▶ Eight-drive backplane with RAID expansion (SAS expander) for 2.5-inch drives
Use this backplane with the other 2.5-inch drive backplanes (one 8-drive, one 4-drive, or two 4-drive backplanes). This configuration results in the ability to connect up to 16 2.5-inch drive bays to one RAID controller. This solution is useful if you want to create a single array (volume) from all 16 drives.

Figure 1-11 shows the 8x 2.5-inch backplane with RAID expansion. The SAS expander daughter card and connectors are attached to the backplane.

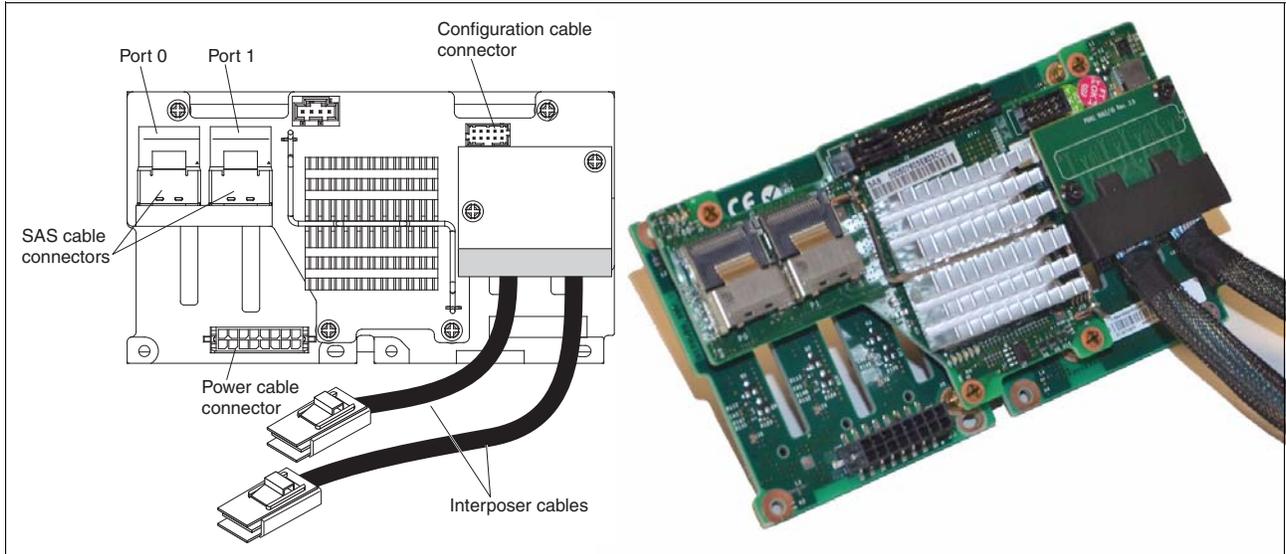


Figure 1-11 8x 2.5-inch drive backplane with controller expander rear view

- ▶ Eight-drive backplane for 1.8-inch SSDs (eXFlash Pack)
Up to four backplanes can be installed, each requiring two SAS connections. One RAID controller is needed per backplane. The 8-drive backplane for 1.8-inch SSDs cannot be connected to the 8-drive backplane with a SAS expander.

Figure 1-12 shows the 8x 1.8-inch backplane.

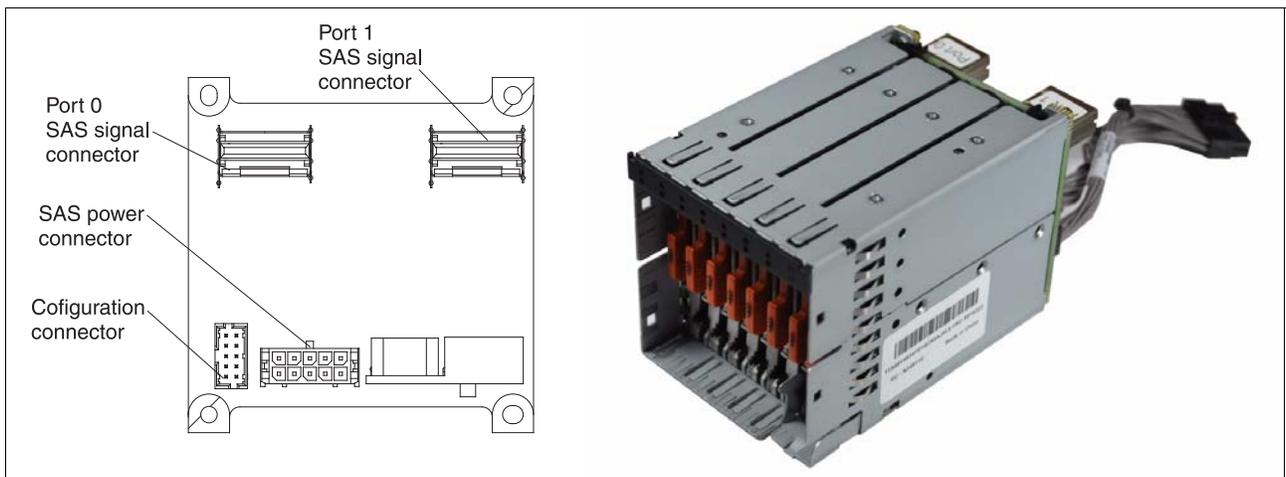


Figure 1-12 8x 1.8-inch drive backplane rear view

When you are building drive backplane configurations, all 1.8-inch SSD drive backplanes must be installed to the right of all 2.5-inch HDD or 2.5-inch SSD backplanes (when you are looking at the front of the server). Every four drives use an SAS signal cable, except when the 8-drive backplane with an SAS expander is used. When the 8-drive backplane with SAS expander is used, the other backplanes connect to the SAS expander backplane with the supplied cables. Then, the SAS expander backplane is connected to the single RAID controller with two cables.

1.9.2 Backplane combinations

The System 3750 M4 can support up to 16 2.5-inch drives, up to 32 1.8-inch drives, or a combination of both 2.5-inch and 1.8-inch hot-swap drives. You can mix drives in the same server if you do not mix drives on the same array.

Backplane installation: Consider the following points about backplane installations:

- ▶ All 2.5-inch backplanes must be installed to the left of all 1.8-inch backplanes.
- ▶ The 8x2.5-inch hot-swap drive backplane with controller expander always must be installed in backplane slots 3 and 4.

The following subsections are sorted by the number of drive bays offered:

- ▶ “Backplane configuration for four drives” on page 30
- ▶ “Backplane configurations for eight drives” on page 31
- ▶ “Backplane configurations for 12 drives” on page 32
- ▶ “Backplane configurations for 16 drives” on page 33
- ▶ “Backplane configurations for 20 drives” on page 37
- ▶ “Backplane configurations for 24 drives” on page 38
- ▶ “Backplane configuration for 28 drives” on page 39
- ▶ “Backplane configuration for 32 drives” on page 39

Tip: All of the necessary SAS cables are included with the backplanes that are used in each configuration.

Backplane configuration for four drives

The configuration that is shown in Figure 1-13 contains one 4x 2.5-inch drive backplane.

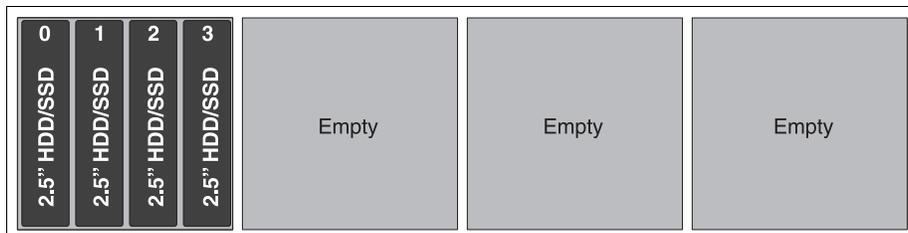


Figure 1-13 One 4x 2.5-inch drive backplane

The following components are required for a 4x 2.5-inch drive backplane configuration:

- ▶ One SAS signal cable
- ▶ One SFF-8087 mini SAS port
- ▶ One RAID controller

Backplane configurations for eight drives

The configuration that is shown in Figure 1-14 contains two 4x 2.5-inch drive backplanes.



Figure 1-14 Two 4x 2.5-inch drive backplanes

The following components are required for a configuration with two 4x 2.5-inch drive backplanes:

- ▶ Two SAS signal cables
- ▶ Two SFF-8087 mini SAS ports
- ▶ One RAID controller

The configuration that is shown in Figure 1-15 contains one 8x 2.5-inch drive backplane.



Figure 1-15 One 8x 2.5-inch drive backplane

The following components are required for a configuration with one 8x 2.5-inch drive backplane:

- ▶ Two SAS signal cables
- ▶ Two SFF-8087 mini SAS ports
- ▶ One RAID controller

The configuration that is shown in Figure 1-16 contains one 8x 1.8-inch drive backplane.

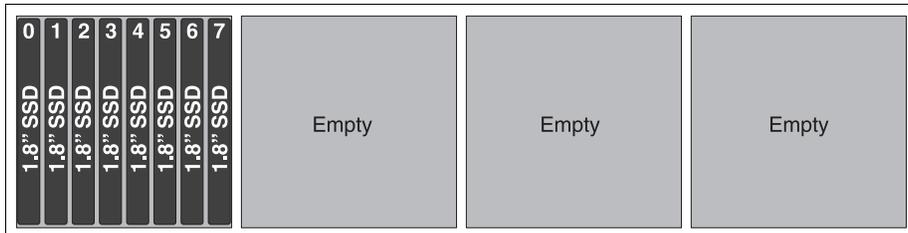


Figure 1-16 One 8x 1.8-inch drive backplane

The following components are required for a configuration with one 8x 1.8-inch drive backplane:

- ▶ Two SAS signal cables
- ▶ Two SFF-8087 mini SAS ports
- ▶ One RAID controller

Backplane configurations for 12 drives

The configuration that is shown in Figure 1-17 has three 4x 2.5-inch drive backplanes.

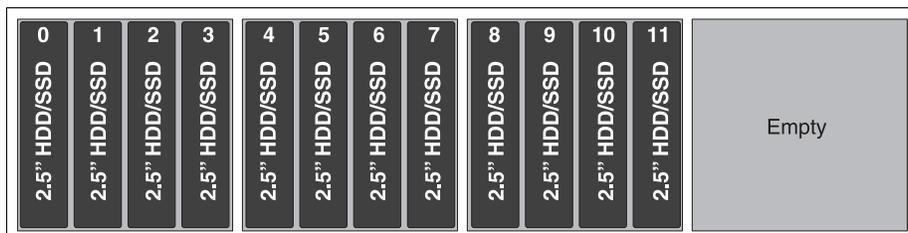


Figure 1-17 Three 4x 2.5-inch drive backplanes

The following components are required for a configuration with three 4x 2.5-inch drive backplanes:

- ▶ Three SAS signal cables
- ▶ Three SFF-8087 mini SAS ports
- ▶ Two RAID controllers

The configuration that is shown in Figure 1-18 has one 8x 2.5-inch and one 4x 2.5-inch drive backplanes.

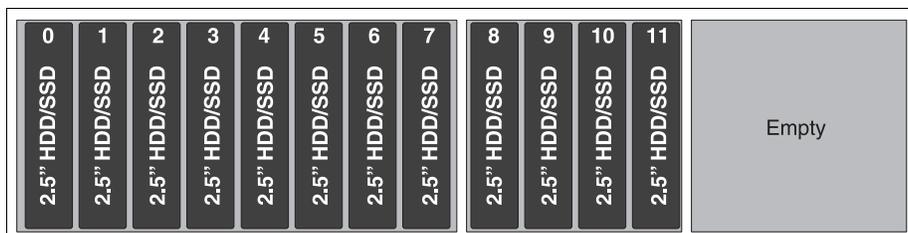


Figure 1-18 One 8x 2.5-inch and one 4x 2.5-inch drive backplanes

The following components are required for a configuration with one 8x 2.5-inch and one 4x 2.5-inch drive backplanes:

- ▶ Three SAS signal cables
- ▶ Three SFF-8087 mini SAS ports
- ▶ Two RAID controllers

The configuration that is shown in Figure 1-19 has one 4x 2.5-inch and one 8x 1.8-inch drive backplanes.

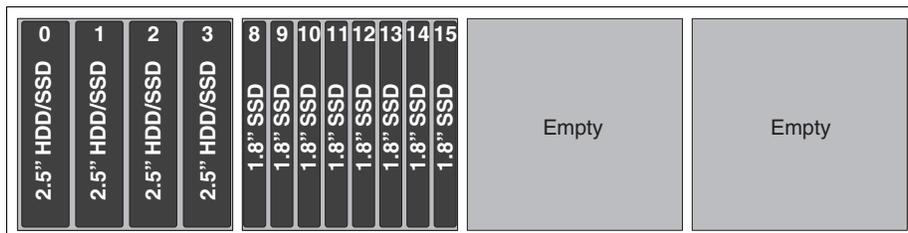


Figure 1-19 One 4x 2.5-inch and one 8x 1.8-inch drive backplanes.

The following components are required for a configuration with one 4x 2.5-inch and one 8x 1.8-inch drive backplanes:

- ▶ Three SAS signal cables
- ▶ Three SFF-8087 mini SAS ports
- ▶ Two RAID controllers

The configuration that is shown in Figure 1-20 has one 4x 2.5-inch and one 8x 2.5-inch drive backplane with controller expander.

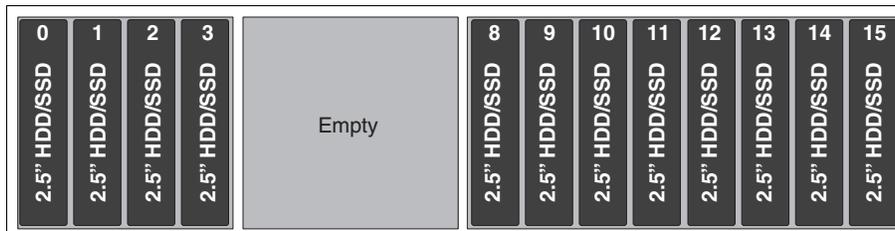


Figure 1-20 One 4x 2.5-inch and one 8x 2.5-inch drive backplane with controller expander

The following components are required for a configuration with one 4x 2.5-inch and one 8x 2.5-inch drive backplane with controller expander:

- ▶ Two SAS signal cables
- ▶ Two SFF-8087 mini SAS ports
- ▶ One RAID controller

Backplane configurations for 16 drives

The configuration that is shown in Figure 1-21 has four 4x 2.5-inch drive backplanes.

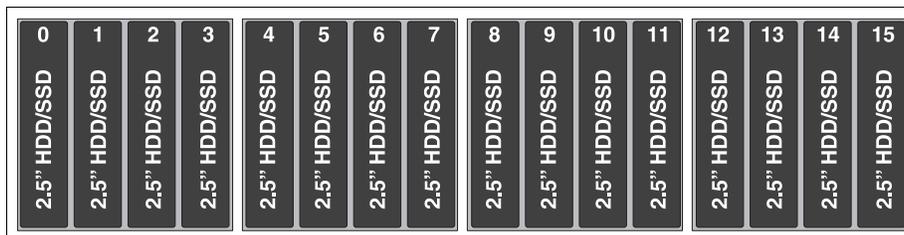


Figure 1-21 Four 4x 2.5-inch drive backplanes

The following components are required for a configuration with four 4x 2.5-inch drive backplanes:

- ▶ Four SAS signal cables
- ▶ Four SFF-8087 mini SAS ports
- ▶ Two RAID controllers

The configuration that is shown in Figure 1-22 has one 8x 2.5-inch and two 4x 2.5-inch drive backplanes.

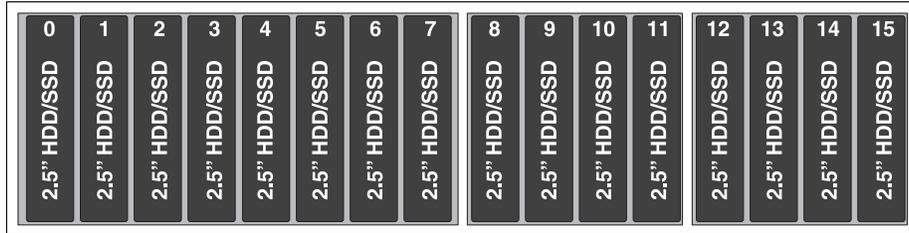


Figure 1-22 One 8x 2.5-inch and two 4x 2.5-inch drive backplanes

The following components are required for a configuration with one 8x 2.5-inch and two 4x 2.5-inch drive backplanes:

- ▶ Four SAS signal cables
- ▶ Four SFF-8087 mini SAS ports
- ▶ Two RAID controllers

The configuration in Figure 1-23 has two 4x 2.5-inch and one 8x 2.5-inch drive backplanes.

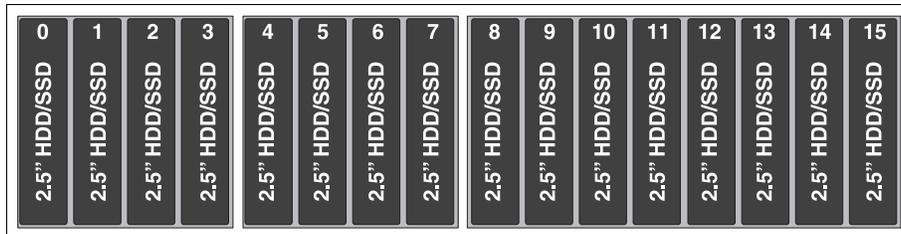


Figure 1-23 Two 4x 2.5-inch and one 8x 2.5-inch drive backplanes

The following components are required for a configuration with two 4x 2.5-inch and one 8x 2.5-inch drive backplanes:

- ▶ Four SAS signal cables
- ▶ Four SFF-8087 mini SAS ports
- ▶ Two RAID controllers

The configuration that is shown in Figure 1-24 has two 8x 2.5-inch drive backplanes.

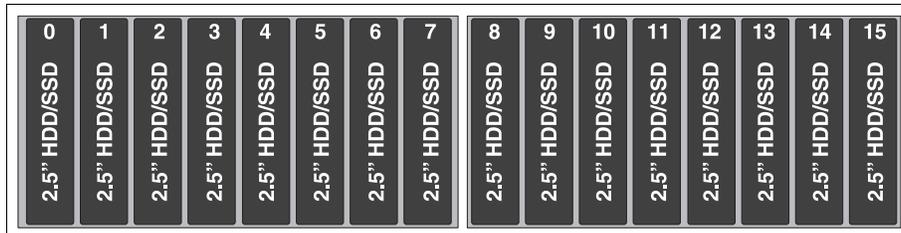


Figure 1-24 Two 8x 2.5-inch drive backplanes

The following components are required for a configuration with two 8x 2.5-inch drive backplanes:

- ▶ Four SAS signal cables
- ▶ Four SFF-8087 mini SAS ports
- ▶ Two RAID controllers

The configuration that is shown in Figure 1-25 has one 8x 2.5-inch and one 8x 2.5-inch drive backplane with controller expander.

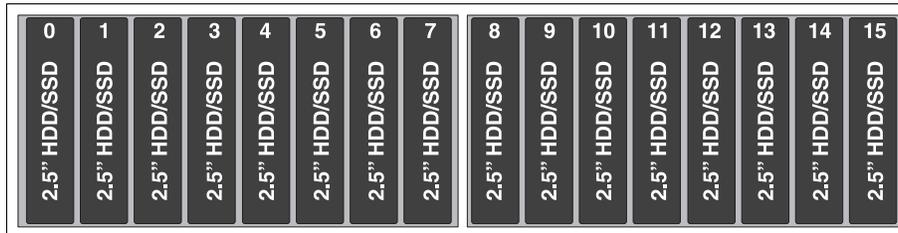


Figure 1-25 One 8x 2.5-inch and one 8x 2.5-inch drive backplane with SAS expander

The following components are required for a configuration with one 8x 2.5-inch drive and one 8x 2.5-inch drive backplane with SAS expander:

- ▶ Two SAS signal cables
- ▶ Two SFF-8087 mini SAS ports
- ▶ One RAID controller

The configuration that is shown in Figure 1-26 has two 4x 2.5-inch and one 8x 1.8-inch drive backplanes.

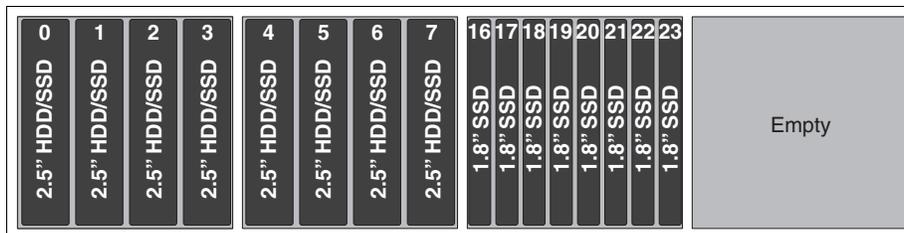


Figure 1-26 Two 4x 2.5-inch and one 8x 1.8-inch drive backplanes

The following components are required for a configuration with two 4x 2.5-inch and one 8x 1.8-inch drive backplanes:

- ▶ Four SAS signal cables
- ▶ Four SFF-8087 mini SAS ports
- ▶ Two RAID controllers

The configuration that is shown in Figure 1-27 has one 8x 2.5-inch and one 8x 1.8-inch drive backplanes.

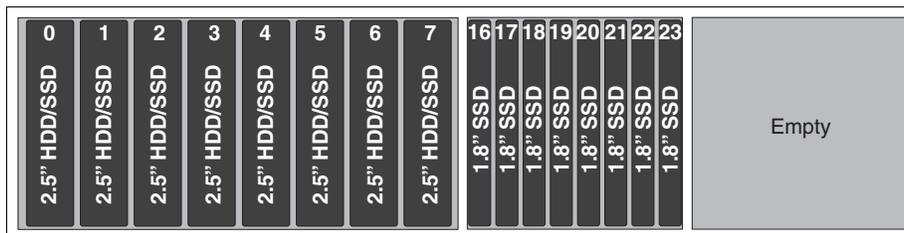


Figure 1-27 One 8x 2.5-inch and one 8x 1.8-inch drive backplanes

The following components are required for a configuration with one 8x 2.5-inch and one 8x 1.8-inch drive backplanes:

- ▶ Four SAS signal cables
- ▶ Four SFF-8087 mini SAS ports
- ▶ Two RAID controllers

The configuration that is shown in Figure 1-28 has two 8x 1.8-inch drive backplanes.

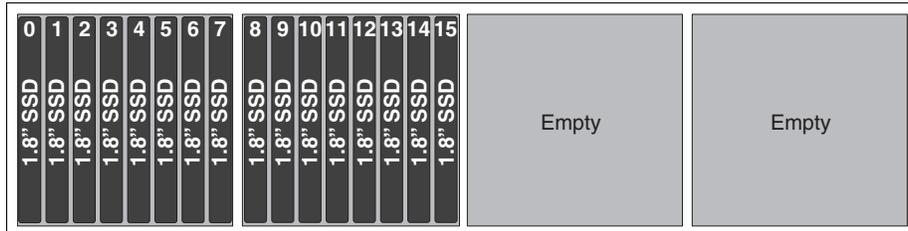


Figure 1-28 Two 8x 1.8-inch drive backplanes

The following components are required for a configuration with two 8x 1.8-inch drive backplanes:

- ▶ Four SAS signal cables
- ▶ Four SFF-8087 mini SAS ports
- ▶ Two RAID controllers

The configuration that is shown in Figure 1-29 has two 4x 2.5-inch and one 8x 2.5-inch drive backplane with controller expander.

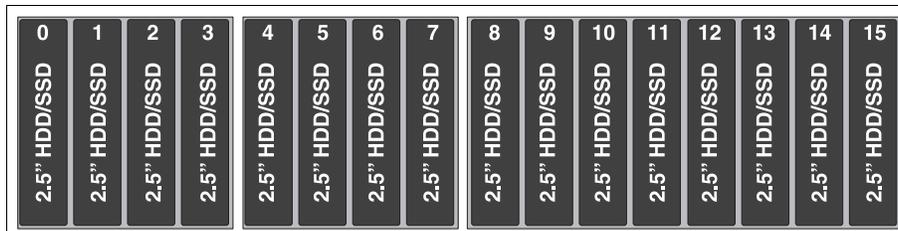


Figure 1-29 Two 4x 2.5-inch and one 8x 2.5-inch drive backplane with controller expander

The following components are required for a configuration with two 4x 2.5-inch and one 8x 2.5-inch drive backplane with controller expander:

- ▶ Two SAS signal cables
- ▶ Two SFF-8087 mini SAS ports
- ▶ One RAID controller

Backplane configurations for 20 drives

The configuration that is shown in Figure 1-30 has three 4x 2.5-inch and one 8x 1.8-inch drive backplanes.

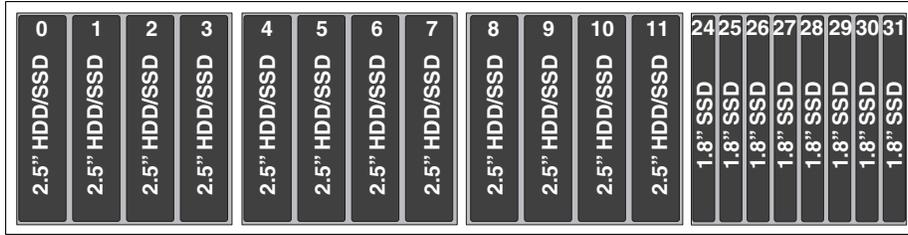


Figure 1-30 Three 4x 2.5-inch and one 8x 1.8-inch drive backplanes

The following components are required for a configuration with three 4x 2.5-inch and one 8x 1.8-inch drive backplanes:

- ▶ Five SAS signal cables
- ▶ Five SFF-8087 mini SAS ports
- ▶ Three RAID controllers

The configuration that is shown in Figure 1-31 has one 8x 2.5-inch, one 4x 2.5-inch, and one 8x 1.8-inch drive backplanes.

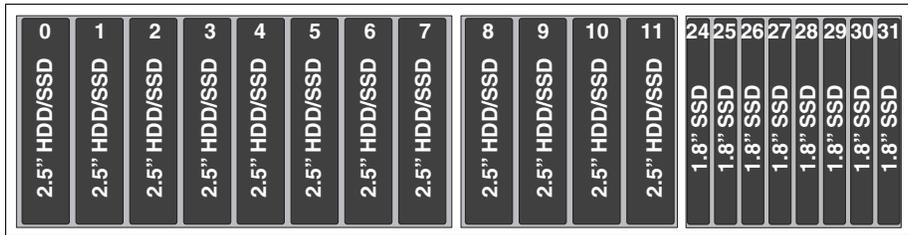


Figure 1-31 One 8x 2.5-inch, one 4x 2.5-inch, and one 8x 1.8-inch drive backplanes

The following components are required for a configuration with one 8x 2.5-inch, one 4x 2.5-inch, and one 8x 1.8-inch drive backplanes:

- ▶ Five SAS signal cables
- ▶ Five SFF-8087 mini SAS ports
- ▶ Three RAID controllers

The configuration that is shown in Figure 1-32 has one 4x 2.5-inch and two 8x 1.8-inch drive backplanes.

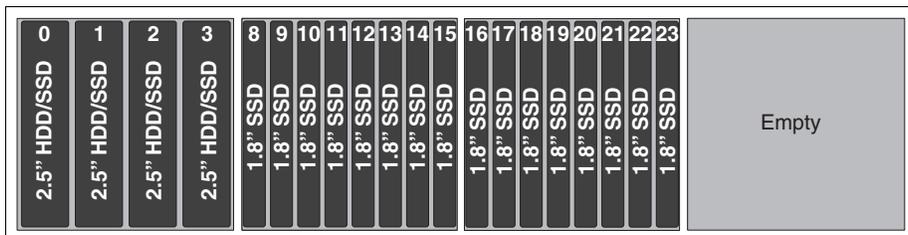


Figure 1-32 One 4x 2.5-inch and two 8x 1.8-inch drive backplanes

The following components are required for a configuration with one 4x 2.5-inch and two 8x 1.8-inch drive backplanes:

- ▶ Five SAS signal cables
- ▶ Five SFF-8087 mini SAS ports
- ▶ Three RAID controllers

Backplane configurations for 24 drives

The configuration that is shown in Figure 1-33 has two 4x 2.5-inch and two 8x 1.8-inch drive backplanes.

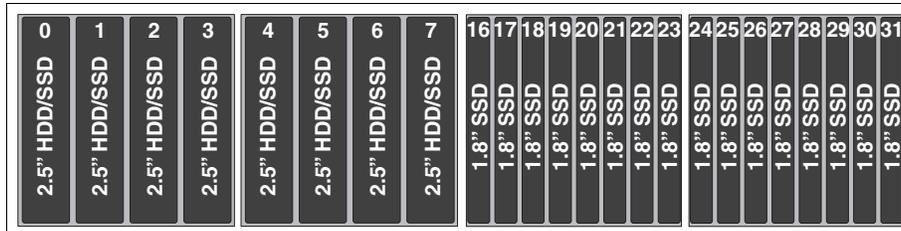


Figure 1-33 Two 4x 2.5-inch and two 8x 1.8-inch drive backplanes

The following components are required for a configuration with two 4x 2.5-inch and two 8x 1.8-inch drive backplanes:

- ▶ Six SAS signal cables
- ▶ Six SFF-8087 mini SAS ports
- ▶ Three RAID controllers

The configuration that is shown in Figure 1-34 has one 8x 2.5-inch and two 8x 1.8-inch drive backplanes.

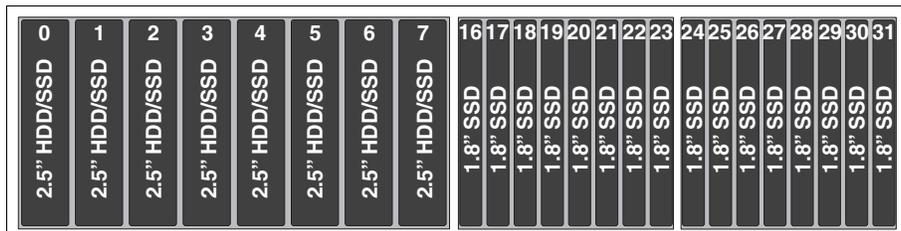


Figure 1-34 One 8x 2.5-inch and two 8x 1.8-inch drive backplanes

The following components are required for a configuration with one 8x 2.5-inch and two 8x 1.8-inch drive backplanes:

- ▶ Six SAS signal cables
- ▶ Six SFF-8087 mini SAS ports
- ▶ Three RAID controllers

The configuration that is shown in Figure 1-35 has three 8x 1.8-inch drive backplanes.

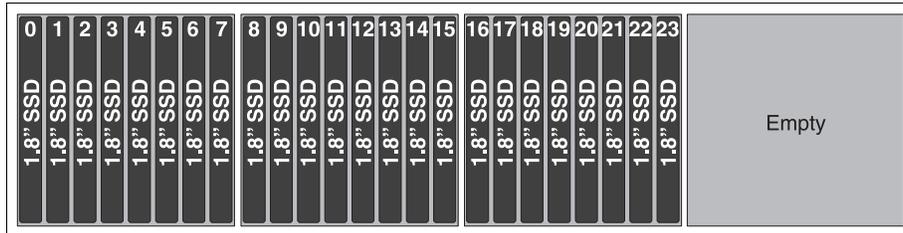


Figure 1-35 Three 8x 1.8-inch drive backplanes

The following components are required for a configuration with three 8x 1.8-inch drives:

- ▶ Six SAS signal cables
- ▶ Six SFF-8087 mini SAS ports
- ▶ Three RAID controllers

Backplane configuration for 28 drives

The configuration that is shown in Figure 1-36 has one 4x 2.5-inch and three 8x 1.8-inch drive backplanes.

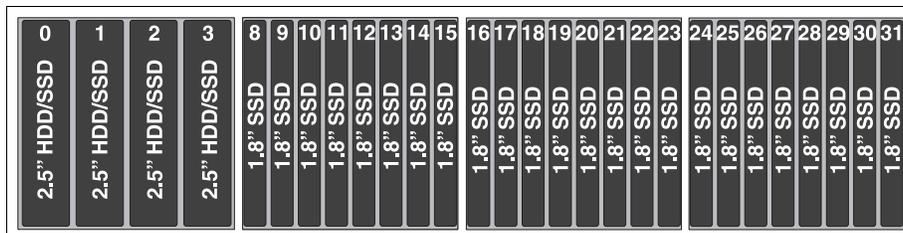


Figure 1-36 One 4x 2.5-inch and three 8x 1.8-inch drive backplanes

The following components are required for a configuration with one 4x 2.5-inch and three 8x 1.8-inch drive backplanes:

- ▶ Seven SAS signal cables
- ▶ Seven SFF-8087 mini SAS ports
- ▶ Four RAID controllers

Backplane configuration for 32 drives

The configuration that is shown in Figure 1-37 has four 8x 1.8-inch drive backplanes.

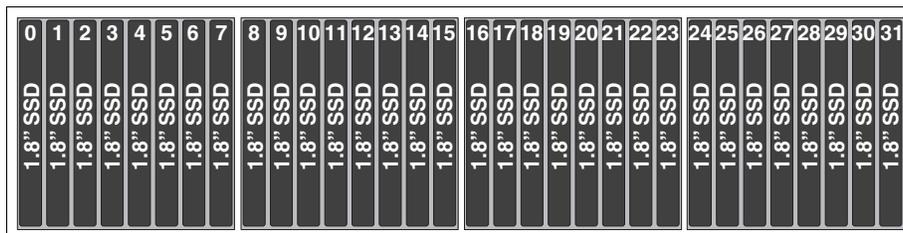


Figure 1-37 Four 8x 1.8-inch drive backplanes

The following components are required for a configuration with four 8x 1.8-inch drive backplanes:

- ▶ Eight SAS signal cables
- ▶ Eight SFF-8087 mini SAS ports
- ▶ Four RAID controllers

1.9.3 RAID controllers

The x3750 M4 has the following integrated SAS RAID controller:

- ▶ Machine type 8752 (with Xeon E5-4600 v2 processors): SeveRAID M5210e SAS/SATA Controller, which is based on LSI SAS3108 chipset
- ▶ Machine type 8722 (with Xeon E5-4600 processors): SeveRAID M5210e SAS/SATA Controller, which is based on LSI SAS2208 chipset

The server also supports more RAID controllers when more than eight drives are installed. Table 1-16 lists the RAID controllers that are used for the internal disk storage of the x3750 M4 server.

Table 1-16 RAID controllers for internal storage

Part number	Feature code	Description	Supported on 8752 (E5-4600 v2)	Supported on 8722 (E5-4600)	Maximum supported
None	A3QA	SeveRAID M5210e SAS/SATA Controller	Yes	No	Integrated
46C9110	A3YZ	ServeRAID M5210 SAS/SATA Controller	Yes	No	3
None	A2N2	ServeRAID M5110e SAS/SATA Controller	No	Yes	Integrated
81Y4481	A347	ServeRAID M5110 SAS/SATA Controller	No	Yes	3
81Y4448	A1MZ	ServeRAID M1115 SAS/SATA Controller	No	Yes	1
46M0912	3876	6 Gb Performance Optimized HBA	No	Yes	3
46C8988	A3MW	N2115 SAS/SATA HBA for System x	No	Yes	3
47C8675	A3YY	N2215 SAS/SATA HBA for System x	Yes	No	3

For the cable routing requirements, it is necessary to install the supported RAID controllers in PCIe slots 1, 2, and 3. Consider the following general rules when you are connecting the SAS signal cables to the RAID controller and drive backplane:

- ▶ Connect Port 0 on the adapter to port 0 on the drive backplane.
- ▶ Connect Port 1 on the adapter to port 1 on the drive backplane.

For more information about installation, see *System x3750 M4 Installation and Service Guide*, which is available at this website:

<http://www-947.ibm.com/support/entry/portal/docdisplay?lnodocid=MIGR-5090828>

Multiple RAID controllers: If you have multiple RAID controllers that are installed, each one manages the drives and backplanes that are connected to it and does not share configuration information with other adapters. You cannot span arrays (volumes) across RAID controllers other than through the operating system. If you need more drives connected to one controller, consider the use of an 8x 2.5-inch backplane with SAS expander, as described in 1.9.1, “Backplanes” on page 27.

The integrated ServeRAID M5210e SAS/SATA controller is in a chip on the system board for machine type 8752. It has the following specifications:

- ▶ LSI SAS3108 12 Gbps RAID on Chip (ROC) controller
- ▶ Eight internal SAS/SATA ports (the controller supports 12 Gbps, but the supported backplanes and drives operate only at 6 Gbps)
- ▶ Two internal x4 HD Mini-SAS connectors (SFF-8643)
- ▶ PCI Express 3.0 x8 host interface
- ▶ Support for RAID 0, 1, and 10 as standard
- ▶ Support for RAID 5, 50, 6, and 60 with optional upgrades
- ▶ Optional onboard data cache: 1 GB (no battery backup), or 1 GB, 2 GB, or 4 GB with flash backup
- ▶ Optional support for self-encrypting drives (SEDs) with MegaRAID SafeStore

Note: The drives and backplanes that are supported at the time of this writing in the x3750 M4 operate at 6 Gbps speeds, so the M5210e also operates at 6 Gbps.

The integrated ServeRAID M5110e SAS/SATA controller is a chip on the system board for machine type 8722. The ServeRAID M5110 is a PCIe low-profile, half-length (MD2) form factor adapter. The following specifications are featured with both controllers:

- ▶ Two Mini-SAS internal connectors
- ▶ 6 Gbps throughput for each port
- ▶ PCIe x8 Gen three-host interface
- ▶ LSI SAS2208 6 Gbps RAID on Chip (ROC) controller-based
- ▶ Supports RAID levels 0, 1, and 10
- ▶ Supports RAID levels 5 and 50 with optional M5100 Series RAID 5 upgrades
- ▶ Supports RAID 6 and 60 with the optional M5100 Series RAID 6 Upgrade
- ▶ Supports 512 MB battery-backed cache or 512 MB or 1 GB flash-backed cache

For more information about these upgrades, see 1.9.4, “RAID controller upgrades” on page 42.

The ServeRAID M1115 offers a low-cost RAID 0/1/10 solution that can be upgraded to a cacheless RAID 5 with a Features on Demand (FoD) license upgrade. The following specifications are included with the ServeRAID M1115 standard solution:

- ▶ PCI low profile, half-length, MD2 form factor
- ▶ Two internal Mini-SAS connectors (SFF-8087)
- ▶ Eight internal 6 Gbps SAS/SATA ports
- ▶ 6 Gbps throughput per port
- ▶ 533 MHz PowerPC processor with an LSI SAS2008 6 Gbps ROC controller
- ▶ PCI Express 2.0 x8 host interface
- ▶ Support for RAID levels 0, 1, and 10 as standard
- ▶ Support for RAID 5 and 50 with an optional upgrade (81Y4542)
- ▶ Zero Controller Cache, no battery or flash backup

For more information, see the list of Lenovo Press Product Guides in the RAID adapters category:

<https://lenovopress.com/servers/options/raid>

Installation: The supported adapters for internal storage must be installed in slots 1, 2, and 3 because of cable routing requirements. One of the supported riser cards must be installed (for more information, see Table 1-23 on page 49). These slots also require that Processor 2 is installed.

1.9.4 RAID controller upgrades

Table 1-17 lists the optional upgrades that are supported by the ServeRAID controllers.

Table 1-17 RAID controller upgrades

Part number	FC	Description	Upgrade type ^a	Max supt.	Supported on 8752 (E5-4600 v2)	Supported on 8722 (E5-4600)	Models where used
Upgrades for the ServeRAID M5210e and M5210 SAS/SATA controller							
47C8656	A3Z0	ServeRAID M5200 Series 1GB Cache/RAID 5 Upgrade	Hardware	4	Yes	No	-
47C8660	A3Z1	ServeRAID M5200 Series 1GB Flash/RAID 5 Upgrade	Hardware	4	Yes	No	8752-C1x, C2x
47C8664	A3Z2	ServeRAID M5200 Series 2GB Flash/RAID 5 Upgrade	Hardware	4	Yes	No	-
47C8706	A3Z5	ServeRAID M5200 Series RAID 6 Upgrade-FoD	FoD	1 ^b	Yes	No	-
47C8708	A3Z6	ServeRAID M5200 Series Zero Cache/RAID 5-FoD	FoD	1	Yes	No	-
47C8710	A3Z7	ServeRAID M5200 Series Performance Accelerator-FoD	FoD	1 ^b	Yes	No	-
47C8712	A3Z8	ServeRAID M5200 Series SSD Caching Enabler-FoD	FoD	1 ^b	Yes	No	-
Upgrades for the ServeRAID M5110e and M5110 SAS/SATA controller							
81Y4484	A1J3	ServeRAID M5100 Series 512 MB Cache/RAID 5 Upgrade	Hardware	4	No	Yes	8722-D1x (2 std) D2x (4 std)
81Y4487	A1J4	ServeRAID M5100 Series 512 MB Flash/RAID 5 Upgrade	Hardware	4	No ^c	Yes	-
81Y4559	A1WY	ServeRAID M5100 Series 1 GB Flash/RAID 5 Upgrade	Hardware	4	No ^c	Yes	8722-C1x (1 std) C2x (1 std)
81Y4508	A22E	ServeRAID M5100 Series Battery Kit	Hardware	4 ^d	No	Yes	8722-D1x (2 std) D2x (4 std)
81Y4546	A1X3	ServeRAID M5100 Series RAID 6 Upgrade	FoD	1 ^{e,f}	No ^c	Yes	-
90Y4273	A2MC	ServeRAID M5100 Series SSD Performance Key	FoD	1 ^f	No ^c	Yes	8722-D1x, D2x

Part number	FC	Description	Upgrade type ^a	Max supt.	Supported on 8752 (E5-4600 v2)	Supported on 8722 (E5-4600)	Models where used
90Y4318	A2MD	ServeRAID M5100 Series SSD Caching Enabler	FoD	1 ^f	No ^c	Yes	-
81Y4544	A1X2	ServeRAID M5100 Series Zero Cache/RAID 5 Upgrade	FoD	1 ^f	No	Yes	-
Upgrades for the ServeRAID M1115 SAS/SATA Controller							
81Y4542	A1X1	ServeRAID M1100 Series Zero Cache/RAID 5 Upgrade	FoD	1 ^f	No	Yes	-

- a. FoD is Features on Demand, a software license that is applied to the server via IMMv2 or other management tools.
- b. M5200 RAID upgrades require one of the cache upgrades (47C8656, 47C8660, or 47C8664).
- c. Supported on 8752, but only with the ServeRAID M5120 for external disk connectivity..
- d. The ServeRAID M5100 Series Battery Kit (81Y4508) is supported by the ServeRAID M5100 Series 512 MB Cache/RAID 5 Upgrade (81Y4484) only.
- e. The ServeRAID M5100 Series RAID 6 Upgrade (81Y4546) requires RAID 5 upgrades with cache (81Y4484, 81Y4487, or 81Y4559 only).
- f. Only one ServeRAID FoD upgrade is required for each system, regardless of the number of adapters that are installed.

The ServeRAID M5100 and M5200 series adapters provide support for RAID 0, 1, and 10 as standard capabilities. More functional upgrades are available as an option to expand the standard capabilities. Some upgrades do not depend on other upgrades and can be applied to the default settings of the controllers (which are referred to as *primary upgrades*). Certain upgrades cannot be applied to the default settings of the controllers and require that other upgrades are enabled before these upgrades are applied (which are referred to as *secondary upgrades*).

There are two types of available upgrades: Hardware (HW) and FoD. HW upgrades contain physical parts (for example, cache module or battery). FoD upgrades are software licenses.

Table 1-18 lists the available primary upgrades for the M5200 adapters, their capabilities, and types.

Table 1-18 ServeRAID M5100 Series primary upgrades and their features

Part number	Description	Type	RAID 5, RAID-50	SED	Cache	Flash-backed cache
47C8708	ServeRAID M5200 Series Zero Cache/RAID 5-FoD	FoD	Yes	Yes	None	No
47C8656	ServeRAID M5200 Series 1GB Cache/RAID 5 Upgrade	HW	Yes	Yes	1 GB	No
47C8660	ServeRAID M5200 Series 1GB Flash/RAID 5 Upgrade	HW	Yes	Yes	1 GB	Yes
47C8664	ServeRAID M5200 Series 2GB Flash/RAID 5 Upgrade	HW	Yes	Yes	2 GB	Yes

Table 1-19 lists the available primary upgrades for the M5100 adapters, their capabilities, and types.

Table 1-19 ServeRAID M5100 Series primary upgrades and their features

Part number	Option description	Type	RAID-5, RAID-50	SED	Cache	Flash-backed cache
81Y4544	ServeRAID M5100 Zero Cache/RAID 5 Upgrade	FoD	Yes	Yes	None	No
81Y4484	ServeRAID M5100 512MB Cache/RAID 5 Upgrade	HW	Yes	Yes	512 MB	No
81Y4487	ServeRAID M5100 512MB Flash/RAID 5 Upgrade	HW	Yes	Yes	512 MB	Yes
81Y4559	ServeRAID M5100 1GB Flash/RAID 5 Upgrade	HW	Yes	Yes	1 GB	Yes

Table 1-20 shows the secondary upgrades of the M5200 series adapters, their capabilities, types, and dependencies. The primary feature upgrades, on which the secondary upgrades depend, are listed in the appropriate columns. *Required* means that the primary upgrade that is listed in the column must be enabled before the secondary feature is enabled that is listed in that particular row.

Table 1-20 ServeRAID M5200 series secondary upgrades, their features, and dependencies

Primary feature upgrades →	Option description		Zero cache/ RAID 5	1GB cache/ RAID 5	1GB flash/ RAID 5	2GB flash/ RAID 5
	Part number		47C8708	47C8656	47C8660	47C8664
	Upgrade type		FoD	HW	HW	HW
↓ Secondary feature upgrades	Part number	Type				
RAID 6, 60	47C8706	FoD	No support	Required	Required	Required
SSD Performance Accelerator (FastPath)	47C8710	FoD	No support	Required	Required	Required
SSD Caching Enabler (CacheCade Pro 2.0)	47C8712	FoD	No support	Required	Required	Required

Table 1-21 shows the secondary upgrades of the M5100 series adapters

Table 1-21 *ServeRAID M5100 series secondary upgrades, their features, and dependencies*

Primary feature upgrades →	Option description:		Zero Cache/ RAID 5	512 MB Cache/ RAID 5	512 MB Flash/ RAID 5	1 GB Flash/ RAID 5
	Part number:		81Y4544	81Y4484	81Y4487	81Y4559
	Upgrade type:		FoD	HW	HW	HW
↓ Secondary feature upgrades	Part number	Type				
Battery Kit (Battery-backed cache)	81Y4508	HW	No support	Required	No support	No support
RAID 6, 60	81Y4546	FoD	No support	Required	Required	Required
SSD Performance Accelerator (FastPath)	90Y4273	FoD	No support	Required	Required	Required
SSD Caching Enabler (CacheCade Pro 2.0)	90Y4318	FoD	No support	Required	Required	Required

For more information, see the following Lenovo Press Product Guides:

- ▶ ServeRAID M5210 and M5210e SAS/SATA Controllers for System x:
<http://lenovopress.com/tips1069>
- ▶ ServeRAID M5110 and M5110e SAS/SATA Controllers for System x:
<http://lenovopress.com/tips0857>

1.9.5 Battery mounting

The onboard RAID controller and some supported RAID adapters support battery backup in the form of an actual battery or a flash power module, depending on the options that are selected. These batteries and flash power modules must be located remotely to prevent them from overheating. In the x3750 M4, these modules are installed in a 4-bay unit at the front of the server, which is above the USB and video ports and the optical drive bay. Figure 1-38 shows the reverse angle.



Figure 1-38 *Battery and flash power module tray (system fans removed)*

The cable from the battery is routed to the ServeRAID adapters that are installed in PCIe slots 1, 2, or 3. For the onboard RAID controller, the battery is routed to the RAID cache card that is installed in a dedicated holder on the system board, next to the RAID controller, as shown in Figure 1-39.

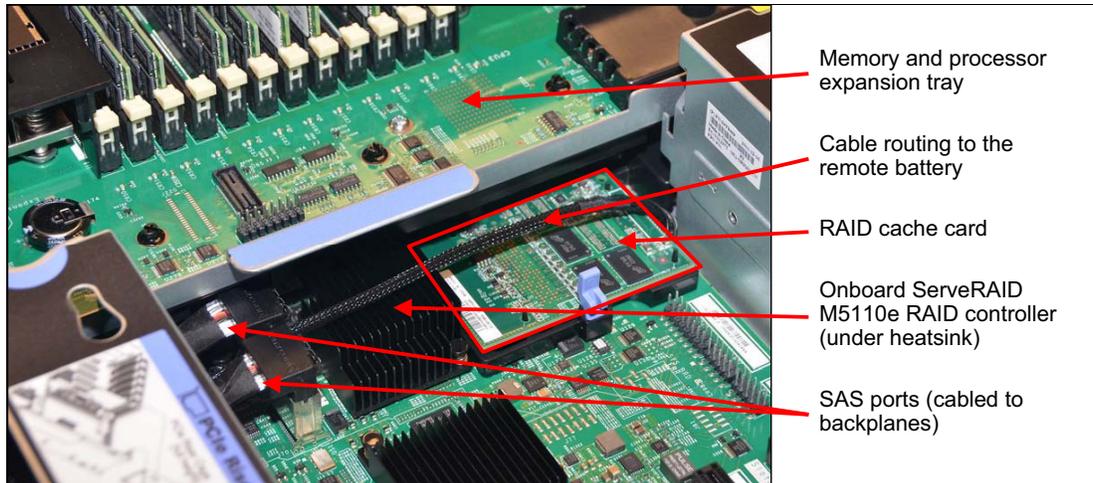


Figure 1-39 Location of the ServeRAID M5110e RAID controller and cache option

1.9.6 Battery versus flash

When the flash power module is attached to flash memory, it eliminates the need for a battery and the requirement to replace that battery every 12 months.

The battery with cache and the flash power module with flash allow a write-back policy for more performance. In the write-back mode, data is cached in dynamic random access memory (DRAM) and then periodically flushed out to the storage media.

During a power-loss event, these devices work in the following manner:

- ▶ The Battery Back Up (BBU) hardware puts the DRAM in self-refresh and the battery provides the power to the DRAM modules; therefore, it keeps it alive. The battery pack can keep the DRAM powered for a few days (the industry standard is 48 hours).
- ▶ For the flash power module with flash, the power isolation logic begins the cache offload process by turning on the flash power module and isolating this power source for the server. The flash power module keeps the RAID subsystem powered during a power-loss event only and the RAID controller starts copying the data contents from cache to non-volatile flash. After this offload is complete, the cache is fully protected by being copied onto the non-volatile flash.

1.9.7 Drive selection

The System 3750 M4 offers the following types of drives from which to choose:

- ▶ SATA drives
- ▶ SAS drives
- ▶ SSD drives

The sizing of the drives is a complicated topic because it depends on which application you installed or which workload it generates. You must know how many IOPS are handled. For more information about the best sizing for your specific environment, application, and I/O needs, contact your Business Partner or Lenovo representative.

The server supports up to 16 2.5-inch drives, up to 32 1.8-inch drives, or a combination of 2.5-inch and 1.8-inch hot-swap drives, as described in 1.9.2, “Backplane combinations” on page 30.

SSD drives have the following advantages over HDDs:

- ▶ No moving parts
- ▶ Higher performance
- ▶ Lower power consumption
- ▶ Less susceptible to physical shock and vibration
- ▶ Lower access times and latency rates

For more information about drive technology, see *System x Server Disk Drive Interface Technology*, REDP-4791, which is available at this website:

<http://lenovopress.com/redp4791>

All supported drives are listed in 1.9.8, “Drive options” on page 47.

You can mix SATA, SAS, and SSD drives if they are not on the same array. This configuration can move the hot data on SSD drives. Alternatively, you can use the ServeRAID M5100 or M5200 SSD Caching Enabler feature to use SSD drives as a cache for SAS drives and to increase overall system performance. This and other ServeRAID features are described in 1.9.4, “RAID controller upgrades” on page 42.

1.9.8 Drive options

For a list of drive options see the Lenovo Press product guides:

- ▶ Machine type 8753: <https://lenovopress.com/tips1252>
- ▶ Machine type 8752: <http://lenovopress.com/tips1159>
- ▶ Machine type 8722: <http://lenovopress.com/tips0881>

1.10 Optical drive

Table 1-22 lists the optical drive option that is supported by the x3750 M4 server. This option is supported on machine type 8752 (Xeon E5-4600 v2) and machine type 8722 (Xeon E5-4600).

Table 1-22 Optical drives

Part number	Feature code	Description	Maximum supported	Standard models where used
46M0902	4163	UltraSlim Enhanced SATA Multi-Burner	1	-

The UltraSlim Enhanced SATA Multi-Burner (46M0902) supports the following media and speeds for reading:

- ▶ CD-ROM 24X
- ▶ CD-DA (DAE) 20X
- ▶ CD-R 24X
- ▶ CD-RW 24X
- ▶ DVD-ROM (single layer) 8X
- ▶ DVD-ROM (dual layer) 8X
- ▶ DVD-R (4.7 GB) 6X
- ▶ DVD-R DL 4X
- ▶ DVD+R 6X
- ▶ DVD+R DL 4X
- ▶ DVD-RW (4.7 GB) 4X
- ▶ DVD+RW 4X
- ▶ DVD-RAM (4.7/9.4 GB) 4X

The drive also supports the following media and speeds for writing:

- ▶ CD-R 24X
- ▶ CD-RW 4X
- ▶ High Speed CD-RW 10X
- ▶ Ultra Speed CD-RW 16X
- ▶ Ultra Speed Plus CD-RW 16X
- ▶ DVD-R 8X
- ▶ DVD-R DL 6X
- ▶ DVD+R 8X
- ▶ DVD+R DL 6X
- ▶ DVD-RW 6X
- ▶ DVD+RW 8X
- ▶ DVD-RAM 5X

1.11 PCI Express slots

The PCI Express (PCIe) controller on the x3750 M4 is integrated into the processor and is not controlled by a separate chip. For more information, see 1.6, “Architecture” on page 12.

The server supports up to eight PCIe slots, not including the special adapter slot for the integrated Ethernet adapter.

Five slots are on the system board and three or two are through a riser card. The riser card is standard on some models and optional on others. The use of slots 1 - 5 requires processor 2 to be installed.

There are two riser cards available for x3750 M4 machine type 8752 (with Xeon E5-4600 v2 processors) as listed in Table 1-23. One one of them is supported in machine type 8722 (with Xeon E5-4600 processors).

Table 1-23 PCI riser card options

Part number	Feature code	Description	Supported on 8752 (E5-4600 v2)	Supported on 8722 (E5-4600)	Maximum supported
88Y7371	A2A1	x3750 M4 PCIe 3 x8 riser	Yes	Yes	1 ^a
00D2009	A419	x3750 M4 PCIe 1x16 and 1x8 riser	Yes	No	1 ^a

a. These riser cards occupy the same riser slot and are mutually exclusive.

The PCIe 3 x8 riser includes the following slots:

- ▶ Slot 1: PCIe 3.0 x8 (x16 mechanical): Full-height, half-length (requires processor 2)
- ▶ Slot 2: PCIe 3.0 x8 (x16 mechanical): Full-height, half-length (requires processor 2)
- ▶ Slot 3: PCIe 3.0 x8 (x16 mechanical): Full-height, half-length (requires processor 2)

The PCIe 1x16 and 1x8 riser includes the following slots:

- ▶ Slot 1: PCIe 3.0 x16: Full-height, half-length (requires processor 2)
- ▶ Slot 2: PCIe 3.0 x8 (x16 mechanical): Full-height, half-length (requires processor 2)
- ▶ Slot 3: Not present

The following five standard slots are on the system board:

- ▶ Slot 4: PCIe 3.0 x8: Low profile, half-length (requires processor 2)
- ▶ Slot 5: PCIe 3.0 x8: Low profile, half-length (requires processor 2)
- ▶ Slot 6: PCIe 3.0 x8: Low profile, half-length
- ▶ Slot 7: PCIe 3.0 x8: Low profile, half-length
- ▶ Slot 8: PCIe 3.0 x8: Low profile, half-length

The server offers up to eight PCIe Gen 3 slots. You need the following options to use all of the slots:

- ▶ Processor 2 installed
- ▶ x3750 M4 PCIe 3 x8 riser or PCIe 1x16 and 1x8 riser

Table 1-24 shows slot availability if you have processor 2 and a riser card installed.

Table 1-24 PCI slots

PCI slot	Processor 1	Processor 2 without riser card	Processor 2 with PCIe 3 x8 riser card	Processor 2 with PCIe 1x16 and 1x8 riser
1	Not available	Not available	Present	Present
2	Not available	Not available	Present	Present
3	Not available	Not available	Present	Not available
4	Not available	Present	Present	Present
5	Not available	Present	Present	Present
6	Present	Present	Present	Present
7	Present	Present	Present	Present
8	Present	Present	Present	Present

PCIe Gen 3 doubles the bandwidth. PCIe Gen 3 with Intel Data Direct I/O also decreases the latency. Table 1-25 provides a comparison between PCIe Gen 3 and Gen 2.

Table 1-25 PCIe Gen3 in comparison to PCIe Gen2

PCIe architecture	Data interface	Signaling rate	Bidirectional bandwidth	Bandwidth/lane/direction	Total bandwidth for x16 link
PCIe 2.0	Full duplex	5 Gbps	4 Gbps	~500 Mbps	~16 Gbps
PCIe 3.0	Full duplex	8 Gbps	8 Gbps	~1 Gbps	~32 Gbps

1.12 Integrated networking: Type 8753 and 8752

The x3750 M4, machine type 8753 and 8752 (which use Intel Xeon E5-4600 v2 processors) offers a dedicated mezzanine LAN-on-motherboard (mezz LOM or ML2) adapter slot for integrated networking. In addition to the usual networking functions for the installed operating system, this adapter supports shared access to the IMM2 service processor if so enabled in UEFI. Table 1-26 lists the supported ML2 adapters.

Table 1-26 Adapters for the ML2 slot (machine type 8752 and 8753)

Part number	Feature code	Description	Maximum supported
10 Gb Ethernet			
00D2026	A40S	Broadcom NetXtreme II ML2 Dual Port 10 GbaseT for System x	1 ^a
00D2028	A40T	Broadcom NetXtreme II ML2 Dual Port 10 GbE SFP+ for System x	1 ^a
00D1994	A40P	Intel X540 ML2 Dual Port 10 GbaseT Adapter for System x	1 ^a
1 Gb Ethernet			
00D1998	A40R	Intel I350-T4 ML2 Quad Port GbE Adapter for System x	1 ^a

a. These ML2 adapters occupy the same dedicated slot and are mutually exclusive.

Figure 1-40 shows mezzanine LOM adapter that is installed in the server.

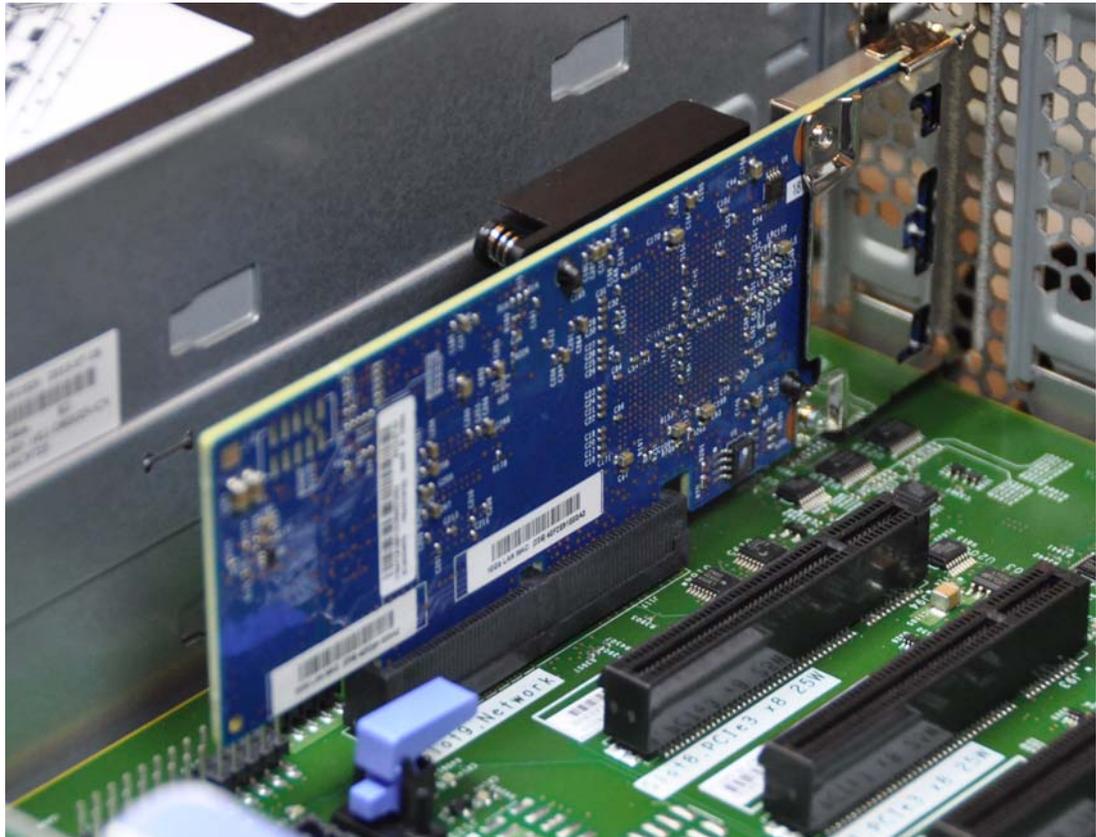


Figure 1-40 Mezzanine LOM networking adapter that is installed in the x3750 M4 type 8752

1.13 Integrated networking: Type 8722

The x3750 M4, machine type 8722 (which uses Intel Xeon E5-4600 processors) has an onboard Emulex BE3 controller. This controller drives two standard Gigabit Ethernet ports and two 10 Gb Ethernet ports.

The Two 10 Gb Ethernet ports are accessed by installing one of the available dual-port 10 Gb Ethernet cards. Table 1-27 lists the options (88Y7429 or 88Y7427). These cards use a dedicated connector on the system board and do not use any of the eight regular PCIe expansion slots. The optional cards are used to route the ports to external RJ45 or SFP+ ports.

Table 1-27 Adapters for the integrated 10 GbE controller (machine type 8722 only)

Part number	Description	Maximum supported
10 Gb Ethernet (does not use a PCI expansion slot)		
88Y7427	x3750 M4 Dual port 10 GB-T Ethernet Adapter	1 ^a
88Y7429	x3750 M4 Dual port 10 GB SFP+ Ethernet Adapter	1 ^a
95Y3760	Emulex VFA III FCoE/iSCSI License for System x (FoD) (an upgrade for 88Y7429, which provides FCoE and iSCSI support)	1

- a. The x3750 M4 Dual port 10 GB-T Ethernet Adapter and x3750 M4 Dual port 10 GB SFP+ Ethernet Adapter occupy the same dedicated slot and are mutually exclusive.

Figure 1-41 shows the optional Dual port 10 GB SFP+ Ethernet Adapter that is installed in the dedicated slot (machine type 8722).



Figure 1-41 10 GbE adapter in the dedicated connector (machine type 8722)

The following features are included with the integrated controller:

- ▶ Emulex BE3 chip
- ▶ Two 1 Gb Ethernet ports and two 10 Gb Ethernet ports with optional adapter (1 Gb and 10 Gb auto-negotiation)
- ▶ Full-duplex (FDX) capability
- ▶ Optional 10 Gb ports that operate in a virtual NIC (vNIC™) or physical NIC (pNIC) mode:
 - vNIC mode: Up to six vNICs (up to three vNICs per one 10 Gb port):
 - Virtual Fabric mode or Switch Independent operational mode
 - Virtual port bandwidth allocation in 100 Mbps increments
 - Up to two vNICs can be configured as an IP-based Small Computer System Interface (iSCSI) or Fibre Channel over Ethernet (FCoE) vNICs (one per port) with optional Advanced Upgrade (90Y9310)
 - pNIC mode: Dual-port 1/10 Gb Ethernet adapter
- ▶ IEEE 802.1Q VLAN tagging
- ▶ VLAN insertion and extraction
- ▶ Jumbo frames up to 9000 bytes
- ▶ Load balancing and failover teaming support, including adapter fault tolerance (AFT), switch fault tolerance (SFT), adaptive load balancing (ALB), and IEEE 802.3ad

- ▶ Supports Serial over LAN (SoL) and concurrent KVM (cKVM)
- ▶ Wake On LAN support (1 Gb ports only)
- ▶ Message Signal Interrupt (MSI-X) support
- ▶ IPv4/IPv6 offload:
 - IPv4 TCP Chimney Offload
 - TCP, UDP checksum offload
 - Large send offload (LSO)
 - Large receive offload (LRO)
 - Receive side scaling (RSS)
- ▶ Enhanced Ethernet (draft) support:
 - Enhanced Transmission Selection (ETS) (P802.1Qaz)
 - Priority-based Flow Control (PFC) (P802.1Qbb)
 - Data Center Bridging eXchange Protocol (DCBX) (P802.1Qaz)

1.14 Network adapters

As described in 1.11, “PCI Express slots” on page 48, the x3750 M4 has five PCIe slots standard on the system board with another three PCIe slots on a riser card (which is optional on some models).

For a list of network adapters currently available, see the Lenovo Press product guides:

- ▶ Machine type 8753: <https://lenovopress.com/tips1252>
- ▶ Machine type 8752: <http://lenovopress.com/tips1159>
- ▶ Machine type 8722: <http://lenovopress.com/tips0881>

1.15 Power supplies

The server supports up to two redundant power supplies. Standard models include one or two power supplies, which is model dependent. One power option includes up to two hot-swap redundant 1400 W AC power supplies (80 PLUS Platinum certification). The other option includes up to 900 W power supplies, which are also available through CTO or special bid.

Installing a second power supply requires that the processor and memory expansion tray (88Y7365) or the power interposer card (88Y7367) are installed. The power interposer card option enables redundancy power support when the processor and memory expansion tray is not installed. If you do not have the processor and memory expansion tray installed and want to install two power supplies, the power interposer card must be installed. Figure 1-42 on page 54 shows the power interposer card.

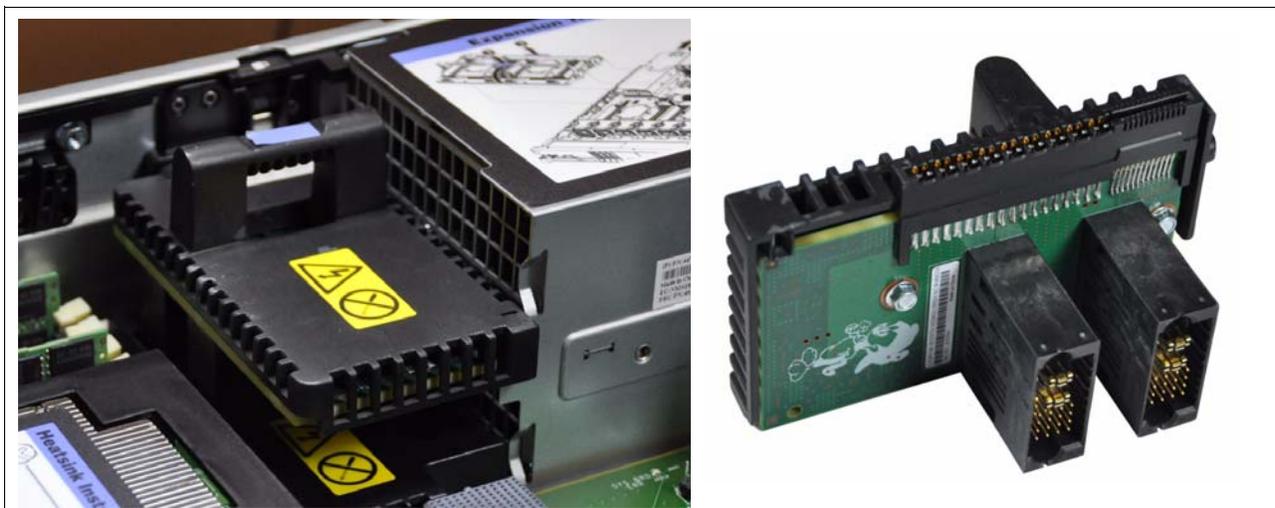


Figure 1-42 Power Interposer for Redundant Power Supply 88Y7367

Table 1-28 lists the power supplies options.

Table 1-28 Power supplies: 8722

Part number	Feature code	Description	Supported on 8752 (E5-4600 v2)	Supported on 8722 (E5-4600)	Maximum supported
44X4150	A54D	1400W HE Redundant Power Supply for altitudes >5000 meters	Yes	No	2
44X4152	A54E	1400W HE Redundant Power Supply	Yes	No	2
88Y7373	A2A6	1400 W HE Redundant Power Supply	No	Yes	2
88Y7431	A2A7	900 W Power Supply	Yes	Yes	2
88Y7433	A2EA	System x 4S- 750 W High Efficiency -48 V DC Power Supply	Yes	Yes	2
88Y7367	A2A0	Power Interposer for Redundant Power Supply	Yes	Yes	1 ^a

a. The power interposer is not needed if the processor and memory expansion tray (88Y7365) are installed.

An AC power supply ships standard with one 2.8 m C13 - C14 power cord.

Two installed power supplies form a redundant pair. Under extreme configurations, it might still be possible to exceed 1400 W DC output with two 1400 W supplies installed. If this condition exists and a power supply fails, the server caps power at 1400 W until the second power supply is back online.

Two 1400 W power supplies are shown in Figure 1-43.

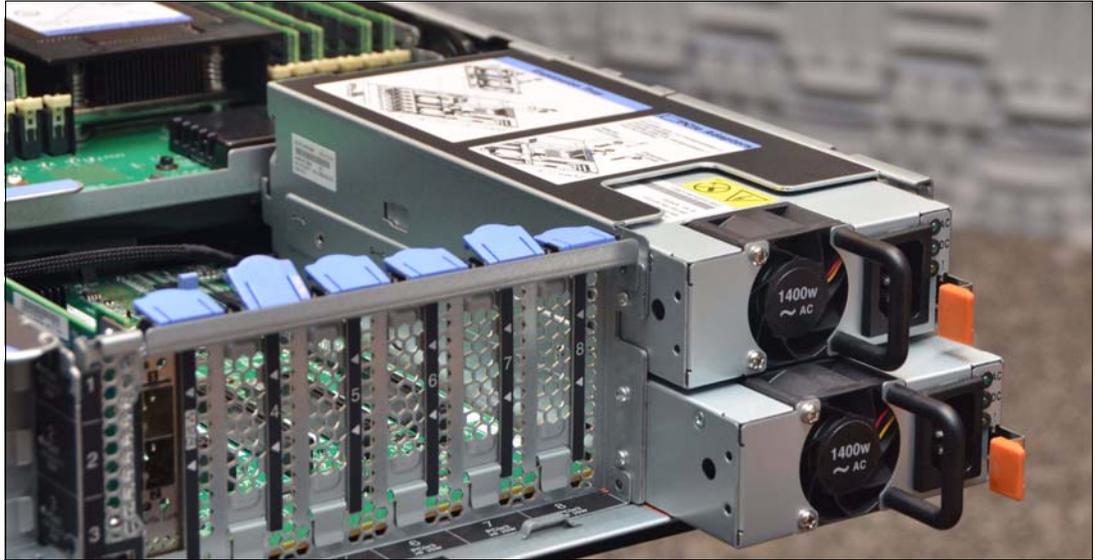


Figure 1-43 Two 1400W power supplies

1.16 Integrated virtualization

The server supports VMware vSphere installed on a USB memory key.

For a list of USB memory keys currently available, see the Lenovo Press product guides:

- ▶ Machine type 8753: <https://lenovopress.com/tips1252>
- ▶ Machine type 8752: <http://lenovopress.com/tips1159>
- ▶ Machine type 8722: <http://lenovopress.com/tips0881>

The server has two USB ports, as shown in Figure 1-44. The USB Hypervisor key can be installed in either socket.

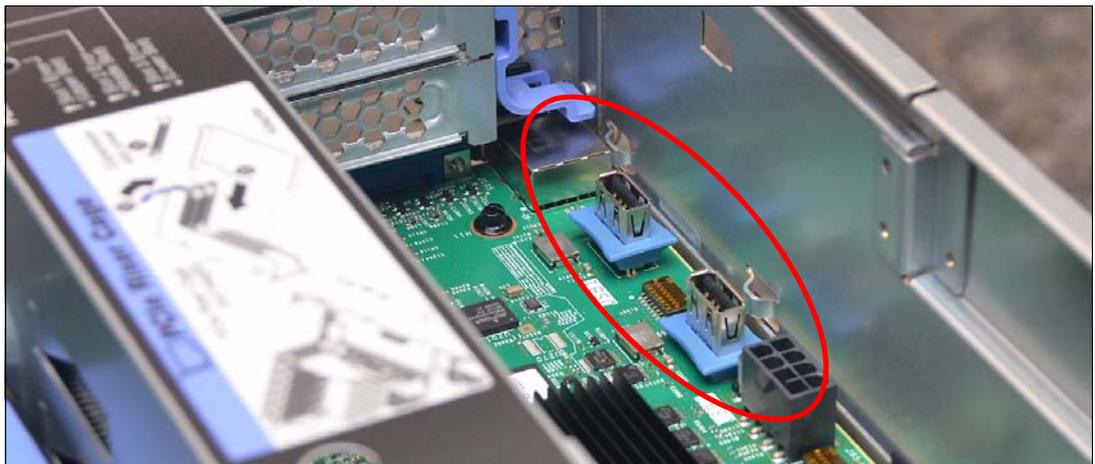


Figure 1-44 Two internal USB ports for VMware vSphere hypervisors

1.17 System management

This section describes some of the systems management features that are available with the x3750 M4.

1.17.1 Light path diagnostic panels

The server features the following three visual indicators that you can use to quickly determine a problem:

- ▶ The Fault LED on the front panel.
- ▶ The Light Path Diagnostics panel, as shown in Figure 1-45.
- ▶ LEDs next to key components on the system board.

The Light Path Diagnostics panel (as shown in Figure 1-45) is accessible by pushing the Eject button on the front of the server.



Figure 1-45 Light Path Diagnostics panel

Each LED in the Light Path Diagnostics panel and its definition is listed in Table 1-29.

Table 1-29 Light Path Diagnostic panel LED definitions

LED	Definition
PS	A power supply failed.
OVER SPEC	The server is using more power than the power supplies can output.
PCI	An error occurred on a PCI bus on the system board or on the PCI riser card. Another LED might be lit on one of the PCI slots on the system board or the riser card.
NMI	A nonmaskable interrupt occurred.

LED	Definition
CONFIG	A hardware configuration error occurred.
LINK	Reserved
CPU	An invalid processor configuration or a processor failed.
MEM	An invalid memory configuration or a memory error occurred.
TEMP	The system temperature exceeded a threshold level.
FAN	A fan failed, is operating too slowly, or was removed.
BOARD	An error occurred on the system board.
HDD	An internal hard disk drive failed or is missing.

Press the round blue Remind button on Light Path Diagnostics panel to place the system-error LED on the operator information panel into Remind mode. In this mode, the system-error LED flashes once every 2 seconds until the problem is corrected, the server is restarted, or a new problem occurs. By placing the system-error LED indicator in Remind mode, you acknowledge that you are aware of the last failure but are not taking immediate action to correct the problem.

The Boot code display that is above the Remind button provides boot error codes that indicate the point at which the system stopped during the boot block and power-on self-test (POST). A boot code is a byte value that is produced by the UEFI. In addition to the UEFI codes, this display provides error codes in the event of a microprocessor error or a power fault. Along with the integrated management module (IMM) event log, the error codes can provide the suggested components to be replaced.

Press the Reset button (which is below the blue Remind button) to reset the server and run the POST. (You might need to use a pen or the end of a straightened paper clip to press the Reset button.)

1.17.2 Integrated Management Module II

The server contains an Integrated Management Module II (IMM2), which is based on the Renesas SH7757 chip. The IMM2 provides advanced service-processor control, monitoring, and an alerting function. If an environmental condition exceeds a threshold or if a system component fails, LEDs on the system board are lit to help you diagnose the problem. The error is then recorded in the event log and you are alerted to the problem.

A virtual presence capability comes standard for remote server management, as described in 1.17.3, “Remote management” on page 58.

The IMM2 provides the following major components as standard features:

- ▶ Intelligent Platform Management Interface (IPMI) v2.0-compliance
- ▶ Remote configuration of the IMM2 and UEFI settings without the need to power on the server
- ▶ Remote access to the system fan, voltage, and temperature values
- ▶ Remote IMM and UEFI update
- ▶ UEFI update when the server is powered off
- ▶ Remote console by way of a serial over LAN

- ▶ Remote access to the system event log
- ▶ Predictive failure analysis and integrated alerting features; for example, via Simple Network Management Protocol (SNMP)
- ▶ Remote presence, including remote control of the server via a Java or Active x client
- ▶ Operating system failure screen (blue screen) capture and display through the web interface
- ▶ Virtual media that allows the attachment of a diskette drive, CD/DVD drive, USB flash drive, or disk image to a server

For more information about the IMM2, see *Integrated Management Module II User's Guide*, which is available at this website:

<http://ibm.com/support/entry/portal/docdisplay?lnocid=MIGR-5086346>

For more information about setting up and using the IMM2, see Chapter 2, "Initial setup of the IMM2" on page 61.

1.17.3 Remote management

Remote server management is provided through the following industry-standard interfaces:

- ▶ IPMI Version 2.0
- ▶ SNMP Version 3
- ▶ Common Information Model (CIM)
- ▶ Web browser

The server also supports the following virtual media and remote control features:

- ▶ Remote viewing of video with graphics resolutions up to 1600 x 1200 at 75 Hz with up to 23 bits per pixel, regardless of the system state.
- ▶ Remote access to the server by using the keyboard and mouse from a remote client.
- ▶ Mapping of the CD or DVD drive, diskette drive, and USB flash drive on a remote client. Also, mapping of the ISO and diskette image files as virtual drives that are available for use by the server.
- ▶ Uploading of a diskette image to the IMM2 memory and mapping it to the server as a virtual drive.
- ▶ Capturing blue-screen errors.

For more information about setting up and using the IMM2, see Chapter 2, "Initial setup of the IMM2" on page 61.

1.18 Physical and electrical specifications

This section describes the physical and electrical specifications that are supported by the System x3750 M4 server.

The server features the following dimensions and weight measurements:

- ▶ Width: 446 mm (17.5 in.)
- ▶ Depth: 734 mm (28.9 in.)
- ▶ Height: 87 mm (3.4 in.)
- ▶ Approximate weight when fully configured: 31.1 kg (68.5 lb.)

The following server environment is supported:

- ▶ Air temperature: 5 - 40 °C (41 - 104 °F) at 0 - 3048 m (10,000 ft); ASHRAE A3 compliant
- ▶ Humidity: 8% - 80%
- ▶ Electrical specifications:
 - 100 - 127 (nominal) V AC, 50 Hz or 60 Hz, and System 20 A (10 A per power supply)
 - 200 - 208 (nominal) V AC, 50 Hz or 60 Hz, and System 10 A
 - 200 - 240 (nominal) V AC, 50 Hz or 60 Hz, and System 9 A
 - Minimum configuration: 0.20 kVA (one power supply)
 - Minimum configuration: 0.26 kVA (two power supplies)
 - Typical configuration: 1.12 kVA (two power supplies)
 - Maximum configuration: 2.16 kVA (two power supplies)
- ▶ BTU output:
 - Minimum configuration: 648 BTU per hour (190 watts)
 - Maximum configuration: 7,336 BTU per hour (2150 watts)
- ▶ Noise level:
 - Idle: 6.5 bels
 - Operating: 6.6 bels

Initial setup of the IMM2

The Integrated Management Module II (IMM2) is the service processor that manages the System x3750 M4. This chapter describes the initial setup tasks to perform after you receive your server.

The following topics are included in this chapter:

- ▶ 2.1, “IMM2 overview” on page 62
- ▶ 2.2, “Setting up IMM2 for remote access” on page 63
- ▶ 2.3, “Initial configuration of IMM2” on page 76

2.1 IMM2 overview

The IMM2 consolidates the service processor functionality, Super I/O, video controller, and remote presence capabilities into a single chip on the server system board.

IMM2 offers three levels of functionality: Basic, Standard, and Advanced. With the System x3750 M4, all levels are enabled by default.

Tip: On many other System x servers, only Standard or Basic functionality is offered by default. Updates to higher levels on those systems require the purchase of a license through Features on Demand (FoD).

The following features are provided with IMM2 Basic:

- ▶ Industry-standard interfaces and protocols
- ▶ Intelligent Platform Management Interface (IPMI) Version 2.0
- ▶ Common Information Model (CIM)
- ▶ Advanced Predictive Failure Analysis (PFA) support
- ▶ Continuous health monitoring
- ▶ Choice of a dedicated or shared Ethernet connection
- ▶ Domain Name System (DNS) server support
- ▶ Dynamic Host Configuration Protocol (DHCP) support
- ▶ Embedded Dynamic System Analysis™ (DSA)
- ▶ LAN over USB for in-band communications to the IMM2
- ▶ Serial over LAN
- ▶ Server console serial redirection

The following features are provided with IMM2 Standard:

- ▶ Access to server vital product data (VPD)
- ▶ Advanced Predictive Failure Analysis (PFA) support
- ▶ Automatic notification and alerts
- ▶ Continuous health monitoring and control
- ▶ Choice of a dedicated or shared Ethernet connection
- ▶ Domain Name System (DNS) server support
- ▶ Dynamic Host Configuration Protocol (DHCP) support
- ▶ Email alerts
- ▶ Syslog logging support
- ▶ Embedded Dynamic System Analysis (DSA)
- ▶ Enhanced user authority levels
- ▶ Event logs that are time stamped, saved in the IMM, and can be attached to email alerts
- ▶ Support for industry-standard interfaces and protocols: IPMI V2.0, CIM, and SNMP
- ▶ OS watchdogs
- ▶ Remote configuration through Advanced Settings Utility™ (ASU)
- ▶ Remote firmware updating
- ▶ Remote power control
- ▶ Secure web server user interface

- ▶ Serial over LAN
- ▶ Server console serial redirection
- ▶ User authentication that uses a secure connection to a Lightweight Directory Access Protocol (LDAP) server

IMM2 Advanced includes all of the features of IMM Standard and the following additions:

- ▶ Remote presence, including remote control of the server through a Java or Active x client
- ▶ Operating system failure screen capture and display through the web interface
- ▶ Virtual media that allows the attachment of a diskette drive, CD/DVD drive, USB flash drive, or disk image to a server

IMM2 features: The System x3750 M4 has all the IMM2 features enabled by default.

The IMM2 offers several improvements over the first-generation IMM:

- ▶ Faster dedicated Ethernet connection
- ▶ Faster boot and power-on self-test (POST) performance
- ▶ Faster power-on sequence
- ▶ Improved remote presence performance
- ▶ New appearance of the web interface
- ▶ Web interface is more responsive

The following major hardware changes are included with the IMM2:

- ▶ Faster and larger memory, upgraded to 256 MB DDR3
- ▶ Faster CPU, upgraded to 576 MHz
- ▶ Higher capacity boot flash and NAND flash
- ▶ Dedicated 1 Gb LAN connection
- ▶ Integrated PCIe 3.0 support
- ▶ Support for USB 2.0

2.2 Setting up IMM2 for remote access

This section describes how to access the IMM2 and how you can configure the IMM2 network interface.

2.2.1 Network interface

You can configure the IMM2 network interface in the Unified Extensible Firmware Interface (UEFI) setup. The UEFI setup is accessed by pressing F1 during the system initializing process, as shown in Figure 2-1 on page 64.

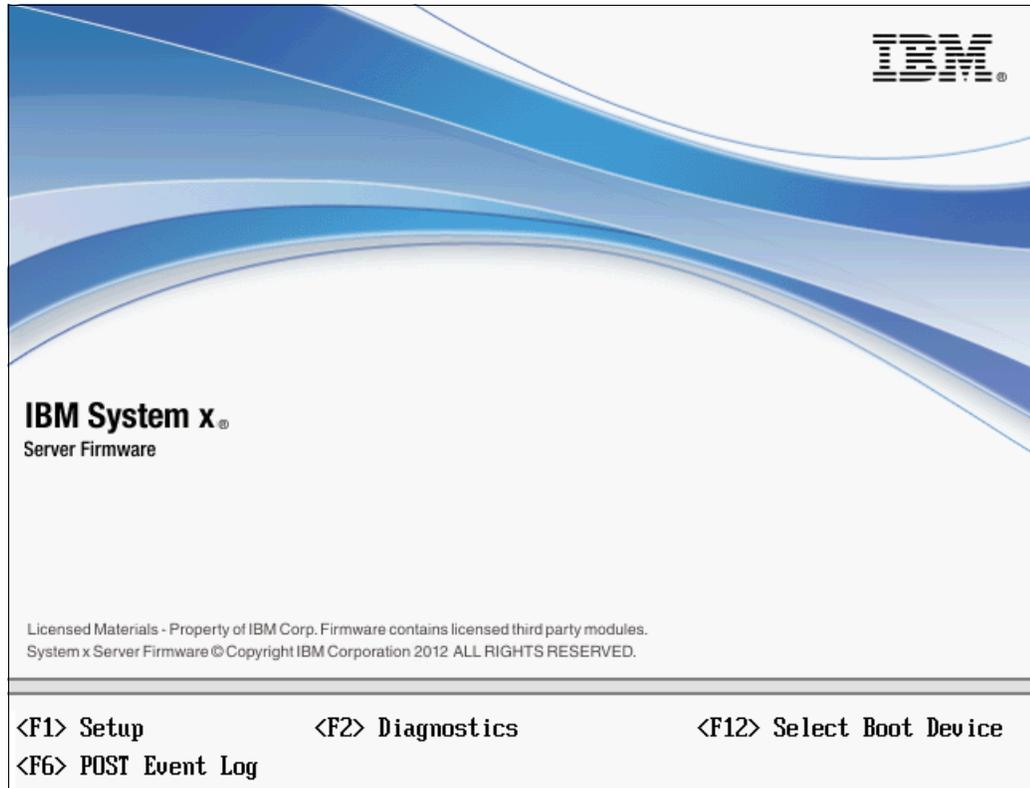


Figure 2-1 UEFI panel on system startup

To configure network settings, select **System Settings** → **Integrated Management Module** → **Network Configuration**. The Network Configuration page, as shown in Figure 2-2 on page 65, is displayed.

The following IMM2 network connection choices are available in the Dynamic Host Configuration Protocol control field:

- ▶ Static IP
- ▶ Dynamic Host Configuration Protocol (DHCP) Enabled
- ▶ DHCP with Failover (default)

The default static IPv4 address for IMM2 is 192.168.70.125. IMM2 does not have a fixed static IPv6 IP address by default.

At Advanced IMM Ethernet Setup, you can set the MTU Size and deactivate the Auto negotiation feature (which is on by default). Remember to save your network settings. You might need to scroll down to view Save Network Settings, which is indicated by the red arrow on the right, as shown in Figure 2-2.

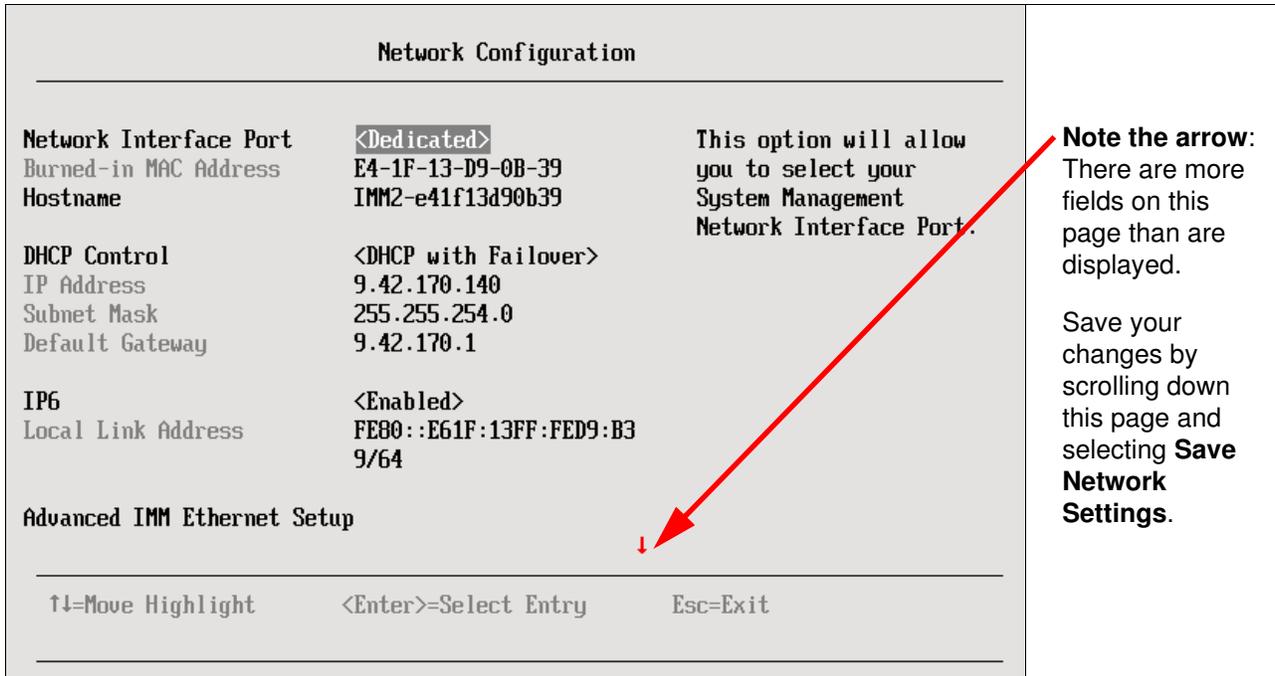


Figure 2-2 IMM2 Network Configuration page in UEFI

2.2.2 Methods to access: Web and command-line interface

This section describes how you can connect to the web interface and command-line interface (CLI) of the IMM2.

Web interface

You can use a web browser to access the IMM2 web interface to obtain all the management functions of the IMM2. The Java plug-in 1.5, or later, and the following web browsers are supported:

- ▶ Internet Explorer 7 or later
- ▶ Firefox 3.5 or later

Start the web browser and enter the IMM2 IP address in the address bar. The default static IPv4 address for IMM2 is 192.168.70.125. The following information provides access to the default user account:

- ▶ User account: USERID (all in uppercase)
- ▶ Password: PASSWORD (with a zero, not the letter O)

The default user has Supervisor access; therefore, change this user during your initial configuration. Figure 2-3 shows the IMM2 Login page.



Figure 2-3 IMM2 Login page

The IMM2 main page is displayed upon a successful login, as shown in Figure 2-4.

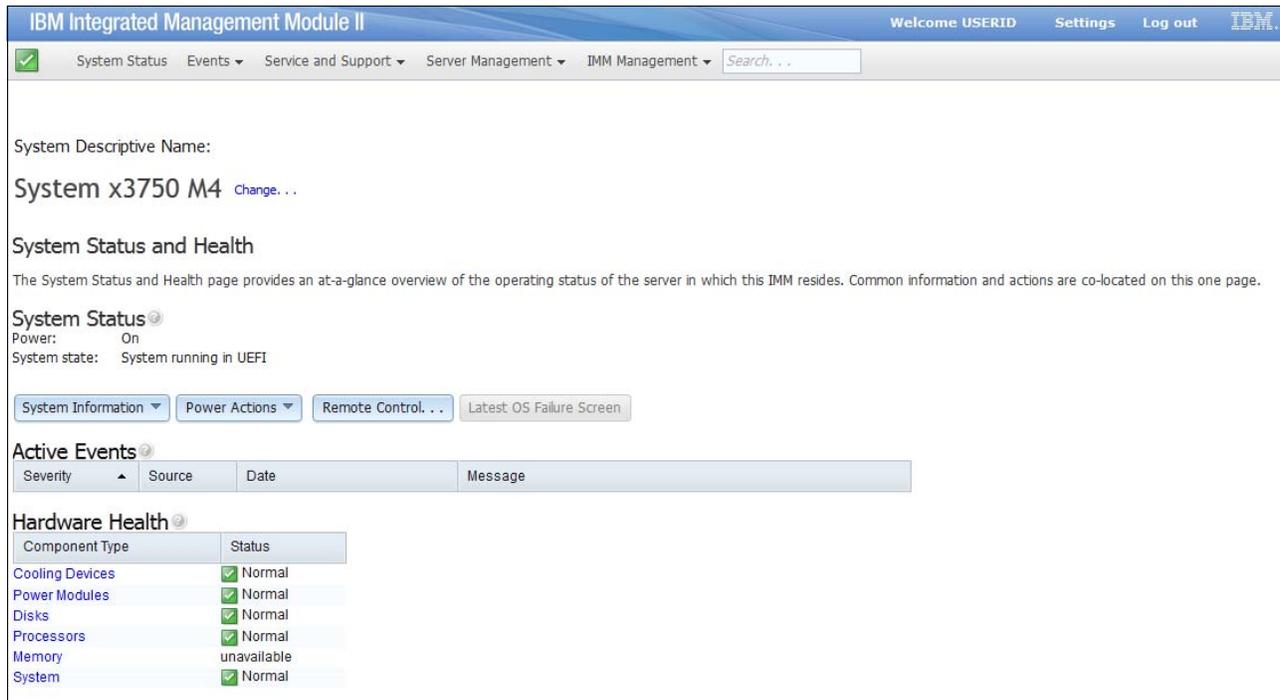


Figure 2-4 IMM2 main page

HTTPS: HTTPS is disabled, by default. You can enable HTTPS by selecting **IMM Management** → **Security**.

For more information, see the *User's Guide for Integrated Management Module II - BladeCenter and System x*, which is available at this website:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5086346>

Command-line interface

You can use the command-line interface (CLI) to access the IMM2 without having to use the web interface. The CLI provides a subset of management functions that are provided by the web interface.

Start a Telnet or Secure Shell (SSH) session with your preferred application to the IMM2 IP address to access the CLI. The default user name is USERID and password is PASSWORD (with a zero, not the letter O). The default user has Supervisor access; therefore, change this user during your initial configuration.

After a successful login, the following window is displayed, as shown in Figure 2-5.

```
login as: USERID
Using keyboard-interactive authentication.
Password:
MYIMM> █
```

Figure 2-5 IMM2 CLI main page

Multiple sessions: Multiple concurrent CLI sessions are allowed with different access methods (Telnet or SSH). At most, two Telnet command-line sessions can be active at any time.

For more information, see the *User's Guide for Integrated Management Module II - BladeCenter and System x*, at this website:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5086346>

2.2.3 IMM2 web interface quick overview

The web interface of the IMM2 has a new appearance compared to the original IMM. The web interface gives you the ability to perform remote management of the server and to configure and control the IMM2.

System Status page

If the login was successful, the IMM2 main page is displayed, as shown in Figure 2-6.

The screenshot shows the IBM Integrated Management Module II (IMM2) main page. The page title is "IBM Integrated Management Module II" and it includes a navigation bar with "Welcome USERID", "Settings", and "Log out". The main content area displays the following information:

- System Descriptive Name:** System x3750 M4 (Change...)
- System Status and Health:** The System Status and Health page provides an at-a-glance overview of the operating status of the server in which this IMM resides. Common information and actions are co-located on this one page.
- System Status:** Power: On; System state: System running in UEFI.
- Active Events:** A table with columns for Severity, Source, Date, and Message.
- Hardware Health:** A table with columns for Component Type and Status.

Component Type	Status
Cooling Devices	Normal
Power Modules	Normal
Disks	Normal
Processors	Normal
Memory	unavailable
System	Normal

Figure 2-6 IMM2 main page

The IMM2 main page or System Status page provides an overview of the operating status of the server, common information, and actions. At a glance, the following important details are shown on the IMM2 main page:

- ▶ **System Status**
Advises if the system is powered off, running, or booting into the operating system.
- ▶ **System Information**
Provides a quick view of common system information, such as the machine type, serial number, and Server State. Figure 2-7 on page 69 shows the System Information Quick View.

System Information Quick View	
Name	Value
Machine Name	System x3750 M4
Machine Type	8722
Model	AC1
Serial Number	23D6526
UUID	A4D93888624911E1B3E3E41F13DA13BA
Server Power	On
Server State	Booting OS or in unsupported OS
Total hours powered-on	481
Restart count	72
Ambient Temperature	71.60 F / 22.00 C
Enclosure Identify LED	Off Change...
Check Log LED	Off

Close

Figure 2-7 System Information Quick View

- ▶ Active events

Events that are still active but are not recovered or returned to a normal state.

- ▶ Hardware health

Status of the hardware components in the system. You can click each component to view details.

Menu bar

On every web page in the IMM2 interface, a menu bar is displayed, as shown in Figure 2-8. The menu bar provides tabs with easy access to most of the features and functions of the IMM2.

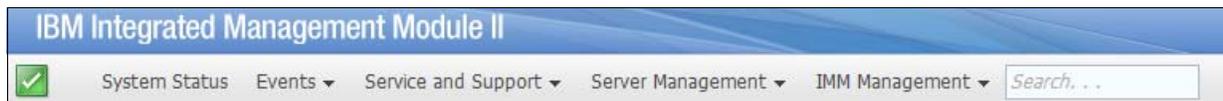


Figure 2-8 Menu bar

Figure 2-9 shows the functions that are available in the Events tab.

Events	Service and Support	Server Management	IMM Management
Event Log	Full log history of all events		
Event Recipients	Add and modify E-Mail and SysLog notifications		

Figure 2-9 Events tab

Figure 2-10 shows the functions that are available in the Service and Support tab.

Service and Support ▾	Server Management ▾	IMM Management ▾	Search
Download Service Data		Obtain a compressed file of relevant service data	

Figure 2-10 Service and Support tab

Figure 2-11 shows the functions that are available in the Server Management tab.

Server Management ▾	IMM Management ▾	Search . . .
Server Firmware	View firmware levels and update firmware	
Remote Control	Allows you access into the operating system of your system	
Server Properties	Various properties and settings related to your system	
Server Power Actions	Power actions such as power on, power off, and restart	
Cooling Devices	Cooling devices installed in your system	
Power Modules	Power modules installed in your system	
Disks	Hard disk drives installed directly in your system	
Memory	RAM installed in your system	
Processors	Physical CPUs installed in your system	
Server Timeouts	Configure watchdogs, etc.	
PXE Network Boot	Settings for how your system performs boot from PXE server	
Latest OS Failure Screen	Windows systems only. View an image of the most recent failure screen.	

Figure 2-11 Server Management tab

Figure 2-12 shows the functions that are available in the IMM Management tab.

IMM Management <input type="text" value="Search..."/>	
IMM Properties	Various properties and settings related to the IMM
Users	Create and modify user accounts that will have access to the IMM console
Network	Network settings such as SNMP and LDAP used by the IMM
Security	Configure security protocols such as SSL and SSH
IMM Configuration	View a summary of the current configuration settings.
Restart IMM	Restart the IMM. Typically only needed when experiencing problems with the IMM
Reset IMM to factory defaults...	Sets all current configuration settings back to default values
Activation Key Management	Add and remove activation keys for additional functionality

Figure 2-12 IMM Management tab

There is also a search field to quickly find IMM2 functions.

Event Logs and critical events

You can see the status of your system on the left side of the menu bar. If everything is OK, you see a green check mark. Otherwise, additional information is included in the title bar, as shown in Figure 2-13.

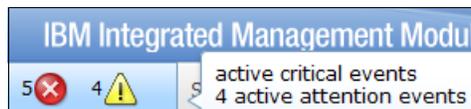


Figure 2-13 Active alerts information

Event Log

To see the Event Log, go to **Events** → **Event Log** as shown in Figure 2-14.

Events ▾ Service and Support ▾ Server Management ▾ IMM Managemen	
Event Log	Full log history of all events
Event Recipients	Add and modify E-Mail and SysLog notifications

Figure 2-14 Event Log tab

This page displays the contents of the IMM event log. All events in the log are time-stamped, by using the IMM2 date and time settings. You can export the events in the event log to a text file. By default, the most recent log entry is the first.

You can filter the log entries on multiple fields; however, you can sort the log entries on only one field at a time. The IMM2 log has a limited capacity and when it is full, the oldest entry is deleted.

Figure 2-15 displays the contents of the IMM event log.

Search Events. . . <input type="button" value="Go"/>			
Severity	Source	Date	Message
 Informational	System	24 May 2012 03:45:36.162 PM	Remote Login Successful. Login ID: USERID from webgui at IP address 9.42.170.197.
 Informational	System	24 May 2012 03:05:36.320 PM	Power cap throttling occurred.
 Warning	Processors	24 May 2012 03:05:14.987 PM	processor 4 is operating in a Degraded State.
 Warning	Processors	24 May 2012 03:05:14.921 PM	processor 3 is operating in a Degraded State.
 Warning	Processors	24 May 2012 03:05:14.779 PM	processor 2 is operating in a Degraded State.
 Warning	Processors	24 May 2012 03:05:14.727 PM	processor 1 is operating in a Degraded State.
 Informational	System	24 May 2012 03:02:46.076 PM	LED Identify state changed to Blinking by Identify Button.
 Informational	Cooling	24 May 2012 03:02:30.392 PM	Non-redundant. Insufficient Resources for cooling zone 131 has deasserted.
 Informational	Cooling	24 May 2012 03:02:28.362 PM	Numeric sensor Fan 3B Tach going low (lower critical) has deasserted.
 Informational	Cooling	24 May 2012 03:02:26.310 PM	Numeric sensor Fan 3A Tach going low (lower critical) has deasserted.
 Informational	Power	24 May 2012 03:02:18.146 PM	Host Power has been turned on.
 Informational	System	24 May 2012 03:02:09.453 PM	LED Fault state changed to Blinking by Remind Button.
 Informational	Power	24 May 2012 02:59:28.815 PM	Redundancy Lost for power unit / power domain 129 has deasserted.
 Informational	System	24 May 2012 02:59:23.750 PM	Power Supply 2 has returned to a Normal Input State.
 Informational	Power	24 May 2012 02:44:02.211 PM	Host Power has been turned off.
 Critical	Cooling	24 May 2012 02:43:56.075 PM	Numeric sensor Fan 3B Tach going low (lower critical) has asserted.
 Critical	Cooling	24 May 2012 02:43:54.101 PM	Non-redundant. Insufficient Resources for cooling zone 131 has asserted.
 Critical	Cooling	24 May 2012 02:43:54.033 PM	Numeric sensor Fan 3A Tach going low (lower critical) has asserted.

Figure 2-15 Event Log

Save Service Data

To create the Service Data for Support, click **Service and Support** → **Download Service Data**. Then click **Download Now**, as shown in Figure 2-16.

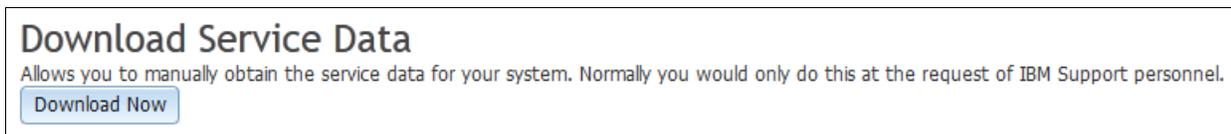


Figure 2-16 Download Service Data

It takes a few minutes to create the Service Data. A progress bar is displayed, as shown in Figure 2-17.

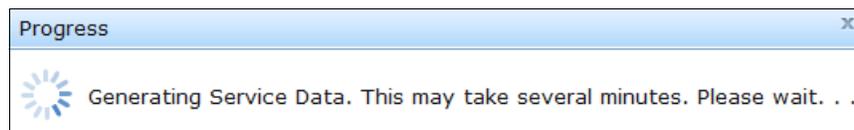


Figure 2-17 Progress bar

Save the `ffdc.tgz` file and send it to your Support personnel.

2.2.4 Remote presence

The IMM2 provides high speed remote control of the system during all stages of server operation, including the POST and boot process. IMM2 also supports remote media connectivity.

Remote control

The remote control feature takes over the keyboard and mouse functions of the remote server. This function also enables you to view what is displayed on the panel from any network connected location. Any operating system that is supported on the server is supported by the remote control feature.

Table 2-1 shows the supported video modes.

Table 2-1 Supported video modes

Resolution	Refresh rate
640 x 480	60, 72, 75, 85 Hz
800 x 600	60, 72, 75, 85 Hz
1024 x 768	60, 72, 75, 85 Hz
1440 x 900	60 Hz
1280 x 1024	60, 75 Hz
1680 x 1050	60
1600 x 1200	60, 75 Hz

Encryption: You have the option of encrypting all remote control traffic. Select **Server Management** → **Remote Control** and select the check box **Encrypt disk and KVM data during transmission**.

Both Windows and Linux clients support single-user or multiuser modes. The IMM2 has both Java based and ActiveX-based remote clients and the ability to mount the disk image, the disk, or the selected remote drive:

► Java based

The Sun Java Runtime Environment 1.5 or later is required for the Java based client. The Java client displays the remote video on the client and sends the keyboard and mouse information to the IMM2.

You also have the virtual media feature to mount a select drive, disk, or disk image to the server.

► ActiveX-based

The ActiveX client runs only in Internet Explorer 7 or later. The ActiveX client displays the remote video on the client and sends the keyboard and mouse information to the IMM2. You also have the virtual media feature to mount a select drive, disk, or disk image to the server.

You might need to install an add-on when you run the ActiveX client for the first time. Figure 2-18 shows the ActiveX add-on.

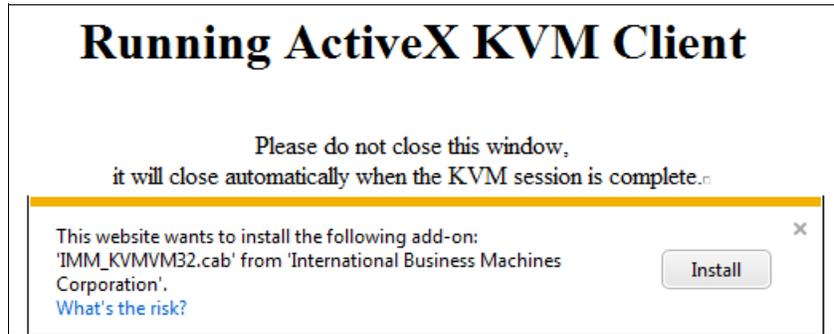


Figure 2-18 ActiveX add-on

Virtual media

The virtual media feature can mount a local drive, optical disk, or ISO image to the remote server. This feature requires USB support. Therefore, only the following operating systems are supported on the server:

- ▶ Microsoft Windows Server 2008
- ▶ Microsoft Windows Server 2003
- ▶ Red Hat Enterprise Linux 4.0 and 5.0
- ▶ SUSE Linux Enterprise Server 10.0
- ▶ Novell NetWare 6.5

To start the Virtual Media feature, select **Tools** → **Launch Virtual Media**, as shown in Figure 2-19.

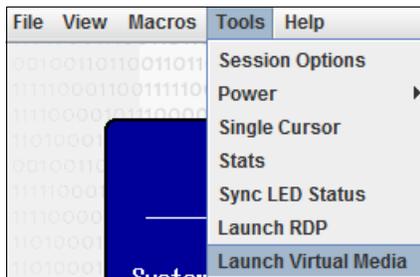


Figure 2-19 Launch Virtual Media feature

Follow these steps to mount a drive through IMM2:

1. Connect to the IMM2 with your web browser.
2. Click **Server Management** → **Remote Control**.
3. Choose between a Java based client or an ActiveX-based client.
4. If you want to allow other users to have remote control access during your session, click **Start remote control in multi-user mode**. Otherwise, click **Start remote control in single-user mode**.

- In the Video Viewer window, click **Tools** → **Launch Virtual Media**. Figure 2-20 shows the Virtual Media window.

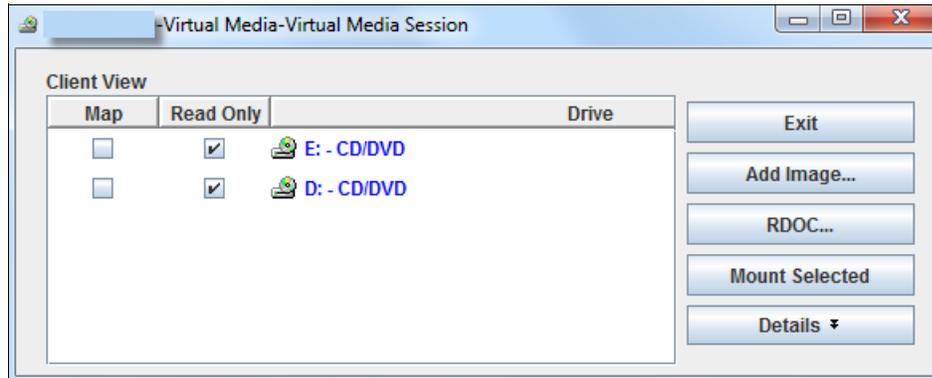


Figure 2-20 The Virtual Media Session window

- Select the **Virtual Media Session** window.
- Click **Add Image** if you want to map an IMG or ISO image file.
- Select the check box next to the drive that you want to map and click **Mount Selected**, as shown in Figure 2-21.

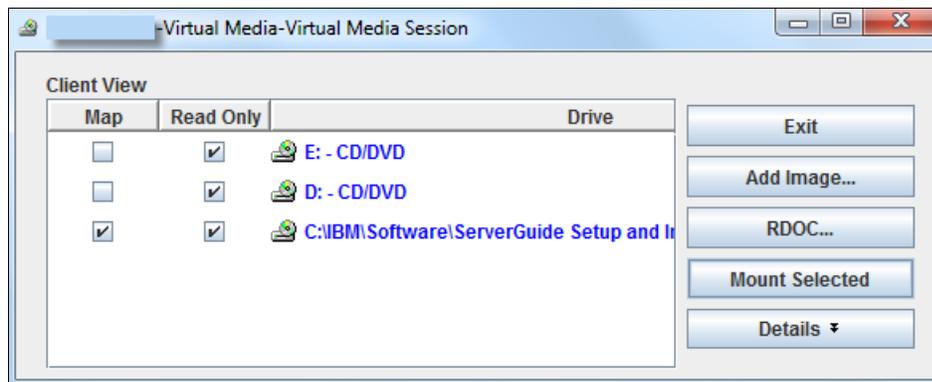


Figure 2-21 Overview of the selected drive

- The image drive is now accessible by the system.

Closing the session: If you close the Virtual Media Session window when a remote disk is mapped to the system, it causes the machine to lose access to the remote disk.

Tip: The Virtual Media function requires your computer to be running Microsoft Windows, Red Hat Linux, or SUSE Linux. Computers that run Linux must be configured to have the operating system emulate a Small Computer System Interface (SCSI) CD-ROM for an integrated development environment (IDE) CD-ROM.

2.3 Initial configuration of IMM2

For the initial configuration of the IMM2, there is a good starting point to use the Initial Setup Wizard. You can find the wizard in **IMM Management** → **IMM Configuration** as shown in Figure 2-22.



Figure 2-22 Initial Setup Wizard

The Initial Setup Wizard provides the following settings that can be changed or added:

- ▶ System descriptive name
- ▶ Contact person
- ▶ Location of the system
- ▶ Host name
- ▶ IP Configuration
- ▶ Authentication Method
- ▶ User Account changes and additions
- ▶ Event Notification

Figure 2-23 shows the Welcome page of the Initial Setup Wizard.

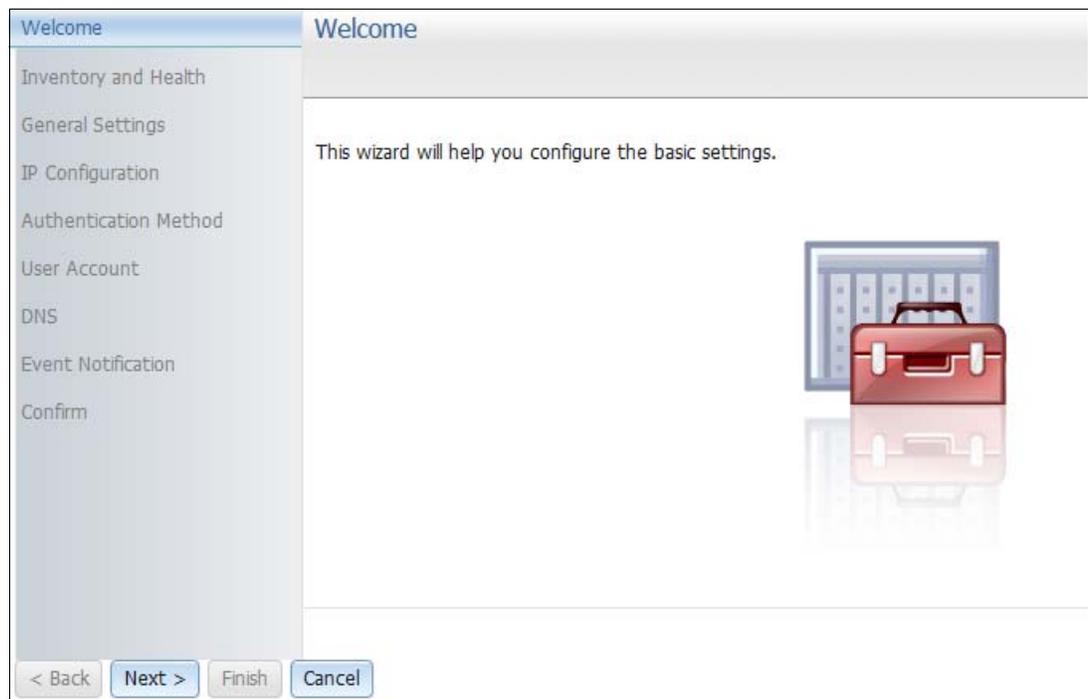


Figure 2-23 Welcome page

Confirm your choices by clicking **Finish**. The IMM2 then saves all settings and performs a reboot. During this reboot, you lose the connection to the IMM2.

Initial setup of UEFI and hardware components

This chapter provides an overview of the options available for configuration and system optimization when setting up an x3750 M4 for the first time.

This chapter describes the process to change the settings in the Unified Extensible Firmware Interface (UEFI) firmware and the different ways to apply these changes. We also describe how the System Configuration and Boot Management menu structure can be used to configure the additional hardware, such as the RAID controllers and network interfaces.

We also describe where to find and how to select the configurations when you choose to implement a system for performance or power savings.

In this chapter, the following topics are described:

- ▶ 3.1, “System Settings in UEFI” on page 78
- ▶ 3.2, “Configuration to enable memory mirroring or sparing” on page 90
- ▶ 3.3, “Configuring the RAID adapter by using UEFI” on page 93
- ▶ 3.4, “Configuring embedded Ethernet using UEFI” on page 106

3.1 System Settings in UEFI

To boot into the System Configuration and Boot Management menu in the UEFI, press **F1** for Setup at the splash window before the power-on self-test (POST). From the System Configuration and Boot Management menu, you can view System Information. Information you can view includes the components and system level Vital Product Data (VPD) along with the base system hardware firmware versions that are installed. You can configure and make settings changes to all onboard system devices and optional hardware additions. Figure 3-1 shows the System Configuration and Boot Management window.

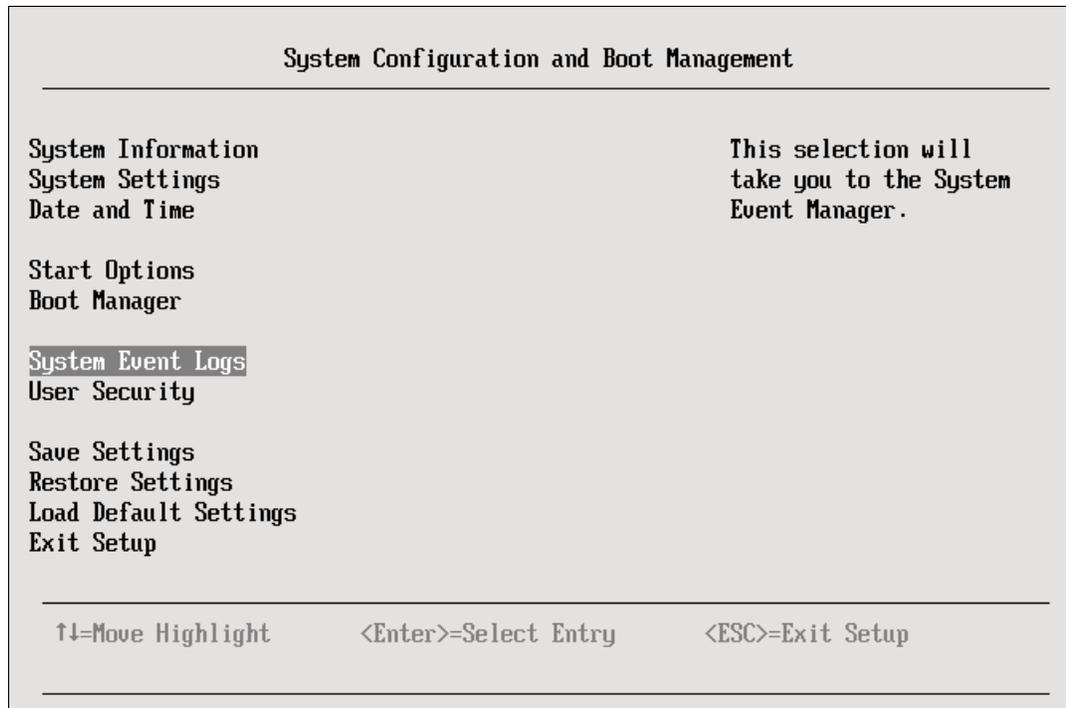


Figure 3-1 System Configuration and Boot Management window

The main menu allows for setting the system date and time. By using the Start Options and Boot Manager menus, many standard and advanced options are available to target booting from a device or setting a standard boot order.

3.1.1 Boot Manager menu

The Boot Manager also allows for configuring to boot using the UEFI option or Legacy BIOS boot. The Legacy BIOS boot configuration is needed for certain OSs or external bootable options, such as a legacy boot from storage area network (SAN).

The Boot Manager menu spans two screens as indicated by the small red arrow on each window. See Figure 3-2 and Figure 3-3.

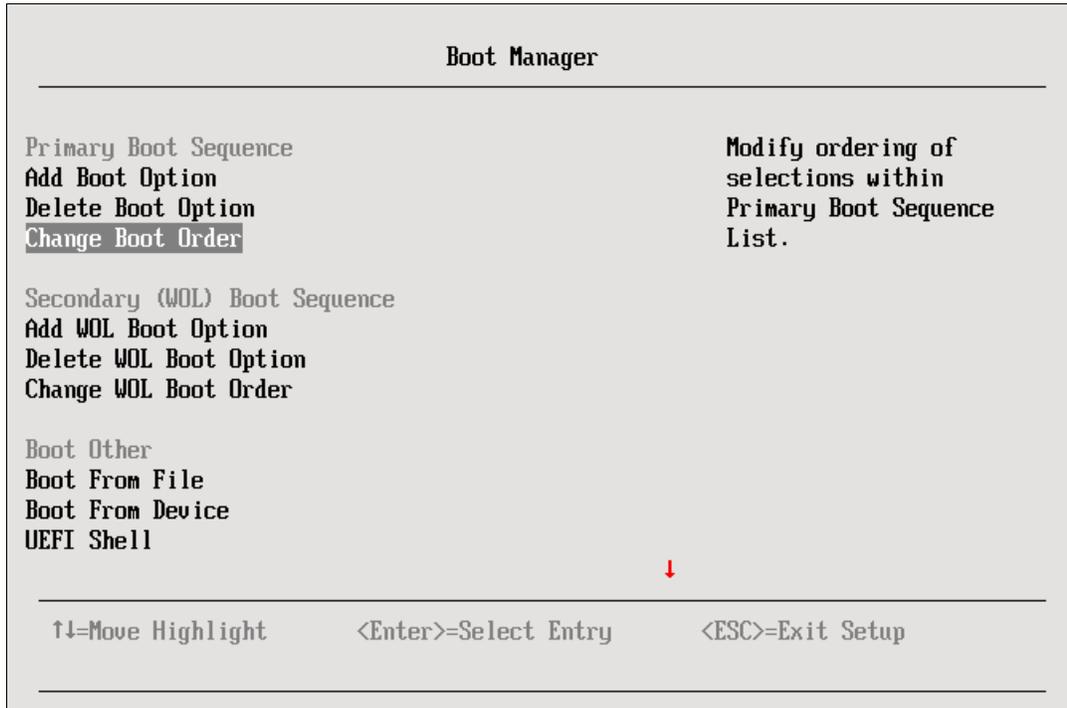


Figure 3-2 Boot Manager (panel 1 of 2)

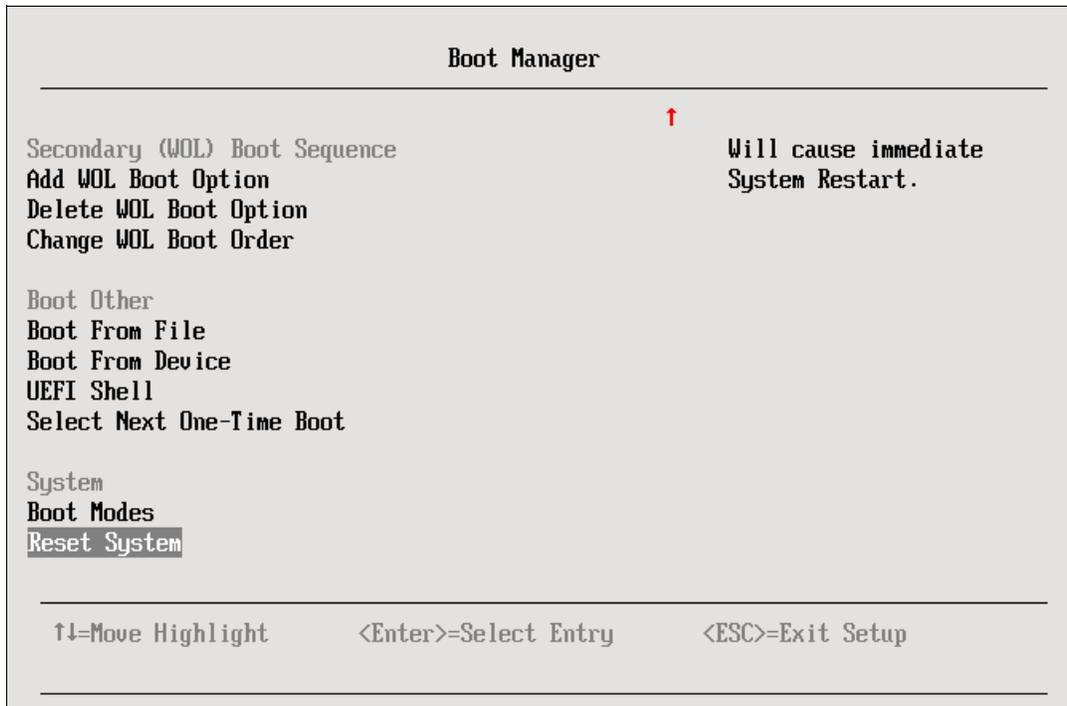


Figure 3-3 Boot Manager (panel 2 of 2)

The selections in the Boot Manager menu include the option to set a primary boot order by adding options, deleting options, and changing the boot order. This configuration can be done for the primary boot order and the Wake On LAN (WOL) boot order. The remaining options in the Boot Manager include the following selections:

- ▶ **Boot From File:** Browse to a bootable file or partition to add to the boot order.
- ▶ **Boot From Device:** Select a system-recognized device in which to boot directly.
- ▶ **UEFI Shell:** Boot into a full UEFI shell to run commands when needed or directed by Support.
- ▶ **Select Next One-Time Boot:** Select the next boot device without adding it to the boot order.
- ▶ **Boot Modes:** A submenu in Boot Manager that provides additional options for how the system boots:
 - **System Boot Mode** gives three options to choose from: UEFI and Legacy, UEFI Only, and Legacy Only. The System Boot Mode controls which Option ROM the boot manager uses. The UEFI Only option only allows UEFI-aware adapters and boot loaders. This option must have a UEFI-aware OS to boot. The Legacy Only option skips UEFI-aware drivers and runs option ROMs to boot a legacy operating system.
 - **Optimized Boot** can be enabled or disabled. This option is enabled by default and limits the number of drivers loaded in POST.
 - **Quiet Boot** is enabled by default. If this function is disabled, a system summary is shown before the system logo.
 - **Driver Health Check** applies only to boots through the POST. When in the F1 function, the Driver Health Check applies to all drivers. The two options available for the Driver Health Check are: Select all drivers or just the drivers for onboard devices.
- ▶ **Reset System:** Perform an immediate system reset.

3.1.2 System Event Log

The System Event Log provides a viewer for POST Events, the base system logs, and the option to clear the system event log when needed. Figure 3-4 shows the System Event Log viewer.

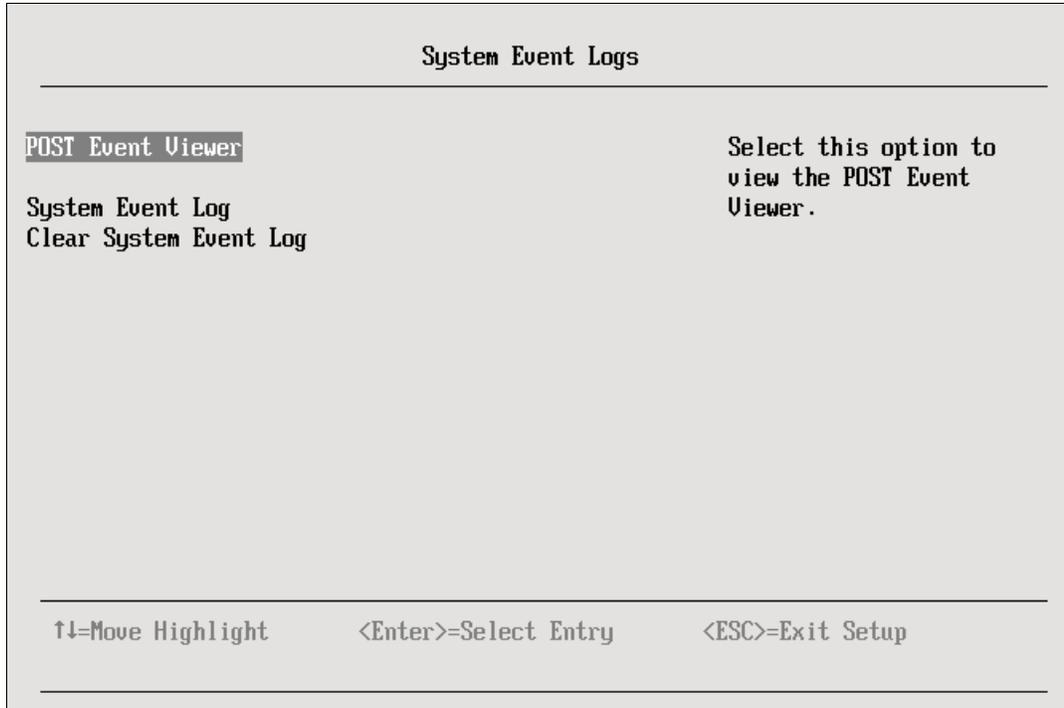


Figure 3-4 System Event Log viewer

3.1.3 User Security

The User Security options can also be modified from the System Configuration menu, which adds a layer of security to the UEFI firmware. The extra security is available because there is an option to set power-on and admin passwords. The admin password sets a privileged access password. Figure 3-5 on page 82 shows the User Security window.

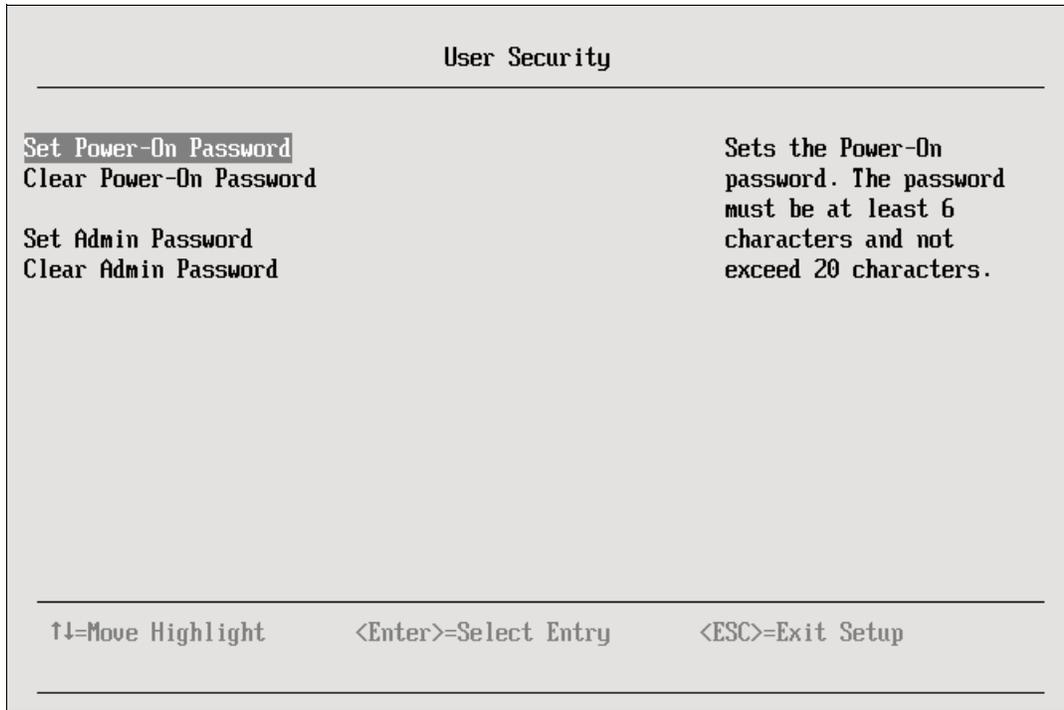


Figure 3-5 User Security window

Before exiting the System Configuration and Boot Management menu, be sure to use the selections to either Save Settings, Restore Settings to those last saved, or Load Default Settings. Loading the default settings from the F1 menu is useful when the settings changes might appear to cause unwanted system behavior.

Default settings: Restoring the default settings reverts all the custom boot settings and system settings changes made in the UEFI, back to their initial values. However, it does not disrupt settings made to the RAID controller, the IMM, or the network adapters from the System Settings.

3.1.4 Default UEFI settings and Advanced Settings Utility values

The Advanced Settings Utility (ASU) tool is used to change or reset the UEFI settings and system information. The change options that are available in this tool include: machine type and model, serial number, and asset tag number.

The ASU tool allows for the modification of the UEFI settings from the command line of multiple OS platforms. It provides these options for modification without the need to reboot the machine to access the System Configuration menu by using F1. The tool can be used to set complementary metal oxide semiconductor (CMOS) values and Vital Product Data (VPD). The tool also interacts with the IMM2 and allows for a remote mount of media through a network accessible IMM2. The ASU tool also offers management for Feature on Demand (FoD) key activation and many more functions.

The ASU tool can also be used on a larger scale for multiple deployments.

For the latest version of the ASU tool for your supported OS, see this website:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=T00L-ASU>

For more information about the use of the ASU tool and command-line options, see the *Advanced Settings Utility, v9.21 User's Guide*:

<http://ibm.com/support/entry/portal/docdisplay?lnocid=MIGR-5085890>

3.1.5 Operating modes

When the system settings in the UEFI need to be modified, many options are available for the optimization of performance or for considering efficiency with either performance or power consumption.

By selecting **System Settings** → **Operating Modes**, you can choose pre-set operating modes. Figure 3-6 shows the Operating Modes window.

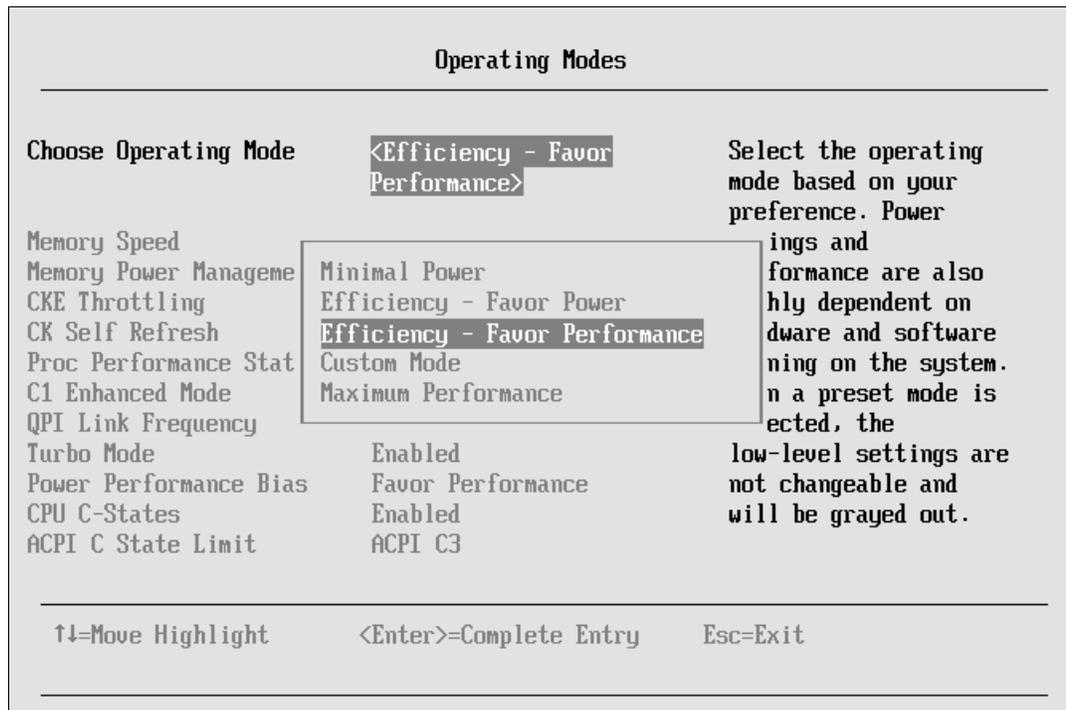


Figure 3-6 Operating Modes window

These pre-set options do not allow for the modification of the remaining low-level settings in Operating Modes, but allow for a base level of operation. Power consumption and performance still depend heavily on the hardware and software options used with the x3750 M4.

The default operating mode is *Efficiency - Favor Performance*.

The following operating mode choices are available in the Operating Modes window:

- ▶ Minimal Power

This preset operating mode allows for only a base system operation. This mode provides absolute minimal power usage. It cannot configure the memory or the processors for performance beyond a base function. The options that are shown in gray text in Figure 3-7 on page 84 show the settings for the processors and memory. These same settings are shown in gray text in the System Settings section on processors (as shown in Figure 3-10 on page 87) and in memory, as shown in Figure 3-11 on page 88.

Figure 3-7 shows the Minimal Power operating mode window.

Operating Modes		
Choose Operating Mode	<Minimal Power>	Select the operating mode based on your preference. Power savings and performance are also highly dependent on hardware and software running on the system. When a preset mode is selected, the low-level settings are not changeable and will be grayed out.
Memory Speed	Minimal Power	
Memory Power Management	Automatic	
CKE Throttling	N/A	
CK Self Refresh	N/A	
Proc Performance States	Enabled	
C1 Enhanced Mode	Enabled	
QPI Link Frequency	Minimal Power	
Turbo Mode	Disabled	
Power Performance Bias	Minimal Power	
CPU C-States	Enabled	
ACPI C State Limit	ACPI C3	
<hr/> ↑↓=Move Highlight <Enter>=Select Entry Esc=Exit		

Figure 3-7 Minimal Power Operating Mode window

► Efficiency - Favor Power

This preset option does allow for a more balanced operation. However, it favors power in the function of the memory and the processors. Figure 3-8 shows the Efficiency - Favor Power operating mode window.

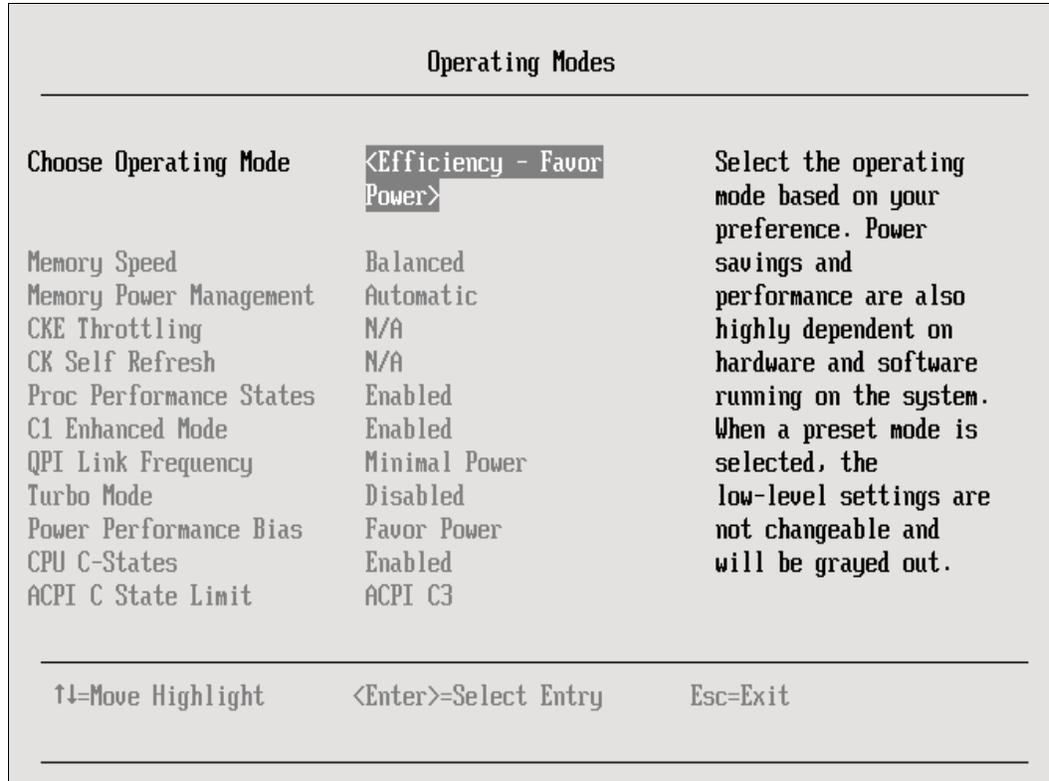


Figure 3-8 Efficiency - Favor Power window

► Efficiency - Favor Performance (default)

This preset operating mode allows for greater performance when balancing for power and performance. Favor performance is the default operating mode. Figure 3-9 shows the Efficiency - Favor Performance window.

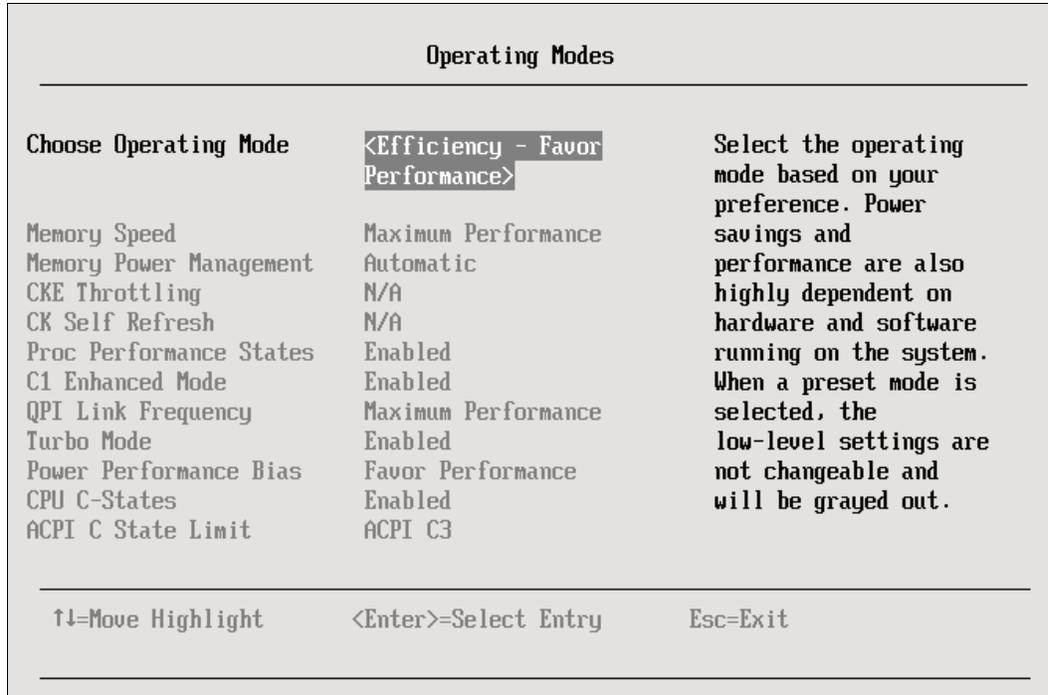


Figure 3-9 Efficiency - Favor Performance window

► **Custom mode**

Choosing this operating mode allows for all previously inaccessible gray fields on the Operating Mode page, along with the Memory and Processor pages under System Settings, to be customized. If the Custom Mode is chosen, all the settings, by default, favor performance until modified.

► **Maximum performance**

This operating mode sets all the subsystems to maximum performance, with no power-saving options enabled. The maximum performance setting does not allow for the modification of Active Energy Manager Power Capping, which is found by selecting **System Configuration and Boot Management** → **System Settings** → **Power**.

3.1.6 Settings for performance of individual subsystems

When considering the performance settings, many options are available for the modification of the individual subsystems.

The System Settings menu has subcategories that allow for configuration of the operation of processors, memory, and power.

Processors

The Processors section of the System Settings in the UEFI allows you to change the function of the installed processors. The changes to these settings can affect performance and power savings. The options in gray text can be modified only by changing the system operating mode. To change the mode, see 3.1.5, “Operating modes” on page 83.

Figure 3-10 shows the processor settings in the UEFI.

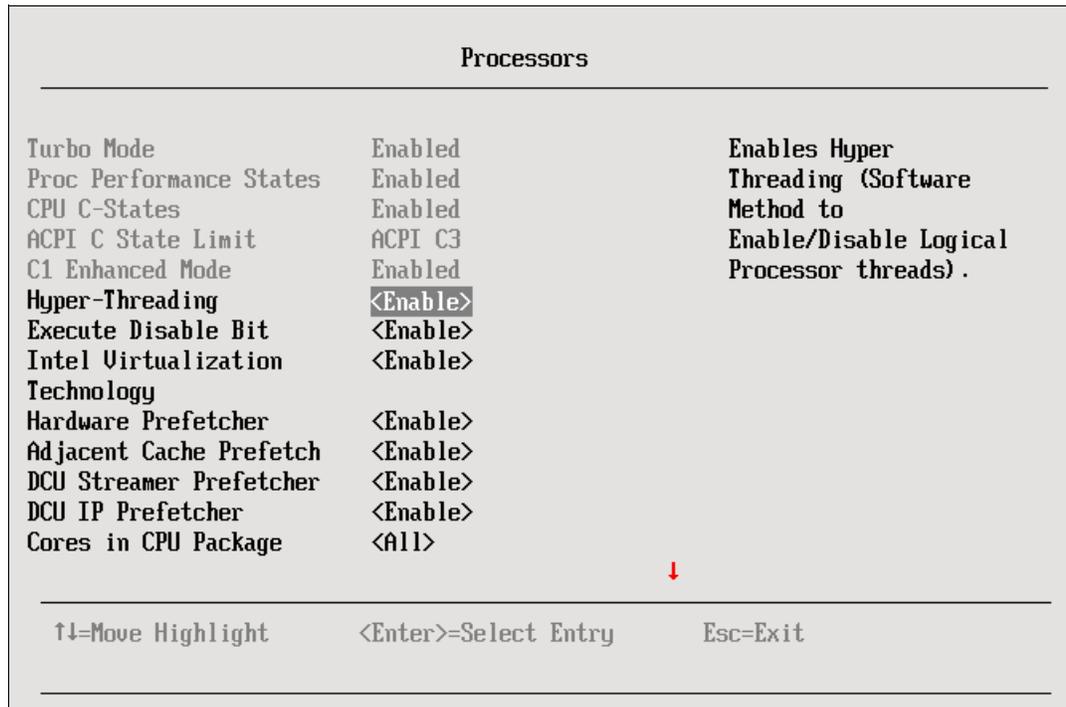


Figure 3-10 System Settings for processors

Memory

The Memory settings page in the UEFI, as shown in Figure 3-11 on page 88, serves many purposes. The page allows you to view the installation order of the DIMMs, per processor, that is installed in the x3750 M4 without having to remove the cover of the machine.

The options in gray text can be modified only by changing the system operating mode. To change the mode, see 3.1.5, “Operating modes” on page 83.

The System Memory Details option shows this installation order. If a DIMM fails, which causes the DIMM slot to become disabled, the DIMM details page allows for the identification and enablement of the disabled DIMM slot on the processor.

Figure 3-11 shows the System Memory Details.

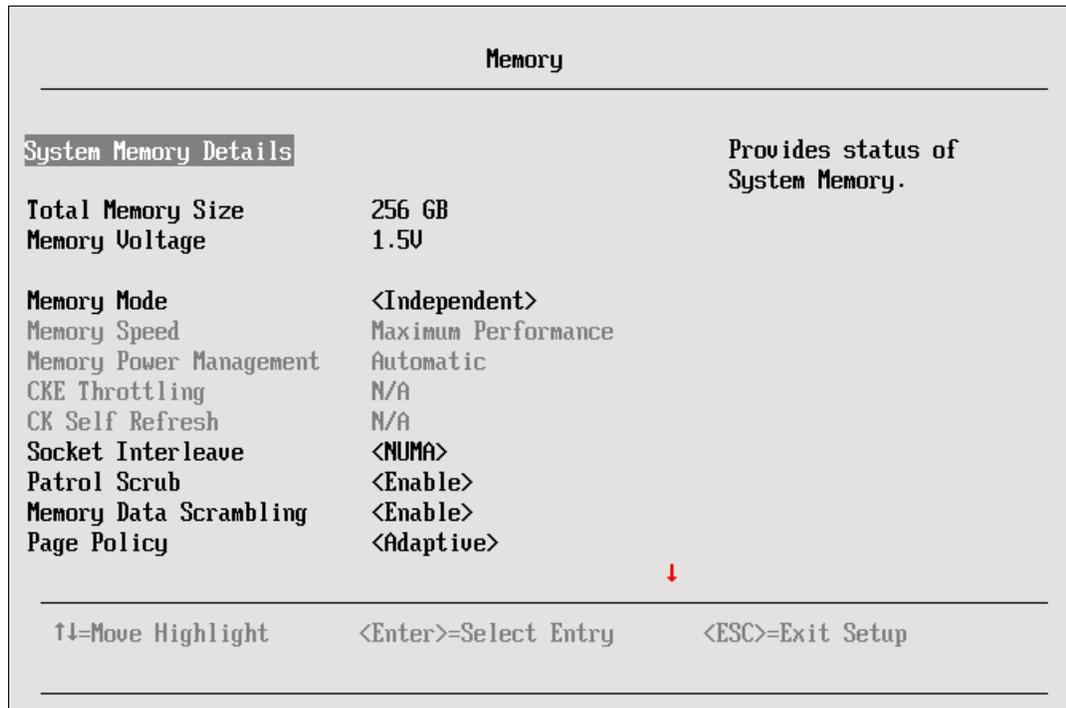


Figure 3-11 System Settings for Memory

Determine what is the best use of the memory in your machine, based on the environment, production needs, or applications that are used. The memory must then be installed in the correct order to fit the configuration (see the installation guide).

To see the current DIMM installation order, select **System Configuration and Boot Management** → **System Settings** → **Memory** → **System Memory Details**. From there, choose which memory to view by the processor for which it is installed. DIMMs are displayed as Present, Enabled, or [EMPTY] if no memory failures are present.

Power

When you are considering power savings over performance, more settings changes can be made from the Power menu in System Settings. To get to the Power settings menu, boot into F1, then select **System Settings** → **Power**. The Power settings window, shown in Figure 3-12 on page 89 is displayed.

The following options are available in the Power settings menu:

- ▶ By default, the Active Energy Manager has Capping Enabled. This default allows Systems Director Active Energy Manager to cap the power consumption of the system if the power policies require it.
- ▶ The Power Performance Bias reflects what the Operating Mode was set to in the System Settings for Operating Modes.
- ▶ For Workload Configuration, the default setting is Balanced, but must be changed to *I/O Sensitive* if I/O expansion cards require high bandwidth when the CPU cores are idle.

Figure 3-12 shows the Power menu in System Settings.

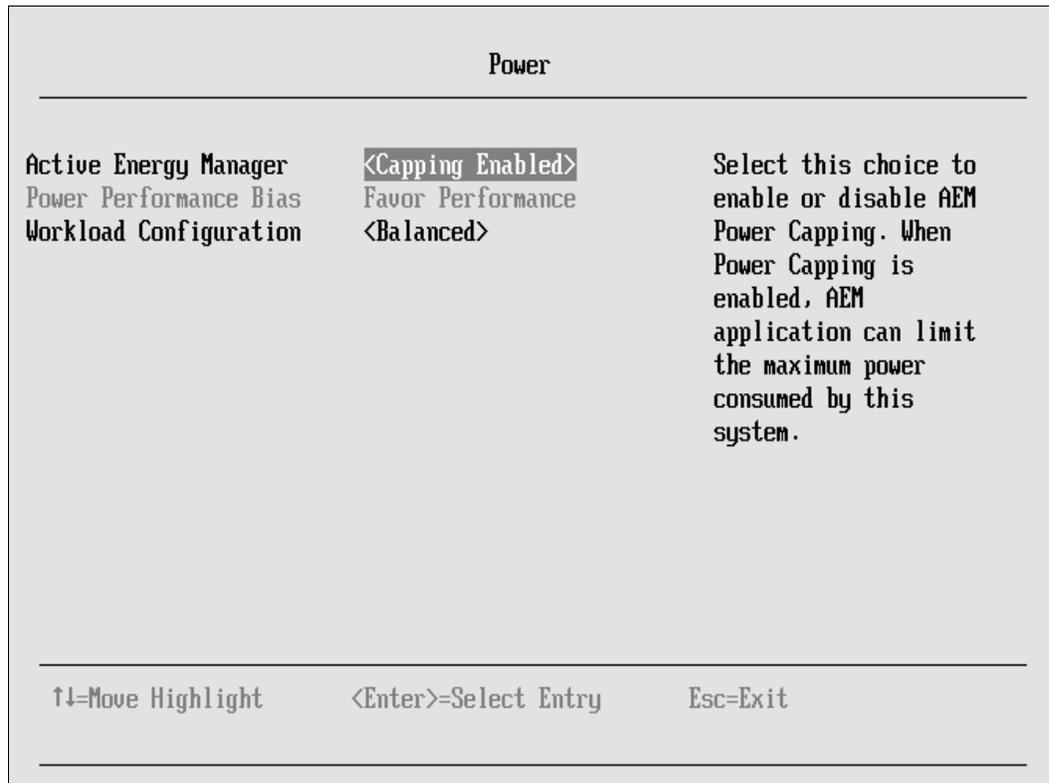


Figure 3-12 Power settings

3.1.7 Optimizing boot time

How to optimize boot time is a common question with UEFI-based machines; this question can be addressed many ways. As with any server, small changes in the configuration of the devices attached to the server externally, and PCI options installed in the server, can increase boot time. One way that boot time is addressed in the x3750 M4 is through a default setting in the UEFI. This setting decreases boot time by limiting the number of drivers that are loaded early in the POST. The setting also decreases boot time by allowing those devices to still be properly used in the OS when the OS driver starts.

This setting can be modified by selecting **Boot Manager** → **Boot Modes** in the UEFI to ensure that Optimized Boot is set to Enabled. Figure 3-13 shows the Boot Mode options.

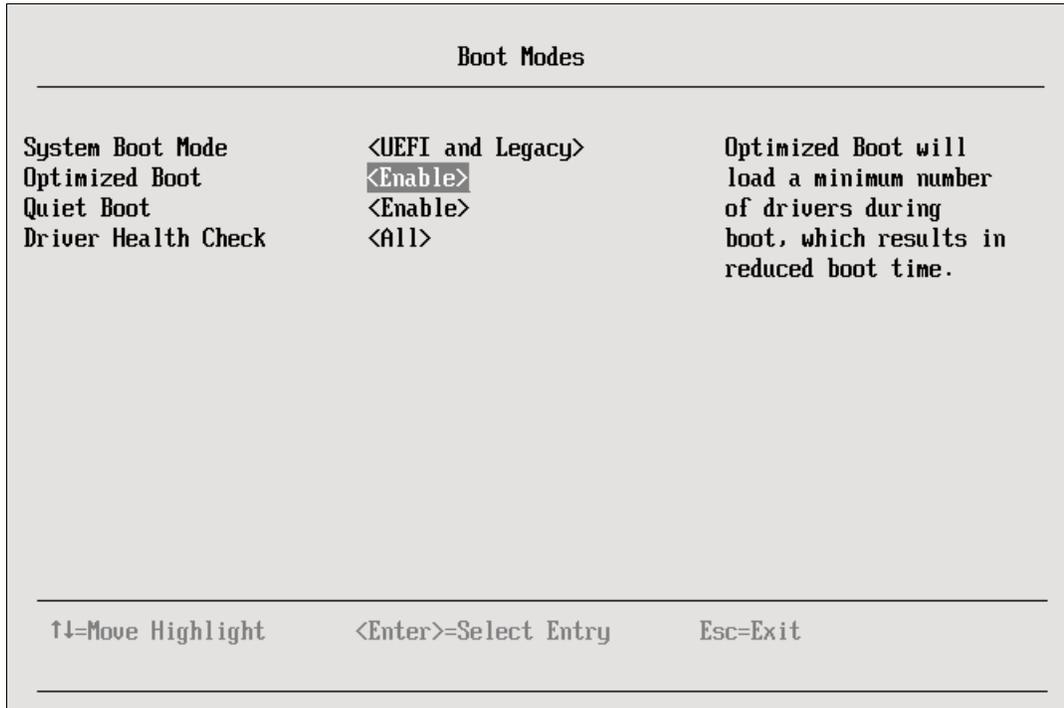


Figure 3-13 Boot Mode options

By having the option ROM enabled for a PCI device, whether the server is booting in the Legacy or UEFI mode, requires the UEFI to load a BIOS for the device in the POST. If this PCI adapter has devices attached to it, additional time is spent discovering these devices before the server starts to boot to the OS.

The following actions can be taken to further improve boot time:

- ▶ Boot into **F1** → **System Settings** → **Devices and I/O Ports** → **Enable / Disable Adapter Option ROM Support**. Then, disable the option ROM for slots or devices that the server does not need in which to boot. The adapter continues to function after the OS driver is started.
- ▶ Boot into **F1** → **Boot Manager** and use the options to either **Change Boot Order** or **Delete Boot Option** to remove or lower the priority of a bootable device.

Tip: Do not disable the option ROMs for any intended bootable devices. This action prevents them from being able to boot.

3.2 Configuration to enable memory mirroring or sparing

The x3750 M4 has three options for how to use the memory that is installed to meet the production needs of the server and the applications installed. The options are: Independent, Memory Mirroring, and Memory Sparing. The Independent mode is the normal operation mode where all memory is available to the OS.

In this section, we describe the use of memory mirroring and memory sparing.

3.2.1 Memory mirroring

Memory mirroring allows for the replication of data across two channels of DIMMs. This function allows for a primary and backup set of DIMMs, with duplicate data across both. If the system fails, the memory controller is able to immediately switch to the backup memory channel without any downtime.

To allow for memory mirroring, the DIMMs cannot be installed one at a time, as in Independent Mode, but as pairs across channels. Data in DIMMs on channel 1 are mirrored to channel 2, and data in DIMMs installed on memory channel 3 are mirrored with channel 4. This configuration is the case with each processor installed.

With memory mirroring, the data is duplicated across pairs on opposite channels to provide a primary and backup source. The usable capacity of the installed DIMMs is half of what is installed.

Before enabling memory mirroring in the UEFI, install the DIMMs in the sequence shown in Table 3-1. The 3750 M4 does not support a three-processor configuration; therefore, the table displays DIMM installations for 1, 2, and 4 processor configurations. See Table 3-1.

Table 3-1 DIMM installation sequence for memory mirroring

DIMMs	Number of processors	DIMM connector population sequence on the system board	DIMM connector population sequence on the optional microprocessor and memory expansion tray
One processor			
1st pair	1	1, 4	None
2nd pair	1	9, 12	None
3rd pair	1	2, 5	None
4th pair	1	8, 11	None
5th pair	1	3, 6	None
6th pair	1	7, 10	None
Two processors			
1st set	2	1, 4, 13, 16	None
2nd set	2	9, 12, 21, 24	None
3rd set	2	2, 5, 14, 17	None
4th set	2	8, 11, 20, 23	None
5th set	2	3, 6, 15, 18	None
6th set	2	7, 10, 19, 22	None

DIMMs	Number of processors	DIMM connector population sequence on the system board	DIMM connector population sequence on the optional microprocessor and memory expansion tray
Four processors			
1st set	4	1, 4, 13, 16	25, 28, 37, 40
2nd set	4	9, 12, 21, 24	33, 36, 45, 48
3rd set	4	2, 5, 14, 17	26, 29, 38, 41
4th set	4	8, 11, 20, 23	32, 35, 44, 47
5th set	4	3, 6, 15, 18	27, 30, 39, 42
6th set	4	7, 10, 19, 22	31, 34, 43, 46

When installation is complete, select **F1** → **System Settings** → **Memory** and change the Memory Mode to **Mirroring**.

3.2.2 Memory sparing

Memory sparing or *memory rank sparing* provides another level of backup to the existing memory. This option is not as redundant as memory mirroring, but less capacity is lost because of redundancy.

In memory sparing, one rank in each memory channel is reserved, reducing the total ranks per channel by one. When a preset limit of recoverable errors on a DIMM rank is reached, the data is copied to the spare rank to continue operation. Because of this requirement, DIMM sparing requires more than one DIMM per channel when using single or dual rank DIMMs. Spare ranks cannot be shared across memory channels or processors. Only the quad-rank DIMMs can handle sparing per DIMM and can then be installed one at a time, if needed.

Consult Table 3-2 for the correct installation order for DIMM sparing.

Table 3-2 DIMM population sequence for the Memory Sparing Mode

DIMMs	Number of processors	DIMM connector population sequence on the system board	DIMM connector population sequence on the optional microprocessor and memory expansion tray
One processor			
1st set	1	1, 2, 3	None
2nd set	1	10, 11, 12	None
3rd set	1	4, 5, 6	None
4th set	1	7, 8, 9	None

DIMMs	Number of processors	DIMM connector population sequence on the system board	DIMM connector population sequence on the optional microprocessor and memory expansion tray
Two processors			
1st set	2	1, 2, 3, 13, 14, 15	None
2nd set	2	10, 11, 12, 22, 23, 24	None
3rd set	2	4, 5, 6, 16, 17, 18	None
4th set	2	7, 8, 9, 19, 20, 21	None
Four processors			
1st set	4	1, 2, 3, 13, 14, 15	25, 26, 27, 37, 38, 39
2nd set	4	10, 11, 12, 22, 23, 24	34, 35, 36, 46, 47, 48
3rd set	4	4, 5, 6, 16, 17, 18	28, 29, 30, 40, 41, 42
4th set	4	7, 8, 9, 19, 20, 21	31, 32, 33, 43, 44, 45

When the memory is installed for Memory Sparing, it must be enabled in the UEFI. Select **F1** → **System Settings** → **Memory** and set Memory Mode to Sparing.

3.3 Configuring the RAID adapter by using UEFI

A new feature to the ServeRAID M5100 series SAS controllers is the ability to configure and monitor arrays and attached devices without leaving the UEFI interface. The traditional methods of using the WebBIOS of the MegaRAID controller are still accessible, either by Ctrl-H in a Legacy POST or through **F1** → **System Settings** → **Adapters and UEFI Driver**. However, the way the BIOS of the storage adapter is integrated into the System Configuration and Boot Management menu, allows for easier movement between the hardware components that need to be configured. This configuration also applies to any additional supported ServeRAID controllers added to the server.

This section describes the many options and configurable settings that use the UEFI to manage and configure the storage adapter. For more information, see the *ServeRAID-M Software User's Guide* at this website:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5086126>

The embedded SAS controller supports RAID 0 and RAID 1. Additional RAID levels and other features, such as the cache and backup options, are available either as hardware upgrades or FoD license upgrades. See 1.9.4, "RAID controller upgrades" on page 42 for a list of upgrades.

To reach the storage adapter to configure or manage it, reboot the server and boot into F1 at the startup window. Select **System Settings** → **Storage** then select the adapter to configure.

The menu options available for the storage controller can be chosen from the Configuration Options menu, as shown in Figure 3-14.

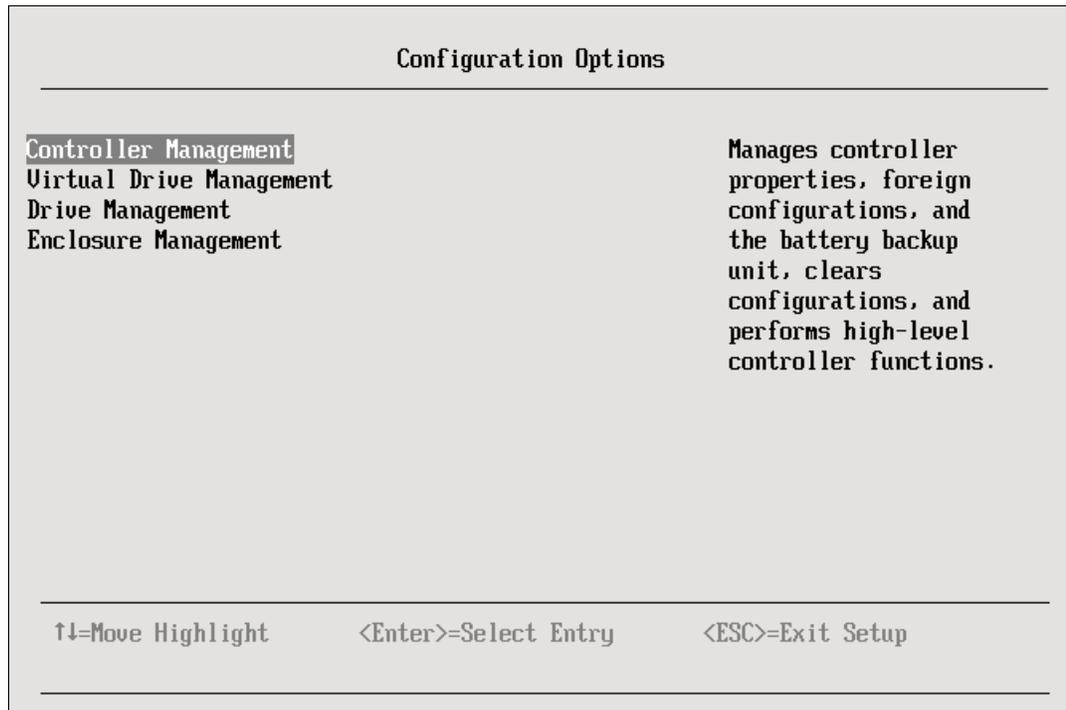


Figure 3-14 Storage Configuration Options

The following menu items are available in the Configuration Options window:

- ▶ Controller Management manages the controller properties, foreign configurations, and the battery backup unit. This tool is used to clear configurations and perform high-level controller functions. For more information, see 3.3.1, “Controller Management” on page 95.
- ▶ Virtual Driver Management creates virtual drives, manages virtual drive properties, and performs operations. This tool can perform operations, such as background initialization, and consistency checks, and can delete, locate, and reconfigure virtual drives. For more information, see 3.3.2, “Virtual Drive Management” on page 100.
- ▶ Drive Management displays drive properties and performs operations, such as assigning and unassigning a hot spare, locating drives, forcing drives offline and online, and rebuilding drivers after a drive failure. For more information, see 3.3.3, “Drive Management” on page 105.
- ▶ Enclosure management displays attached enclosures, enclosure firmware revisions, and drives that are attached to enclosures. For more information, see 3.3.4, “Enclosure Management” on page 106.

3.3.1 Controller Management

Figure 3-15 shows the Controller Management menu.

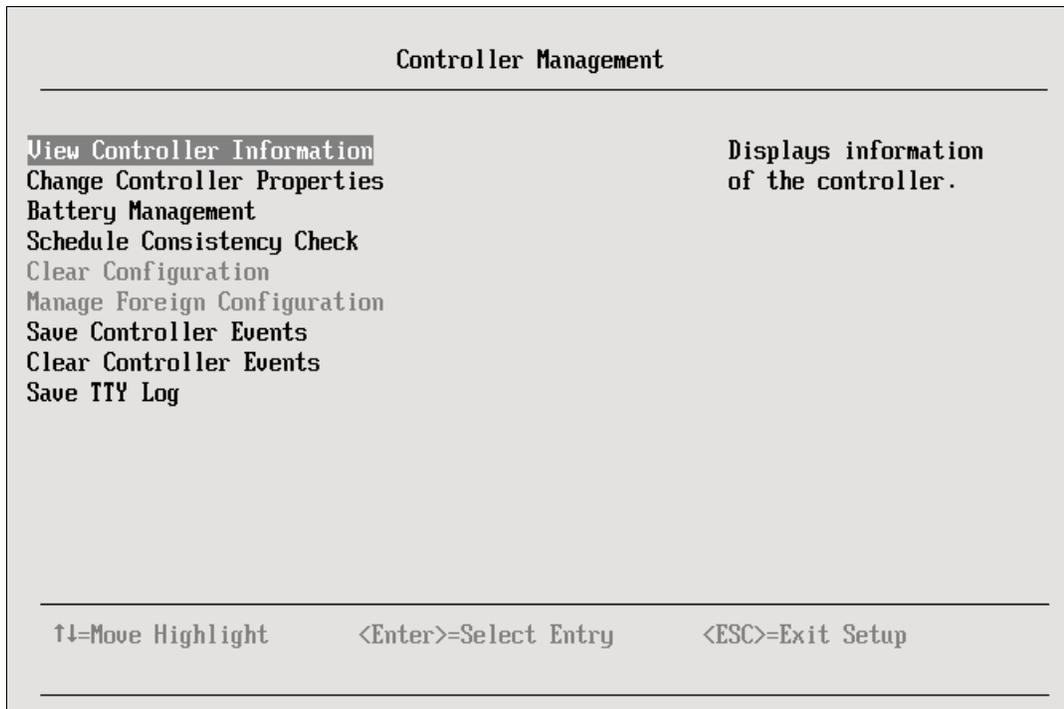


Figure 3-15 Controller Management menu

The following subsections describe the Controller Management menu selections:

- ▶ “View Controller Information” on page 95
- ▶ “Change Controller Properties” on page 97
- ▶ “Battery Management” on page 99
- ▶ “Schedule Consistency Check” on page 99
- ▶ “Clear Configuration” on page 100
- ▶ “Manage Foreign Configuration” on page 100
- ▶ “Event and log options” on page 100

View Controller Information

The following information is displayed in the View Controller Information selection:

- ▶ Figure 3-16 on page 96 shows the Controller Properties.

CONTROLLER PROPERTIES:	
Controller Marketing Name	ServeRAID M5110e
Serial Number	1CH05E
PCI ID	0x100010x005B10x101410x0413
Host Interface	PCI-E
Host Port Count	[2]
PCI Slot Number	Integrated
Drive Count	[4]
Virtual Drive Count	[0]
Encryption Capable	<No>
Protection Capable	<No>
ROC Temperature (C)	[46]
Shield State Supported	<Yes>

Figure 3-16 Controller Properties

- ▶ Figure 3-17 shows the Firmware Properties.

FIRMWARE PROPERTIES:	
Firmware Package Version	23.2.1-0028
Firmware Build Time	Jan 14 2012:15:06:22

Figure 3-17 Firmware Properties

- ▶ Figure 3-18 shows the Memory Properties.

MEMORY PROPERTIES:	
Memory Size (MB)	[512]
NUSRAM Size (MB)	[32]
Metadata Size (MB)	[512]
Minimum Strip Size	<8 KB>
Maximum Strip Size	<1 MB>
CacheCade - SSD Caching	<Disabled>
Write Cache Capable	<No>
Configured Cache Size (GB)	[0]
Maximum Allowed Cache Size (GB)	[0]

Figure 3-18 Memory Properties

- ▶ Figure 3-19 shows the Patrol Read and Emergency Hot Spare Properties.

PATROL READ PROPERTIES:	
Patrol Read State	<Stopped>
Patrol Read Iterations	[0]
EMERGENCY HOT SPARE PROPERTIES:	
Global Hot Spare for Emergency	<Disabled>
Unconfigured Good for Emergency	<Disabled>
Emergency for SMARTer	<Disabled>

Figure 3-19 Patrol Read and Emergency Hot Spare Properties

Change Controller Properties

Select the **Change Controller Properties** option to make and apply settings changes to the controller. Figure 3-20 shows the full list of properties.

Apply Changes	
Set Factory Defaults	
Set Boot Device	<None>
Drive Replace	<Enabled>
Drive Replace on SMART Error	<Disabled>
Cluster Mode	<Disabled>
Rebuild Rate	[30]
Background Initialization (BGI) Rate	[30]
Consistency Check Rate	[30]
Reconstruction Rate	[30]
Controller BIOS Coercion Mode	<Enabled>
	<1 GB>
SMART Polling	[300]
Stop Boot Process on Error	<Enabled>
Drive Power Save	<Disabled>
Stop Consistency Check on Error	<Disabled>
Maintain Drive Fail History	<Enabled>
Load Balance	<Auto>
Enable Auto Import	<Enabled>
Persistent Hot Spare	<Disabled>
Manage Link Speed	
PATROL READ PROPERTIES:	
Patrol Read Mode	<Auto>
Patrol Read Rate	[30]
Resume Patrol Read	
Stop Patrol Read	
Patrol Read Setting for Unconfigured Space	<Enabled>
CACHE PROPERTIES:	
Cache Flush Interval	[4]
Preserved Cache	<No>
Discard Preserved Cache	
EMERGENCY HOT SPARE PROPERTIES:	
Global Hot Spare for Emergency Unconfigured Good for Emergency	<Disabled>
Emergency for SMARTer	<Disabled>

Figure 3-20 Change Controller Properties

The following configurable options are available on the Change Controller Properties option:

- ▶ **Set Boot Device:** Use this feature to select the primary bootable device. The options are None and [any drives detected by the controller].
- ▶ **Drive Replace:** This feature provides the option to copy data back from a hot spare drive to a different drive. The options are Enabled and Disabled.
- ▶ **Drive Replace on SMART Error:** This feature provides the option to start a Drive Replace operation if a Self-Monitoring Analysis and Report Technology (SMART) error is detected on a drive. The options are Enabled and Disabled.
- ▶ **Rebuild Rate:** Use this option to select the rebuild rate for drives connected to the selected controller. The default is 30%. The rebuild rate is the percentage of system resources dedicated to rebuilding a failed drive. The higher the number, the more system resources are devoted to a rebuild.
- ▶ **Background Initialization (BGI):** Use this option to select the amount of system resources dedicated to background initialization of virtual drives connected to the selected controller. The default is 30%.
- ▶ **Consistency Check Rate:** Use this option to select the amount of system resources dedicated to consistency checks of virtual drives connected to the selected controller. The default is 30%.
- ▶ **Reconstruction Rate:** Use this option to select the amount of system resources dedicated to the reconstruction of drives connected to the selected controller. Also use this option when performing a RAID Level Migration (RLM) or Online Capacity Expansion (OCE). The default is 30%.
- ▶ **Controller BIOS:** Use this feature to enable or disable the controller BIOS. If the boot device is on the selected RAID controller, the BIOS needs to be enabled. The options are Enabled and Disabled.
- ▶ **Coercion Mode:** Drive coercion is a tool for forcing drives of varying capacities to the same size so they can be used in a drive group. The coercion mode options are None, 128MB-way, and 1GB-way. The default is None. The number you choose depends on how much the drives from various vendors vary in their actual size. Use the 1GB coercion mode option.
- ▶ **SMART Polling:** Use this option to determine how frequently the controller polls for drives that are reporting a Predictive Drive Failure (SMART error). The default is 300 seconds (5 minutes).
- ▶ **Stop Boot Process on Error:** When set to Enabled, the system stops booting if there is any critical BIOS error that requires user intervention to resolve. The options are Enabled and Disabled.
- ▶ **Stop Consistency Check on Error:** Select Enabled to stop the consistency check operation on a redundant virtual drive if there is an inconsistency found in the data. The options are Enabled and Disabled.
- ▶ **Maintain Drive Fail History:** Use this feature to track any bad physical drives across reboot. The options are Enabled and Disabled.
- ▶ **Load Balance:** This feature can load balance I/O when redundant paths are detected. The options are Auto and Disabled.
- ▶ **Enable Auto Import:** Use this feature to automatically import foreign configurations without any user intervention. The options are Enabled and Disabled.
- ▶ **Persistent Hot Spare:** Select Enabled to turn the system backplane or storage enclosure drive slots into dedicated hot spare slots. If enabled, replacing a hot spare drive in the same slot automatically configures the drive as a hot spare. The options are Enabled and Disabled.

- ▶ **Manage Link Speed:** This feature allows the user to change the link speed between the controller and an expander, or between the controller and a drive that is directly connected to the controller. Devices detected by the controller are listed. The options are: Auto, 1.5 Gbps, 3 Gbps, and 6 Gbps. When finished, select OK to confirm changes or select Cancel to cancel the changes.
- ▶ **Patrol Read Mode:** Use this feature to change the patrol read mode for the controller. Patrol read is an operation that scans and resolves potential problems on configured drives. Select Auto for the patrol read to run continuously on the system based on a schedule. Select Manual to manually start the patrol read. The options are: Auto, Manual, and Disabled.
- ▶ **Patrol Read Rate:** Use this option to select the rate for patrol reads for drives connected to the selected controller. The default is 30%. The patrol read rate is the percentage of system resources dedicated to running a patrol read.
- ▶ **Start Patrol Read:** Use this item to start a patrol read for the selected controller.
- ▶ **Suspend Patrol Read:** Use this item to halt the patrol read after it starts. The user can resume patrol read later with no loss of data.
- ▶ **Resume Patrol Read:** Use this item to continue the patrol read after it is suspended.
- ▶ **Stop Patrol Read:** Use this item to stop the patrol read.
- ▶ **Patrol Read Setting for Unconfigured:** This feature is used to set the patrol read settings for the unconfigured space. The options are Enabled and Disabled.
- ▶ **Cache Flush Interval:** Use this option to control the interval (in seconds) at which the contents of the onboard data cache are flushed. The default is 4 seconds.
- ▶ **Preserved Cache:** This item displays if the controller cache is preserved because of missing or offline virtual drives. The cache is preserved until the virtual drive is imported or the cache is discarded. This information is read only.
- ▶ **Discard Preserved Cache:** Use this item to delete any data in the preserved cache.
- ▶ **Global Hot Spare for Emergency:** Use this feature to allow the use of otherwise incompatible global hot spare drives as emergency hot spare drives. The options are Enabled and Disabled. The default is disabled.
- ▶ **Unconfigured Good for Emergency:** Use this feature to allow the use of unconfigured good drives as emergency hot spare drives. The options are Enabled and Disabled. The default is disabled.
- ▶ **Emergency for SMARTer:** Use this feature to allow the use of Emergency Hot Spare drives for Predictive Failure Analysis (PFA) events. The options are Enabled and Disabled. The default is disabled.

Battery Management

The Battery Management feature displays battery information, battery capacity and status, and the option to run a battery learn cycle. The Battery Management information also displays if a battery replacement is required.

Schedule Consistency Check

The Schedule Consistency Check option is used to schedule Consistency Check (CC) operations. Consistency Check is a background operation that verifies and corrects the mirror or parity data for fault tolerant virtual drives.

The following functions are available with the Schedule Consistency Check option:

- ▶ Consistency Check Frequency: Use this option to select the frequency of the consistency check operations. The options are: Disable, Hourly, Daily, Weekly, and Monthly.
- ▶ Consistency Check Start: Use this option to select a start date for the consistency check schedule.
- ▶ Consistency Check Start Time: Use this option to select a start time for the consistency check schedule.
- ▶ Consistency Check Mode: This feature is used to select the mode for performing a consistency check over multiple virtual drives. Select *concurrent* to perform a consistency check on all drives at the same time. Select *sequential* to perform a consistency check on all drives, one drive at a time. The options are Concurrent and Sequential.

Clear Configuration

The Clear Configuration option is used to clear the configuration; however, it is gray if no configuration is found on the drives.

Manage Foreign Configuration

The Manage Foreign Configuration option is used to import, clear, or manage found foreign configurations on the drives. The option is gray when no foreign configuration is present.

Event and log options

The Save Controller Events option can be used to save the controller events if a files system, in which the logs can be saved, is seen by the controller. These logs are then accessible by navigating the UEFI Shell. This function is an option to save the controller clients; however, the suggested way to collect and view these logs is with the DSA tool.

Select the Clear Controller Events option to clear all events on the controller.

The Save TTY Log menu item that is shown in Figure 3-15 on page 95 can be used to save log information to a local file system, if one is available.

3.3.2 Virtual Drive Management

The Virtual Drive Management menu allows the Create Configuration feature to be a selectable function only until a virtual drive is created.

Figure 3-21 shows the Virtual Drive Management menu with the Create Configuration option selected.

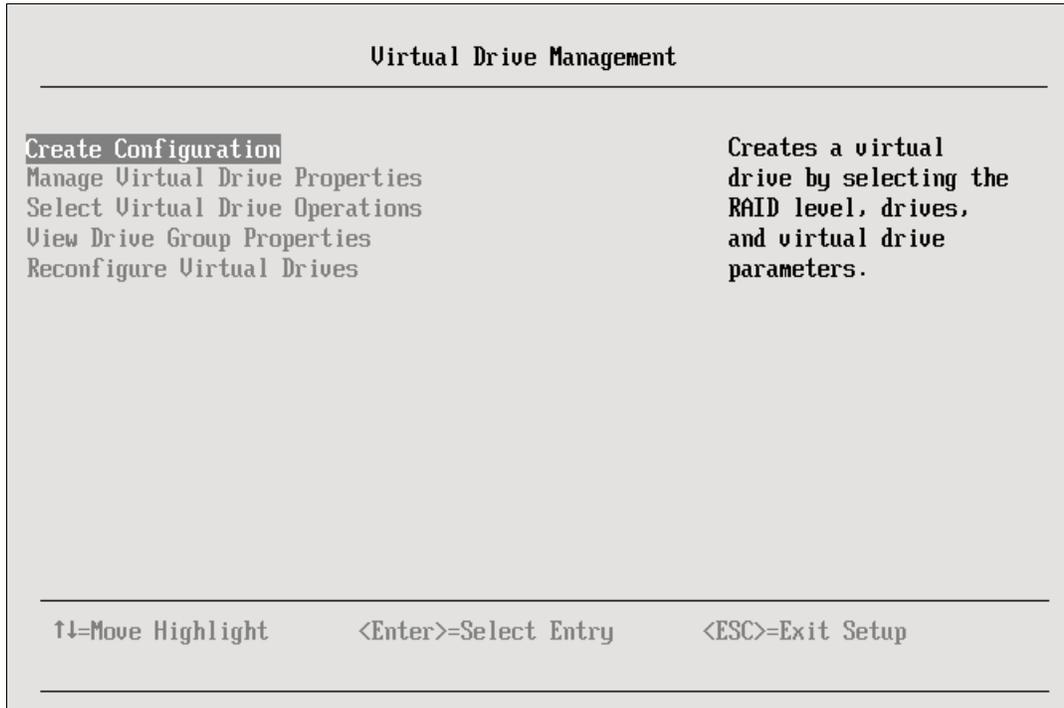


Figure 3-21 Virtual Drive Management window

The following subsections describe the Virtual Drive Management menu items:

- ▶ “Create Configuration” on page 102
- ▶ “Manage Virtual Drive Properties” on page 103
- ▶ “Select Virtual Drive Operations” on page 104
- ▶ “View Drive Group Properties” on page 104
- ▶ “Reconfigure Virtual Drives” on page 104

Create Configuration

The Create Configuration option creates a virtual drive by selecting the RAID level, drives, and virtual drive parameters. Figure 3-22 shows the Create Configuration options.

The screenshot shows a terminal-style window titled "Create Configuration". It contains several sections of text and input fields. The first section is "Save Configuration" with "Select RAID Level" set to "<RAID0>". The second section is "Select Drives From" set to "<Unconfigured Capacity>". The third section is "Select Drives". Below that is a section titled "CONFIGURE VIRTUAL DRIVE PARAMETERS:" with fields for "Virtual Drive Name" (set to "-"), "Virtual Drive Size" (set to "[0]"), "Virtual Drive Size Unit" (set to "<GB>"), "Strip Size" (set to "<128 KB>"), "Read Policy" (set to "<Adaptive>"), "Write Policy" (set to "<WBack>"), "I/O Policy" (set to "<Direct>"), and "Access Policy" (set to "<Read/Write>"). A red arrow points down from the "Access Policy" field. At the bottom, there are three instructions: "↑↓=Move Highlight", "<Enter>=Select Entry", and "Esc=Exit".

Figure 3-22 Create Configuration window

Follow these steps to create a virtual drive (see Figure 3-22):

1. Select the RAID level.
The options that are available are RAID 0 and 1 unless more support for 5, 50, 6, and 60 are added to the server with FoD licenses. .
2. Select the drives from the Unconfigured Capacity or Free Capacity.
3. Highlight **Select Drives** and press Enter.
4. To see the available drives for selection, be sure to properly select the media type HDD, SSD, or BOTH. Highlight and select **Apply Changes** to create the drive. You receive the following Success message, as shown in Figure 3-23.

The screenshot shows a terminal-style window titled "Success". It contains a message: "The operation has been performed successfully. Some operations may take several minutes or hours to complete. View the Operation Progress to check on the percentage complete." Below the message is a small box containing the text "OK".

Figure 3-23 Successful selection of drives

5. To complete the virtual drive creation, be sure to select **Save Configuration** from the Create Configuration menu. Click **Enter** after making any other settings changes to the Virtual Drive Parameters in the Create Configuration menu. Figure 3-24 shows the warning that is displayed. Select **Yes** and click **Enter** to continue.

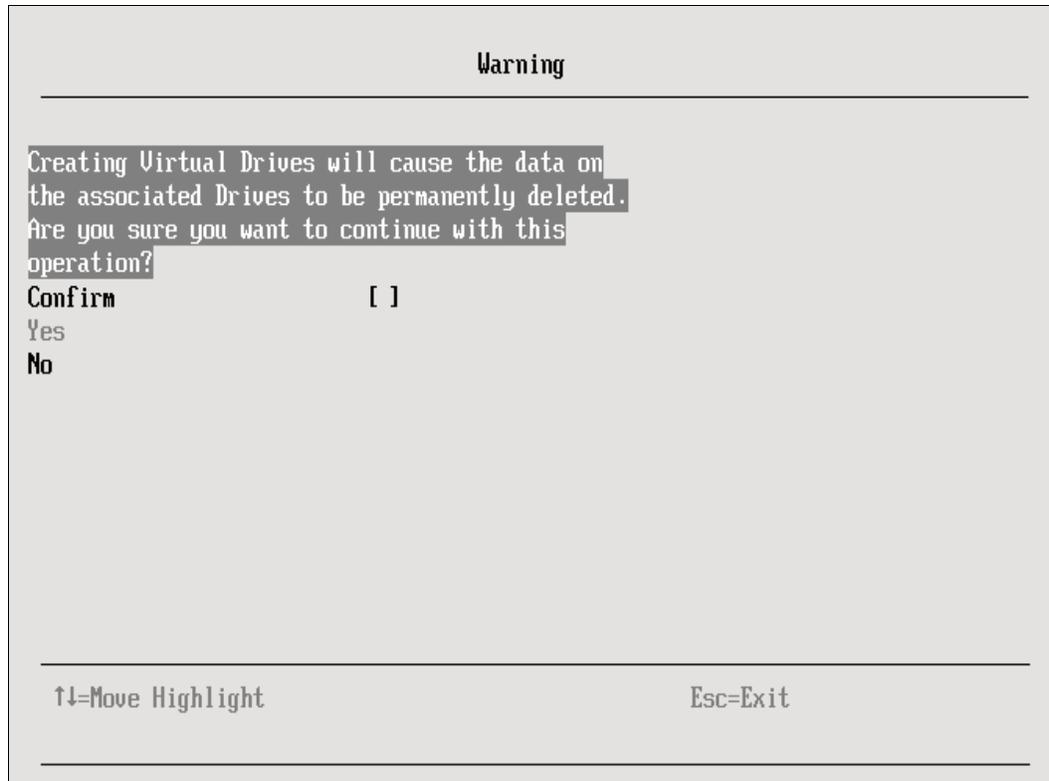


Figure 3-24 Warning creating virtual drives

Manage Virtual Drive Properties

The Manage Virtual Drive Properties function displays and manages virtual drive properties, such as the virtual drive health and status. This feature allows some changes to be made such as virtual drive policies.

Select Virtual Drive Operations

Choose the Select Virtual Drive Operations menu item to run virtual drive operations. Figure 3-25 shows this option.

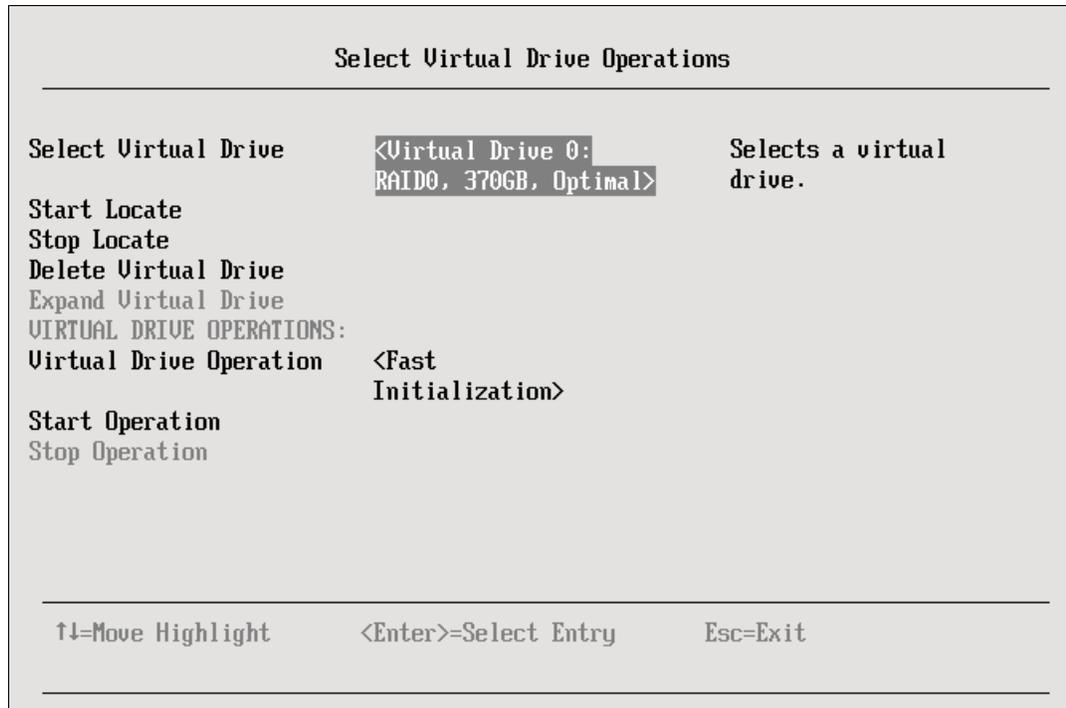


Figure 3-25 Select Virtual Drive Operations

The following options are included in the Select Virtual Drive Operations window:

- ▶ Flash the LEDs on all associated drives (Start Locate, Stop Locate).
- ▶ Delete the virtual drive.
- ▶ Expand the capacity of the virtual drive by adding more drives.
- ▶ Perform operations on the virtual drive: Fast Initialization, Slow Initialization, Check Consistency, and Virtual Drive Erase. The Fast Initialization function overwrites the first and last 8 MB of the virtual drive, clearing any boot records or partition information.

View Drive Group Properties

The View Drive Group Properties function displays the information about the available drive groups, associated virtual drives, and the capacity allocation.

Reconfigure Virtual Drives

Select the Reconfigure Virtual Drives option to perform a RAID Level Migration (RLM). RLM converts a virtual drive to a different RAID level.

Have a backup: Have a known, good, verified backup before attempting RLM on virtual drives with existing data.

3.3.3 Drive Management

Select the Drive Management menu to view drive properties and perform tasks. Choose this function to assign a hot spare, locate specific drives, force drives offline and online, and rebuild drives after a drive failure. Figure 3-26 shows the Drive Management menu options.

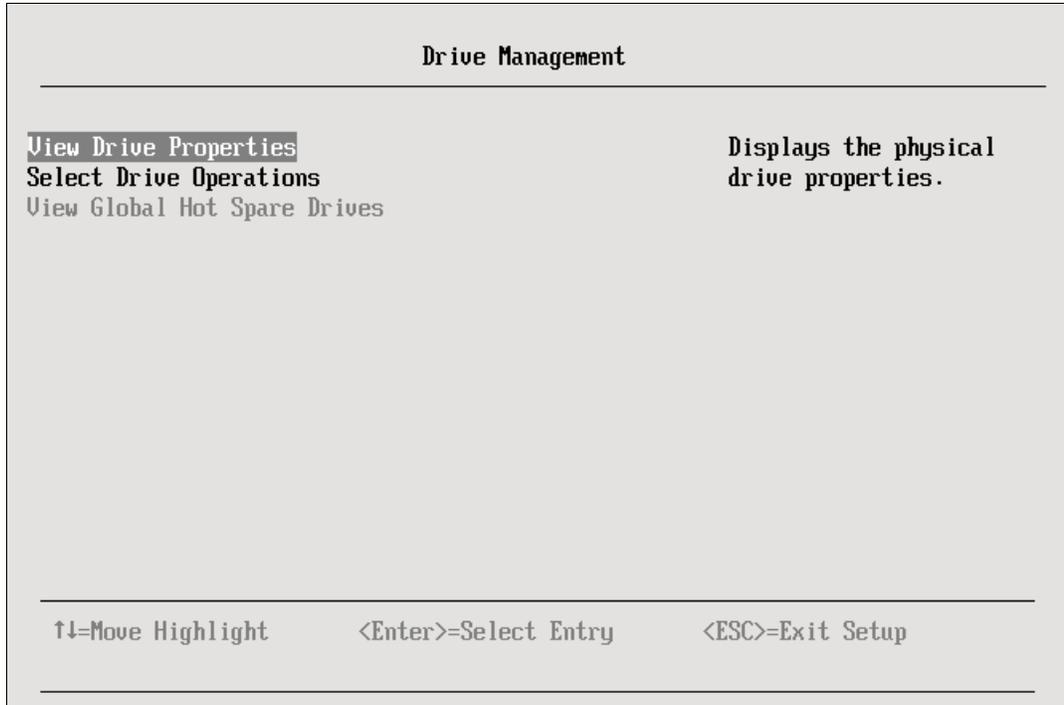


Figure 3-26 Drive Management

The following options are available in the Drive Management window:

- ▶ View Drive Properties: This option displays the physical drive properties of each individual drive that is seen by the controller.
- ▶ Select Drive Operations: This feature runs the drive operations.
- ▶ View Global Hot Spare Drives: This feature becomes an active menu option when a global hot spare is defined.

3.3.4 Enclosure Management

Select the Enclosure Management panel to view any attached enclosures, including information about the enclosure firmware level, and the drives installed in the enclosures. Figure 3-27 shows an example.

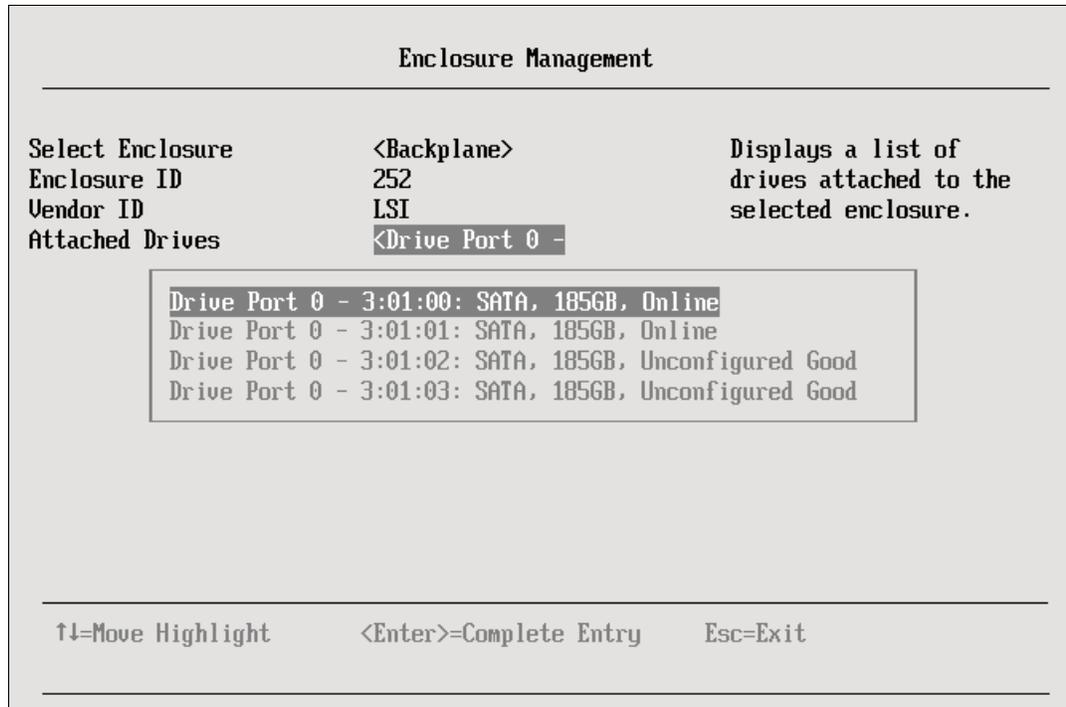


Figure 3-27 Enclosure Management

3.4 Configuring embedded Ethernet using UEFI

As described in 1.14, “Network adapters” on page 53, the x3750 M4 offers two integrated Gigabit Ethernet ports. In addition, two 10 Gb Ethernet ports can be added by installing one of the available dual-port 10 Gb Ethernet cards that are listed. Table 1-27 on page 51 shows these options (88Y7429 or 88Y7427). These cards use a dedicated connector on the system board and do not use a PCI expansion slot. The optional cards are used to route the ports to external RJ45 or SFP+ ports.

Both the Gigabit Ethernet ports and the 10 Gb Ethernet ports are driven by the onboard Emulex BE3 controller.

This section describes the configuration options available for the embedded Emulex network adapter in the UEFI.

To access the settings menu to configure the Ethernet adapters, reboot the server and press **F1** during the boot sequence to enter UEFI. Select **System Settings** → **Network**. Figure 3-28 is displayed.

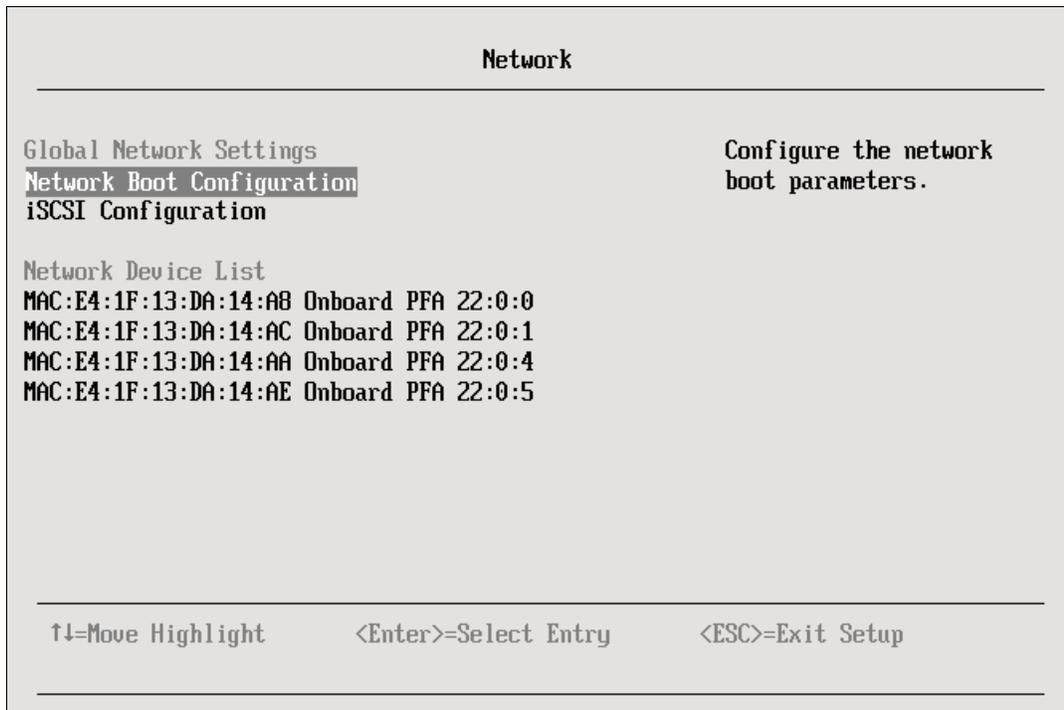


Figure 3-28 Network settings in UEFI

The following options are available on the Network window:

- ▶ **Global Network Settings:** This option is gray until an initial network configuration is completed. When the configuration is complete, this option becomes available.
- ▶ **Network Boot Configuration:** Select this option to configure the network boot parameters. A list of each adapter is displayed. Select the adapter to configure by the MAC address that is listed to enter the boot configuration menu for that adapter. See Figure 3-29 on page 108 for the network boot configuration for an adapter.

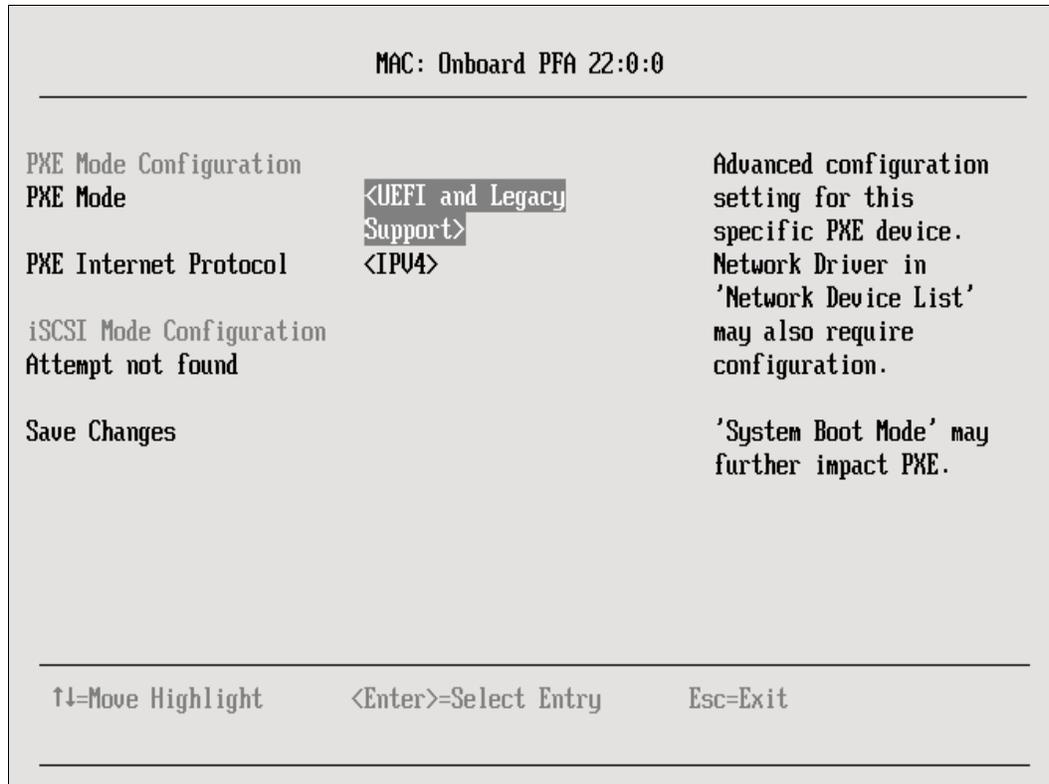


Figure 3-29 Network Boot Configuration for an adapter

- ▶ Select the iSCSI Configuration to configure the Small Computer System Interface (iSCSI) parameters by selecting the options that are shown in Figure 3-30 on page 109.

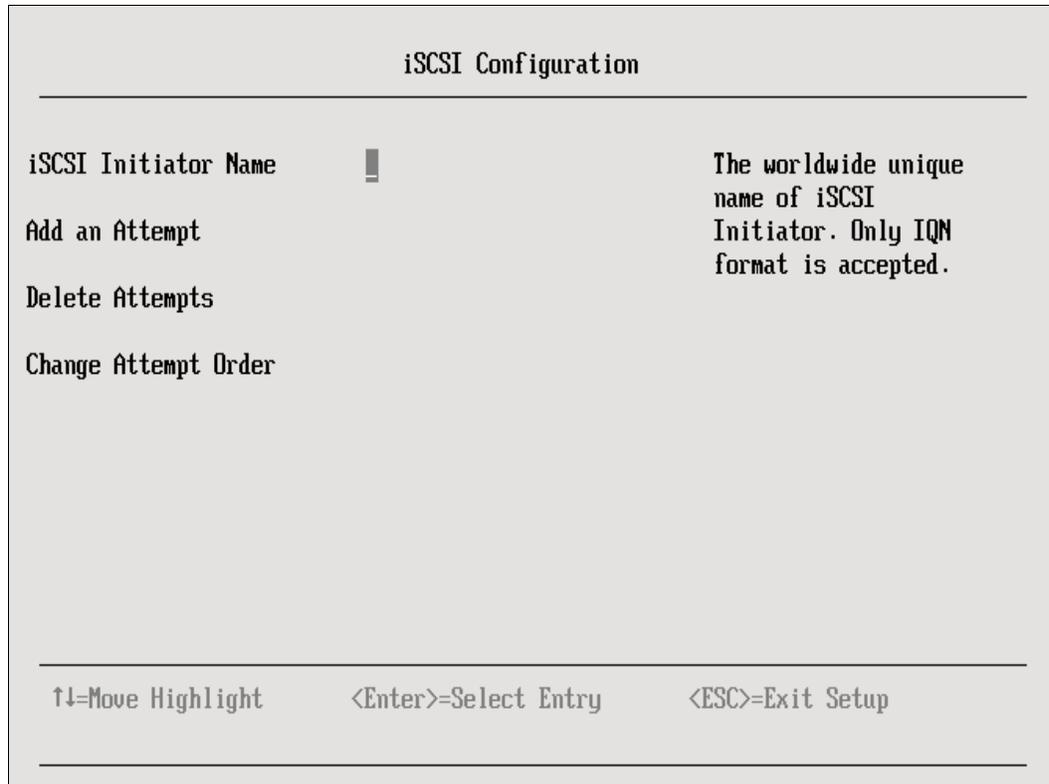


Figure 3-30 iSCSI Configuration

- ▶ The Network Device List is followed by the list of each Ethernet adapter and is represented by the MAC address. Each device can be highlighted and accepted to then configure the network parameters for the individual adapter. Figure 3-31 on page 110 shows the Ethernet adapter configuration.

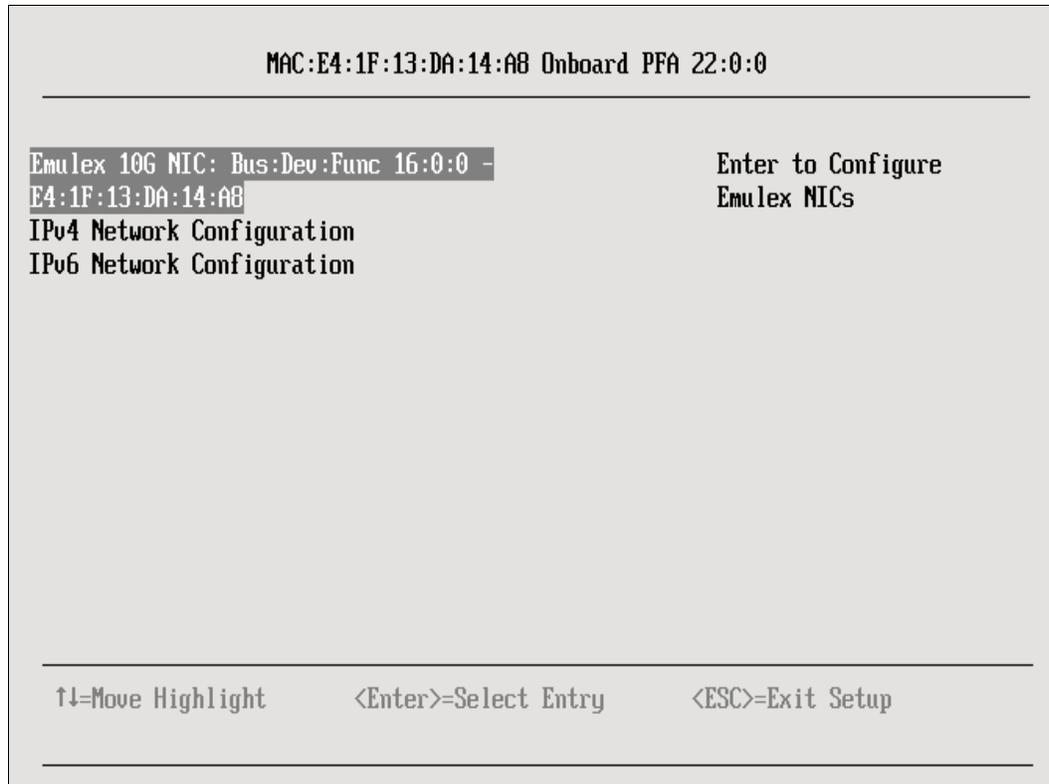


Figure 3-31 Ethernet adapter configuration

The following steps describe how to configure an adapter:

1. Select the adapter, in this case the Emulex 10 Gb, to configure or view the network adapters. Figure 3-32 shows the Emulex network interface controller (NIC) Selection window.

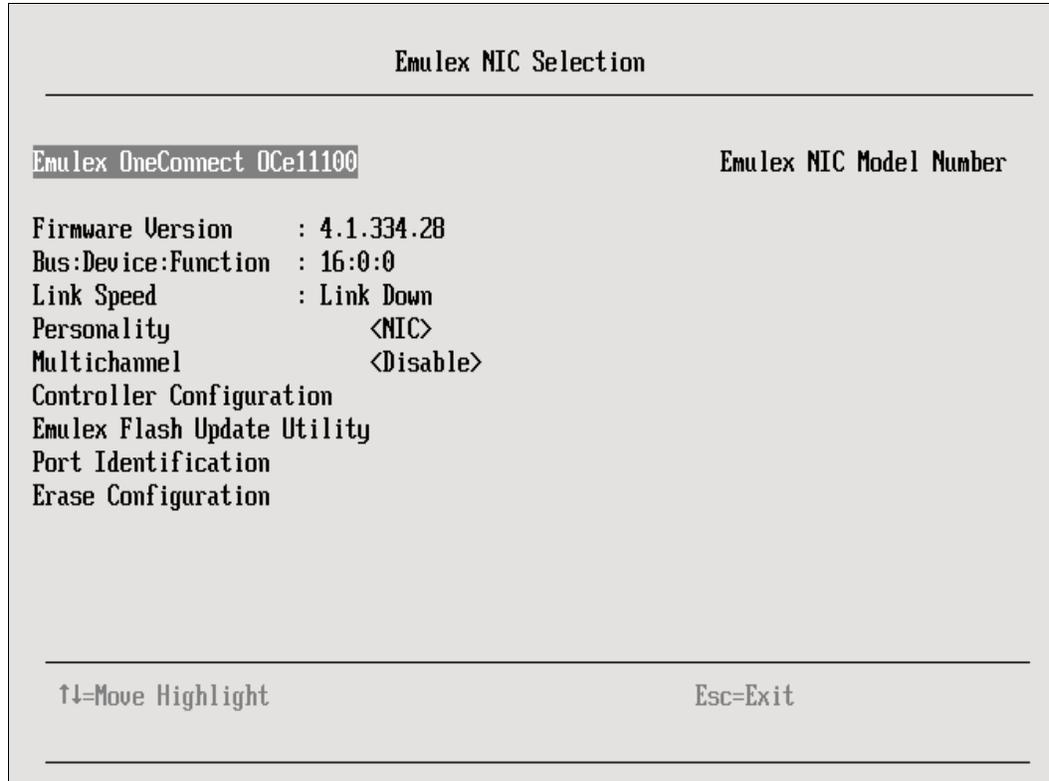


Figure 3-32 Emulex NIC Selection

2. Use the arrow-down key to select the IPv4 Network Configuration or IPv6 Network Configuration to assign the IP settings to the adapter. Figure 3-33 shows the IPv4 configuration.

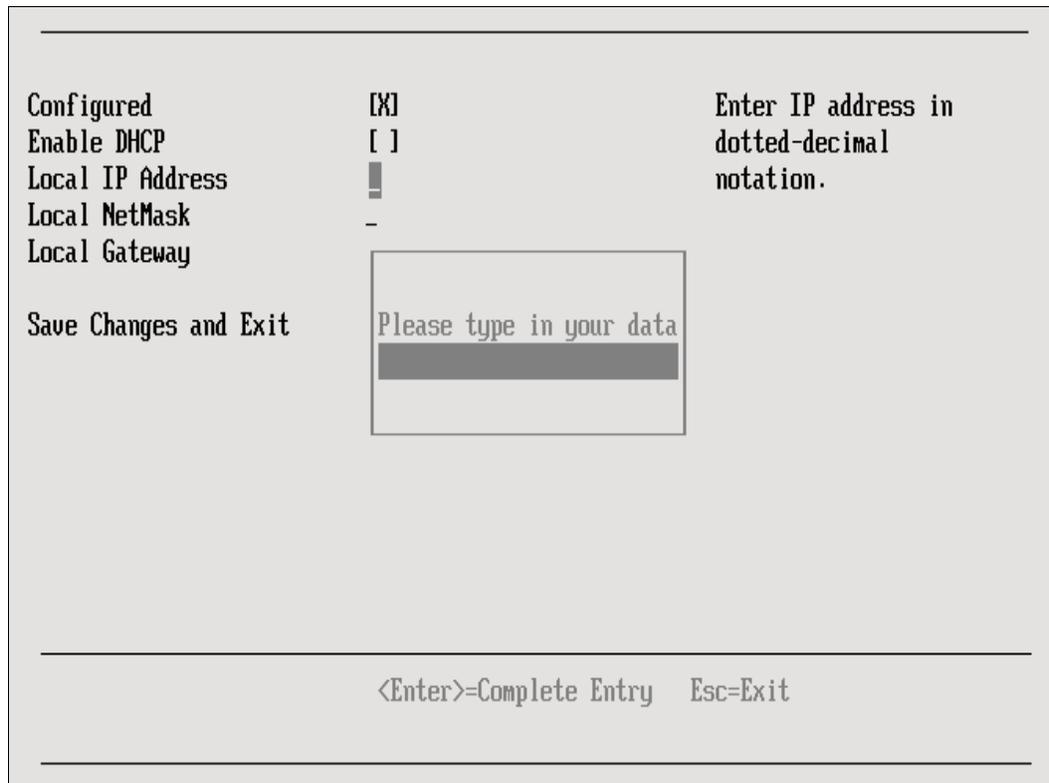


Figure 3-33 IPv4 setting

Installation and management

This chapter describes readying the System x3750 M4 for production.

We describe firmware updates, how to install an operating system (OS), and the currently supported operating systems from which to choose. Hints and tips are also featured.

We also describe post-installation management and the options for remote management and setting up notifications. Information is also included about how to deploy the firmware updates as they become available after a system is added to a production environment.

The following topics are described in this chapter:

- ▶ 4.1, “Firmware considerations” on page 116
- ▶ 4.2, “Supported operating systems” on page 126
- ▶ 4.3, “Installing an operating system” on page 127
- ▶ 4.4, “Post-installation management tools” on page 131

4.1 Firmware considerations

Firmware updates are released for servers to improve functionality, to resolve problems that are found, and to add support for additional hardware and software. During the time a server is purchased and when that server is first powered up, there might be additional firmware updates released for the server and its components.

When planning to install an entire server or options (for example: adapters, drives, or memory) to an existing server, firmware must always be the first place to look to ensure the new hardware that is added is compatible with the existing hardware or the base server. As new options are added to the supported configuration it is possible that it only became supported because of a firmware update.

Before beginning the installation of any OS or additional hardware, it is best to update the firmware. If you are adding hardware after an OS is installed, the suggested first step is to update the driver in the OS for any possible hardware components. This step is best taken before the firmware is updated.

Lenovo provides many tools to make the firmware and driver update less hands-on and more streamlined. These tools range from options to update one server to options for deploying updates across multiple machines.

Before listing the individual tools available for the firmware and driver updates, it is important to know that all of these tools are available at <http://lenovo.com/support>. There are multiple ways to view which updates are needed for one system or one component.

4.1.1 Firmware update tools and preferred practices

The following supported tools are described in detail in this chapter:

- ▶ Bootable Media Creator™
- ▶ UpdateXpress System Pack Installer
- ▶ System Director Update Manager

These tools streamline the firmware and driver update processes by ensuring that prerequisites and installation orders are met to ensure more accurate updating with less chance of failure. For more information about the direction of the firmware update processes, see *System x Update Tools Transition* at this website:

<http://ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5082923>

This link also contains *Firmware Update Best Practices*. It is suggested that you follow this guide when planning any firmware updates to System x servers with the IMM2. The guide goes into detail about considerations to make and how firmware updates can fail when the best considerations are not put into place. If a firmware update fails, review the recovery processes mentioned in this chapter. If a firmware update failure continues, contact Lenovo Support.

4.1.2 Fix Central

Fix Central is the online source to find all the available updates for a system type. This search can be narrowed by a supported component for that machine type or the type of OS installed.

Fix Central is accessible through the Lenovo Support Portal:

<http://www.lenovo.com/support/>

This support portal is a great starting point for finding all the necessary downloads, documentation, upgrade options, and troubleshooting guides for your machine. You can create a login or sign in with an existing IBM.com login to customize a view with any added machine types or previously searched types.

After a machine type is added to your list of searches, Fix Central is accessible by using the Download tab at the top of the Support home page.

Access Fix Central and begin your search by machine type by accessing this website:

<http://www.lenovo.com/support/>

Fix Central can be used to search for, select, and download the following items: fixes, updates, drivers, and management and troubleshooting tools for a specific machine or OS. This feature is further customized by the types of downloads available, such as HTTP, FTP, and Download Director.

Fix Central lets you find fixes in two ways:

- ▶ Select Product - uses filters to narrow down your selection and then browse for the product
- ▶ Find Product - search for a product

Figure 4-1 shows the Fix Central window.

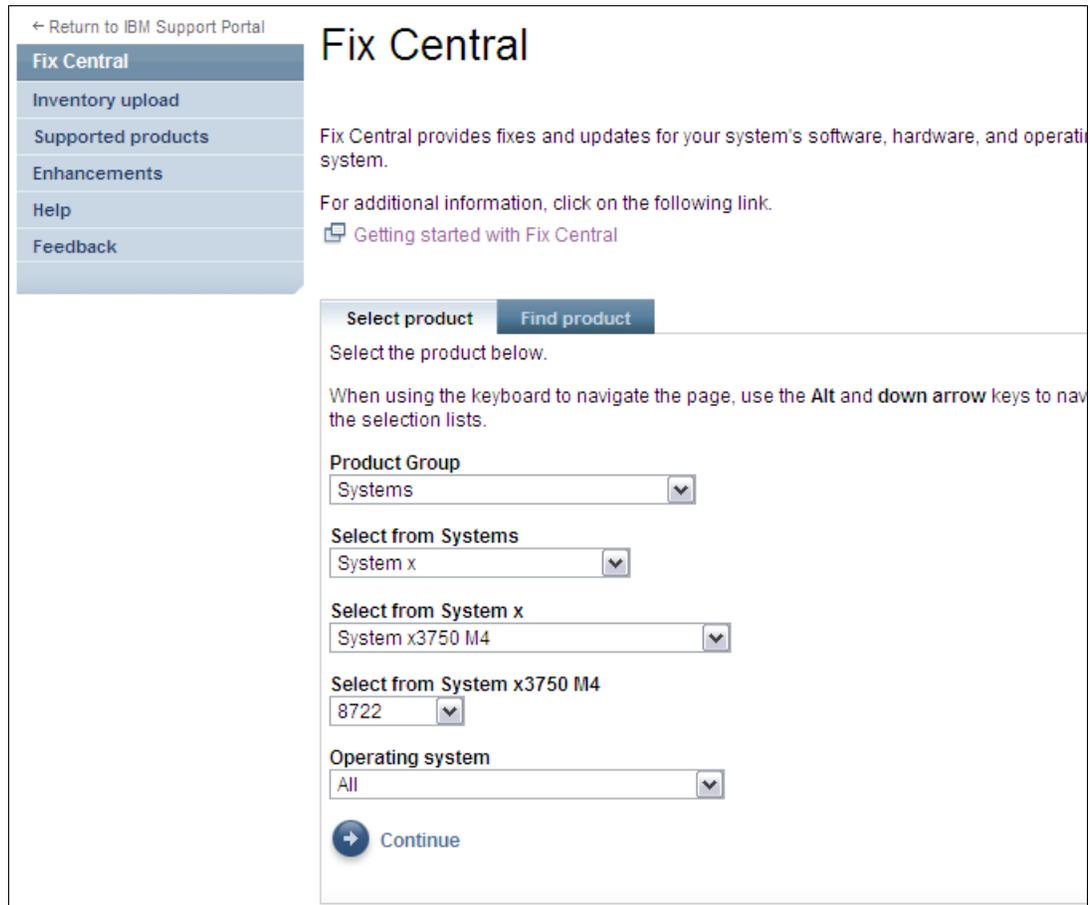


Figure 4-1 Fix Central window

4.1.3 Bootable Media Creator

The Bootable Media Creator (BoMC) is a tool for the creation of a customized optical media that is ideal for systems without a functioning OS. This tool is also for OSs that are not able to run updates or troubleshooting tools.

We describe how to use it to deploy updates on a machine-by-machine basis in an offline environment.

To obtain the latest version of the BoMC tool, see this website:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=T00L-BOMC>

Consider the following functions when using the BoMC tool:

- ▶ Ensure that when you are selecting the BoMC tool for an OS, you choose the OS that matches the server you are running this tool from to create the bootable media. This selection is not necessarily the OS on the target server.
- ▶ When downloading the BoMC tool, review the readme file. This file provides additional instructions on how to use the tool.
- ▶ The tool, when used for the firmware updates, does not update the OS drivers. Drivers and firmware levels often need to be kept in sync; therefore, ensure that *before* doing the firmware updates that the prerequisite drivers are updated.

The following examples are taken from the BoMC tool that runs in Microsoft Windows and are used just to create update media. The result is to create a customized ISO to burn to either a CD or DVD:

1. Download and run the BoMC tool. If asked, be sure to read and accept the license agreement.
2. From the Welcome window, you can choose to load a previously saved configuration or to check for the latest version of this tool. Select either one, if applicable. If not, click **Next**. Figure 4-2 shows the window that is displayed.

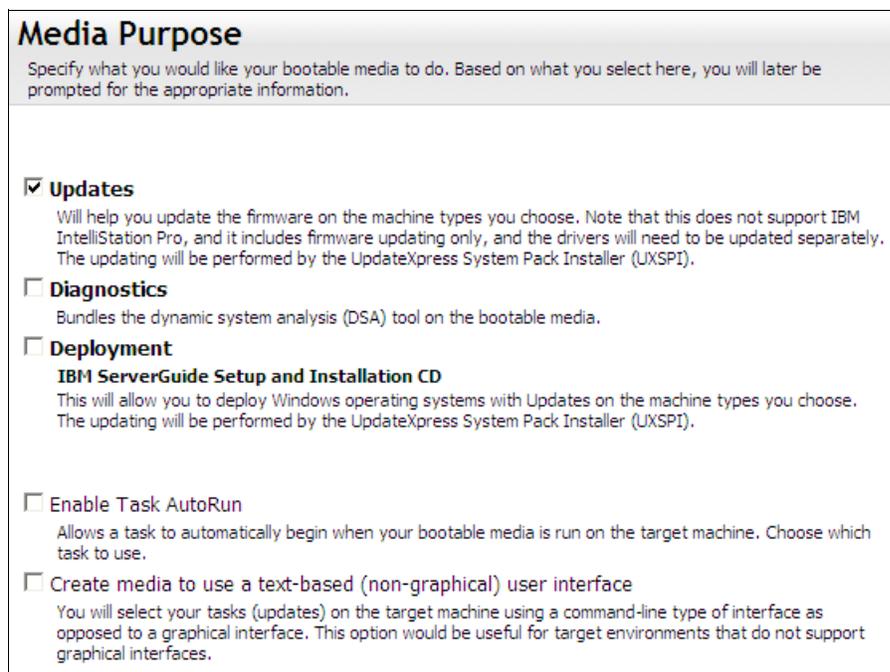


Figure 4-2 BoMC Media Purpose window

3. The Media Purpose by default is set to Updates. However, this tool can be used to create a bootable Diagnostics media, which bundles in the Dynamic System Analysis (DSA) tool. This function can also be used to create a Deployment image, which includes the ServerGuide Setup and Installation Media.

The remaining two options are: Enable Task AutoRun, and Create media to use a text-based (non-graphical) user interface.

We selected **Updates** and clicked **Next**. Figure 4-3 is displayed.

Acquire Location
You can choose to acquire updates, tools and bootable image from IBM web site or from local directory.

- Check the IBM web site** - Downloads the appropriate updates automatically from the IBM site.
 - UpdateXpress System Packs (UXSP's)** - UpdateXpress System Packs contain an integration-tested bundle of online, updateable firmware updates for each System x and BladeCenter server. This is the preferred method to obtain firmware for the server.
 - Latest available individual updates** - Check the IBM web site for the latest individual version of each firmware package. This is the preferred method when you wish to install the latest updates or when IBM support instructs you to install the latest updates. This option may download newer updates and hotfixes, if available, than the UXSP option above.
- Look in a local directory** - Specify a directory on the local file system containing specific individual files to include in the bootable media. The directory should have been populated with the required files either in a previous session of this tool, or manually. See [here](#) for specific requirements on manually obtaining required files.

Figure 4-3 BoMC Acquire Location window

4. The Acquire Location function is important to consider when accessing the updates from IBM.com. By default, the selection is to receive the UpdateXpress System Packs™, which are updates that are grouped together to ensure compatibility. This default is the suggested selection, but it means that you might not have the latest updates available. To ensure that your machine is updated to the latest possible code, choose Latest Available Individual Updates. If the updates are saved to a preconfigured repository, choose to browse to the folder in which they are saved.

- The next step is to set the HTTP Proxy settings (if needed) to allow the tool to access the Internet if it was selected to Check the Lenovo website for downloads. Click **Next** to specify the target system. Figure 4-4 shows the Targeted Systems window.

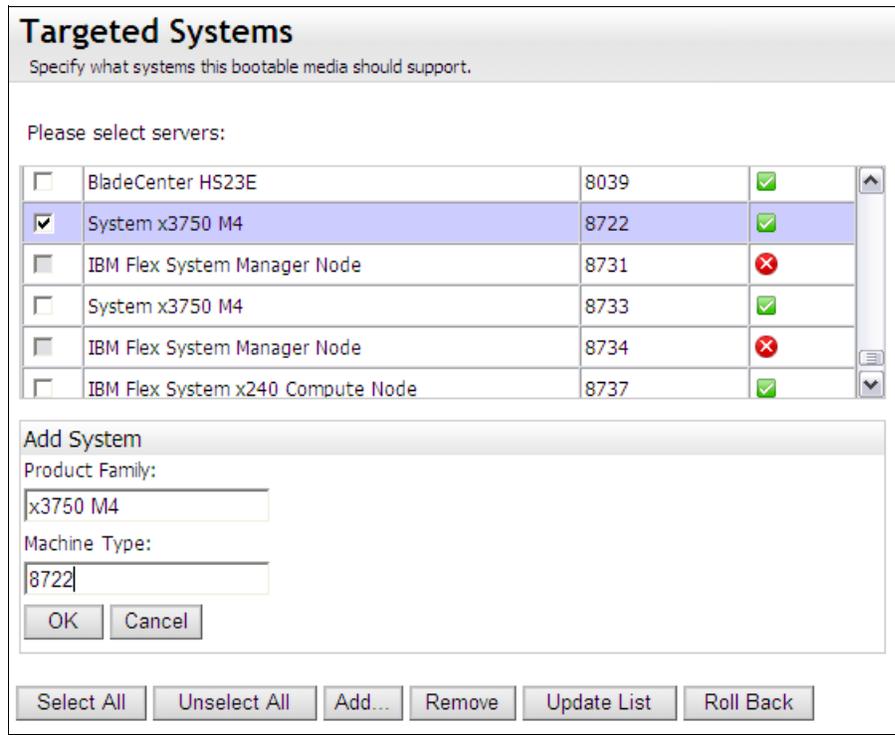


Figure 4-4 Targeted Systems window

- Choose the targeted system for updates. It is possible to choose more than one machine type to use the media to update multiple machines. The systems are listed in numerical order by the four-digit machine type. If the machine type needed is not available in the list, click **Add** and enter the system Product Family then Machine Type and click **OK**. This selection adds the machine to the list to be checked. In our example, we just selected the System x3750 M4. Click **Next**.
- You are now asked to specify a Target Directory where the downloads and created media file are saved. To change the location, click **Browse** to select a different directory. Next you are asked to specify the output media format, as shown in Figure 4-5.



Figure 4-5 Media Format

8. The Media Format window allows you to select which type of bootable device to create:
 - CD/DVD: Creates an ISO file or allows you to burn the files directly to optical media
 - USB: Formats the attached USB drive and makes it bootable
 - PXE: Creates a network bootable media format to deploy updates over the network
 If creating an ISO image file, you can choose the folder in which to write the file.
9. The Unattended Mode Configuration function helps you to configure your created image to run in an unattended mode. In this mode, this feature will upload the log files on to your TFTP server, FTP server, your network file share (NFS or Samba), or your USB drive. This feature then shuts down your clients after the firmware update process. If you want to upload the log files on to the TFTP, FTP server, NFS, or Samba server, ensure that the directory is created and anonymous access is granted. The default setting is the Do Not Use Unattended Mode.
10. Confirm your choices in BoMC and click **Save** at the bottom of the page. This selection saves the settings for the next time they are needed. Figure 4-6 shows the Confirm Choices window.

Confirm Choices

That is all of the information necessary to begin creating your bootable media. Review your selections below and click the Next button when you are ready to begin media creation.

Machine Type	8722
Update Location	IBM Web Site
Use Text-Based User Interface	NO
Update Type	Latest Individual updates
Proxy IP Address	NO
Proxy Port Number	NO
Proxy User Name	NO
Media Purpose	Updates
Enable Task Auto Run	NO
Image File Name	C:\Program Files\Mozilla Firefox\workingdir\bootable.iso
Target Directory	C:\Program Files\Mozilla Firefox\workingdir
Overwrite Existing Image	NO
Include VMWare ESXI Update	NO
Use Unattend Mode	NO
Descriptive Name	Bootable Media - Tue May 22 10:59:54 2012

Save Configuration
You can save your choices to a file and import them into this tool later.

Figure 4-6 Confirm Choices in BoMC

11. The Creation Progress page shows the progress made while downloading the firmware updates and the tools needed to apply them. The progress function is complete when the bootable media is created. When this step is complete, click **Next**, then click **Finish**.

After the bootable media is created, in the case of an ISO, it can be burned to a CD or DVD. Or, if using the Remote Control function of the IMM2, the media can be remotely mounted to boot the machine.

Updating firmware with BoMC

To start running the tool to perform the update, reboot the server to the mounted media, CD or, DVD. This reboot launches a graphical user interface (GUI) for the ToolsCenter™ Customized Media (unless Unattended Mode, Task Autorun, or Text Mode were configured). Figure 4-7 shows the ToolsCenter Customized Media window.



Figure 4-7 Welcome to ToolsCenter

Click the **Updates** link on the panel. Depending on the tasks that you selected in BoMC, you see Updates, Diagnostics, or ServerGuide Installation Media in the window.

To begin the update process, click **Updates** for the tool to display what machines are available to update, then select **click here to start update**. This selection begins the process of the tool by comparing firmware files on the bootable media with the installed code versions on the server.

When the list of suggested updates is compiled, you have 60 seconds to stop the tool from continuing on and applying all the possible updates found. This time gives you the opportunity to clear the items you do not want to update.

The update process completes as the prerequisites require. When the updates are complete the machine can be rebooted.

Reboot tip: When the BoMC creator tool completes the updates to the UEFI firmware, do not shut down the server unless otherwise directed. Exit the tool and allow the machine to reboot for the new version of UEFI to become active.

4.1.4 UpdateXpress System Pack Installer

The UpdateXpress System Pack Installer (UXSPI) is a tool that can be used to update a server directly or to update a remote server (the latter is via a command-line interface (CLI) only). With UXSPI, you can also create a repository of firmware and driver updates to be applied later.

UXSPI is used to apply firmware updates to a server that already has a supported operating system (OS) installed. The tool will collect all the needed firmware updates and apply those updates to the local server. UXSPI can also update drivers, reducing the chance of a firmware update not matching an installed driver and leaving hardware or software in an unsupported configuration.

The following operating systems are supported by UXSPI Version 9.2:

- ▶ Microsoft Windows
- ▶ SUSE Linux Enterprise Server 11 Edition
- ▶ SUSE Linux Enterprise Server 10 Edition
- ▶ Red Hat Enterprise Linux 6 Edition
- ▶ Red Hat Enterprise Linux 5 Edition
- ▶ Red Hat Enterprise Linux 4 Edition
- ▶ VMware ESX 4

Previous versions of UXSPI are also available and support older OS versions.

When UXSPI runs from the server, it is intended to update. It auto-senses the machine type and the OS that is installed. This feature allows for easy gathering of specific updates and drivers. But if this tool is used to create an update repository, additional machine types and OSs can be selected.

Follow these steps to guide you through using the UXSPI tool to download firmware and driver updates, and applying them to the server through the OS:

1. Download the latest version of the UpdateXpress System Pack Installer for the OS that is running on the server that needs the updates:
<http://ibm.com/support/entry/portal/docdisplay?lnocid=SERV-XPRESS>
2. Execute the tool on the server needed. The tool itself does not need to install in the OS.
3. Read and accept the license agreements to begin using the tool.
4. The Welcome panel is displayed. You can then choose to search for the latest version of the tool, then click **Next**.
5. The Update Task panel, shown in Figure 4-8, allows you to choose to update the local machine or create a repository of updates. If you choose to create a repository, no updates are applied at that time. Click **Next**.

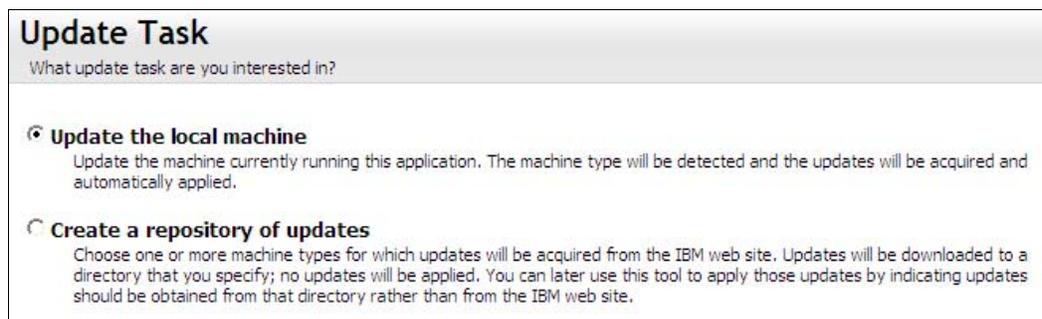


Figure 4-8 UXSPi Update Task

6. Next, choose where to acquire the update files from, as shown in Figure 4-9 on page 124. If a repository is not created, keep the selection **Check the website**. The default action of the tool updates all the drivers and firmware versions that are not current. Click **Next**.

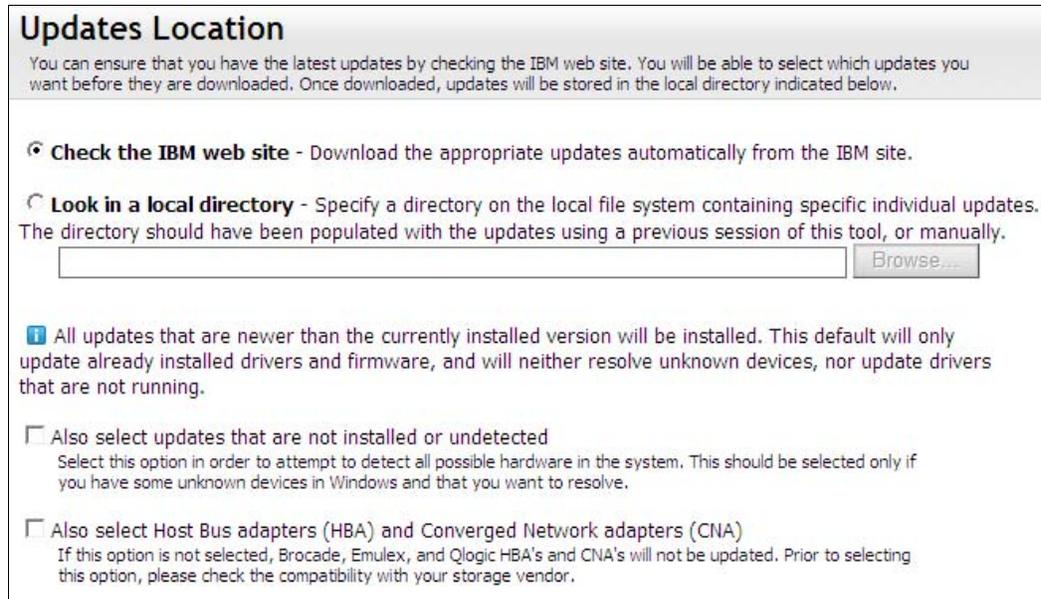


Figure 4-9 UXSPI Updates Location window

7. Next, select the Update Type by choosing between the bundled updates by using the UpdateXpress System Packs, or the Latest Available Individual Updates. Read the selections carefully to assist with your decision and click **Next**. Figure 4-10 shows the Update Type window.



Figure 4-10 UXSPI Update Type window

8. For the HTTP Proxy Settings, check the box only if a proxy server is needed for the tool to access the website and click **Next**.
9. Browse to the Target Directory. This directory contains the updates and tools downloaded by the UpdateXpress System Pack Installer. Click **Next**.
10. The tool now acquires the updates and provides a progress window to follow.
11. Under the Update Recommendation section, any code found to be lower than the downloaded version, is marked for update. Use the check boxes to mark and clear any selection. Click **Next**.
12. When the updates are successfully completed, the Update Selection is displayed. Click **Finish** to exit the UXSPI tool.

For more information about the UpdateXpress System Pack Installer, see the *UpdateXpress System Pack Installer (UXSPI) User's Guide* at the following URL. This guide provides for detail about each use of the tool and the command-line usage.

<http://ibm.com/support/entry/portal/docdisplay?lnid=MIGR-5085892>

4.1.5 Updating firmware through the IMM2

The IMM2 allows for the firmware to update through the web browser and the CLI.

Use as directed: This method is not suggested except when following recovery procedures, or as directed by support personnel. Instead, use BoMC or UXSPI as described in this chapter.

Follow these steps to reach the firmware update section of the IMM2 web interface:

1. Log in to the IMM2 by using the assigned IP address, default 192.168.70.125, and use the assigned login credentials. The default login is USERID and PASSWORD (with a zero instead of the letter O).
2. Click **Server Management** → **Server Firmware** as shown in Figure 4-11.

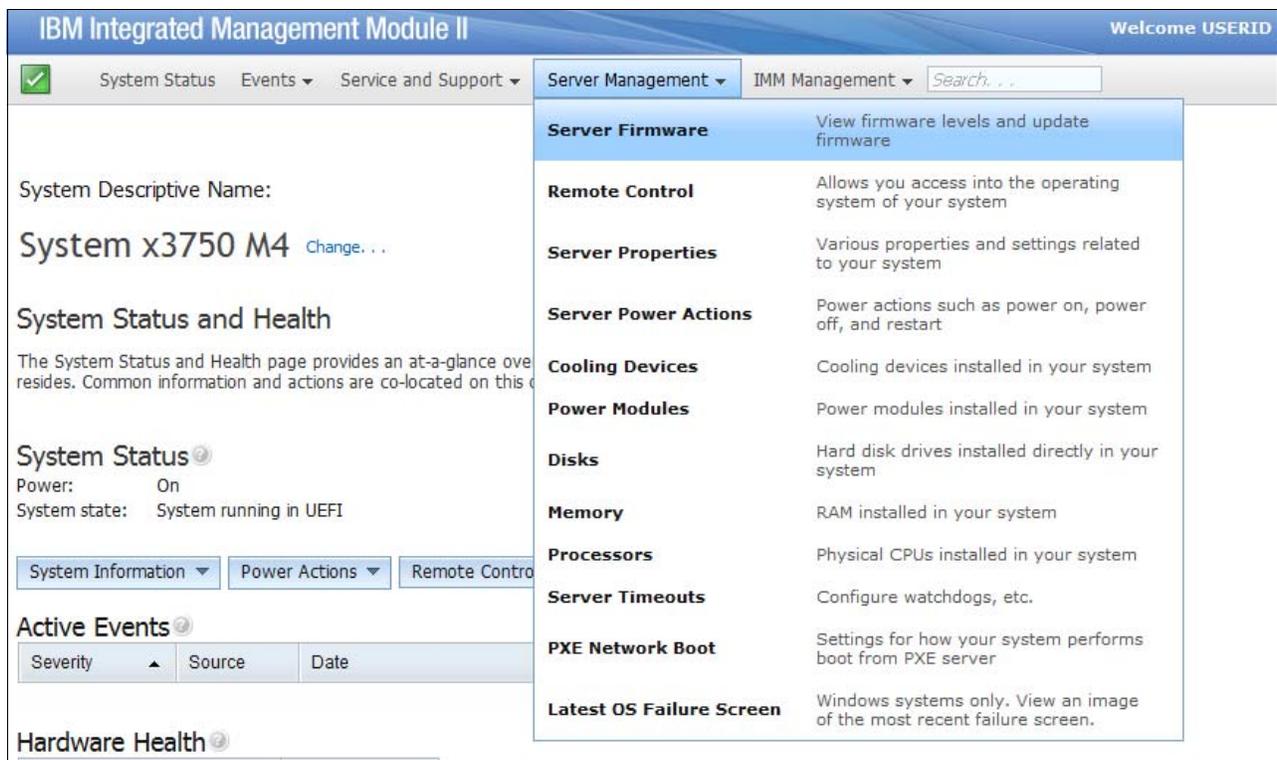


Figure 4-11 IMM2 server firmware

3. Click **Server Firmware** to display the current system firmware levels and to find Update Firmware.

This action displays a pop-up window to begin the firmware update process. The Update Server Firmware window displays a warning message, as shown in Figure 4-12.

Warning message: Please note that updating server firmware via this wizard is intended for recovery purposes only. It is strongly suggested that updates be performed using the UpdateXpress System Pack Installer, available as a download from the Lenovo website.

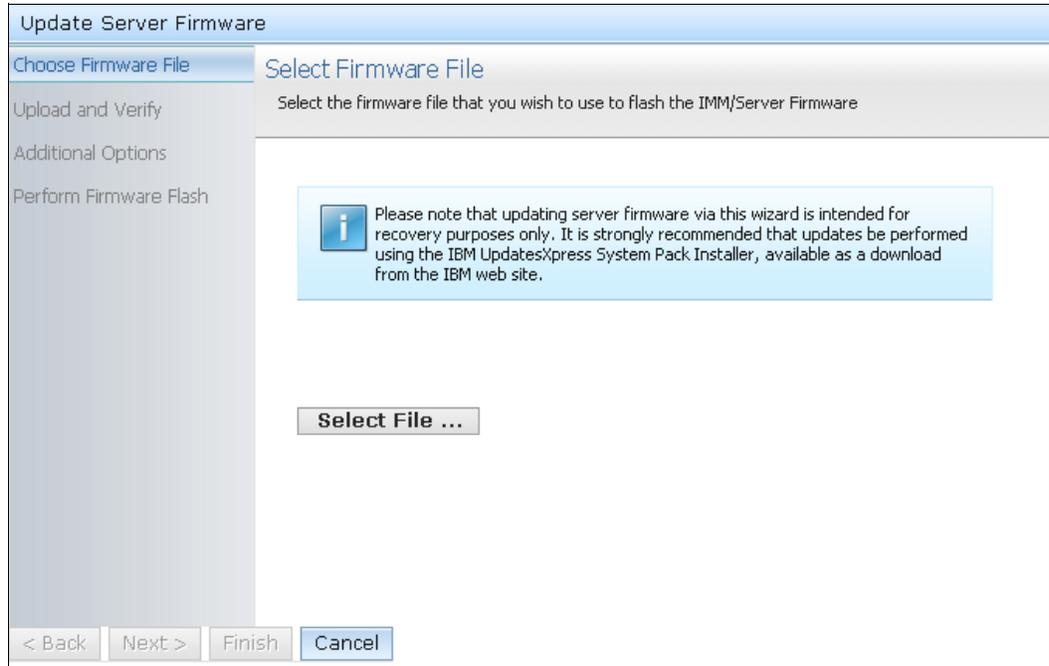


Figure 4-12 IMM Update Server Firmware window

When the capability is needed to recover the firmware, the IMM and the UEFI firmware can be updated via the IMM2. It is imperative to follow all the update procedures and wait times between the updates, exactly. This step increases the chance of a recovery and decreases the chance of firmware failures that might need service or require additional downtime.

Before attempting any firmware update that use the update wizard in the IMM2, be sure to review the *Firmware Update Guides* at this website.

<http://ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5082923>

4.2 Supported operating systems

To see a list of the operating systems supported by the x3750 M4 (machine type specific), see the *Lenovo Operating System Interoperability Guide*:

<http://lenovopress.com/osig>

4.3 Installing an operating system

An OS can be installed on the x3750 M4 with either local or remote media. To install by using a locally attached optical drive, use the UltraSlim Enhanced SATA Multi-Burner optical drive (part number 44W3253).

In the absence of an optical drive or to be able to install media remotely, you can use the remote media feature of the IMM2 to mount the ISO image or optical drive on a management workstation. The ability to install from remote media is ideal for remote servers, OS installation software in ISO form, or deployment to multiple machines.

To boot a server with remote media, follow these steps to use the Remote Control and Remote Disk function of the IMM2. The Remote Control function not only allows control of the server at the OS level, but also at the beginning boot processes through power-on self-test (POST):

1. Log in to the IMM2 web interface with the predefined IP address and login credentials.
2. Click **Remote Control** on the main System Status page. Make sure that any pop-up blockers are turned off. Figure 4-13 shows the IMM2 remote control window.



Figure 4-13 IMM2 remote control

3. Next, choose to use either the ActiveX client or Java client to launch the remote console, then choose **Start Remote Control in Single-User Mode**. This selection is the only of the

two modes that supports a remote disk. Figure 4-14 shows the IMM2 remote control user modes.

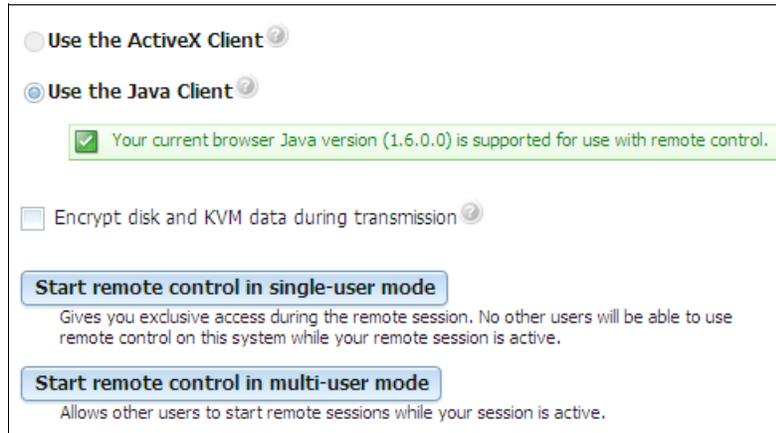


Figure 4-14 IMM2 remote control user modes

4. A new window is displayed that provides access to the Remote Disk and Remote Console functionality. The Remote Disk functionality is launched from the Remote Console window, by clicking **Tools** → **Launch Virtual Media**.
5. The Virtual Media Session window is displayed and senses the optical drive that is mapped in the workstation that is being used to access the IMM2. Either select the check box to map the drive found, or if the media is an image to be mounted, click **Add Image** to browse your workstation for the ISO or IMG file needed. Be sure to check the **Map** box next to the mounted image and click **Mount Selected** on the right side. Figure 4-15 shows the Virtual Media Session window.

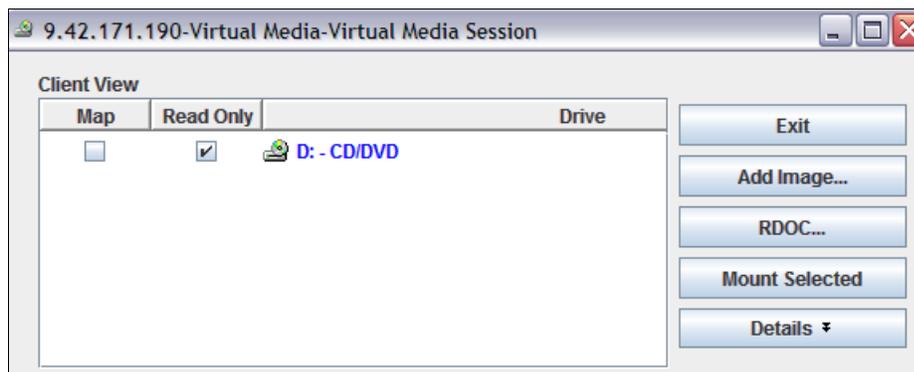


Figure 4-15 Virtual Media Session

6. With media mounted, power on the server and boot into the **F1 System Configuration** menu to review the Start Options to see that CD/DVD are in the boot order.

With the remote media mounted to the machine, it can boot to the machine to allow the OS installation to begin.

4.3.1 Use of ServerGuide

ServerGuide is a tool to assist with setup, configuration, and OS installations on System x servers. The guide can be downloaded from this website:

<http://ibm.com/support/entry/portal/docdisplay?lnocid=SERV-GUIDE>

ServerGuide makes OS installation easier by assisting with RAID configuration and by loading drivers that were not included in the OS install media. The latter is especially helpful when the storage or network adapter drivers are not part of the OS install media. The ability of the tool to load drivers makes it ideal when the OS that is being installed does not contain the driver for the storage adapter or network adapter, for example.

The ServerGuide can be used to assist with the installation of the following supported operating systems:

- ▶ Microsoft Windows Server 2008 R2: Minimum Service Pack 1
- ▶ Microsoft Windows Server 2008, HPC Edition
- ▶ Microsoft Windows Server 2008, Datacenter x64 Edition: Minimum Service Pack 2
- ▶ Microsoft Windows Server 2008, Web x64 Edition: Minimum Service Pack 2
- ▶ Microsoft Windows Server 2008, Enterprise x64 Edition: Minimum Service Pack 2
- ▶ Microsoft Windows Server 2008, Standard x64 Edition: Minimum Service Pack 2

When downloading the ServerGuide ISO file for use on the x3750 M4, make sure to choose the version for 64-bit installations. This ISO image supports both previous and native UEFI installations. The included readme file contains further details about instructions and uses of the ServerGuide.

The ServerGuide Scripting Tool is also available, which allows for easier deployment of OS installations to multiple machines. To access the ServerGuide download matrix for the scripting tool and for more information about its uses, see this website:

<http://ibm.com/support/entry/portal/docdisplay?lnocid=SERV-TOOLKIT>

4.3.2 Installation instructions

Obtain installation instructions for the System x3750 M4 by accessing the Lenovo Support portal:

<http://www.lenovo.com/support>

Or, view Google Search results by accessing this website:

<https://google.com/search?q=allintitle:+x3750+installing+site%3Aibm.com%2Fsupport>

4.3.3 Installing Windows

Here are hints and tips for installing Microsoft Windows Server 2008 on the x3750 M4:

- ▶ Windows Server 2008 is a UEFI-aware OS and can be installed to local disks in the UEFI mode.
- ▶ The OS can be installed with local media or from remote mounted media.
- ▶ The ServerGuide can be used to assist in the installation of Windows by providing the necessary drivers.

- ▶ The Windows Server 2008 install media does not contain any of the following required drivers:

- Emulex BE3 Ethernet controller
- ServeRAID M5100 series SAS controller (LSI SAS2208 chipset)
- Intel chipset

Drivers for these devices are needed at the time of the installation and are available on Fix Central.

- ▶ After the Microsoft Windows Server 2008 OS is installed, the UpdateXpress System Pack Installer can be run to install or update drivers and firmware.

4.3.4 Installing VMware

Here are hints and tips for installing VMware vSphere, ESXi, or ESX on the x3750 M4:

- ▶ To support ESX and ESXi 4.1, the minimum level is Update 2. The VMware installation media for 4.1 does not contain the driver for the Emulex BE3 Ethernet controller.
- ▶ The VMware installation media for vSphere 5 does not contain the driver for the onboard ServeRAID M5100 series SAS controller (LSI SAS2208 chipset).
- ▶ The Lenovo customized media for VMware installation does contain the drivers for key components. Lenovo customized media is only available for x3750 M4 machine type 8753. To download VMware vSphere Hypervisor (ESXi) with Lenovo Customization, see these web page:
 - VMware ESXi custom images for 6.0
https://my.vmware.com/web/vmware/info/slug/datacenter_cloud_infrastructure/vmware_vsphere/6_0#custom_iso
 - VMware ESXi custom images for 5.5
https://my.vmware.com/web/vmware/info/slug/datacenter_cloud_infrastructure/vmware_vsphere/5_5#custom_iso
 - VMware ESXi custom images for 5.1
https://my.vmware.com/web/vmware/info/slug/datacenter_cloud_infrastructure/vmware_vsphere/5_1#custom_iso

4.3.5 Installing Linux

This subsection describes hints and tips for installing RHEL or SLES on the x3750 M4.

Consider the following tips when installing Linux on the System x3750 M4:

- ▶ There are multiple versions of Red Hat Enterprise Linux and SUSE Linux Enterprise Server supported on the x3750 M4.
- ▶ If installing RHEL 5, the minimum level is Update 8.
- ▶ If installing RHEL 6, the minimum level is Update 2.
- ▶ If installing SLES 10, the minimum level is Service Pack 4 and offers only a base system certification. SLES 10 does not support the Romley chipset EP4S features.
- ▶ If installing SLES 11, it must be at least Service Pack 2.
- ▶ The Linux Enterprise OSs can be installed by using local media or remote mounted media.

4.4 Post-installation management tools

Now that your x3750 M4 is running and has a server OS installed, you can consider post-installation management planning and procedures. Lenovo provides many tools for managing the x3750 M4 for one or multiple servers. The tools described in this section are suitable for both local and remote management of the servers in production.

4.4.1 IMM2 email notifications

The IMM2 is useful during all stages of the installation and production management of the x3750 M4 server.

IMM2, as the service processor for the server, constantly monitors the machine for environmental changes, status, and failure of the hardware components within the machine. The IMM2 also maintains a system log to allow for reporting and troubleshooting issues.

IMM2 can be configured to send alerts based on the severity of the IMM messages logged. Therefore, events in the IMM logs are not missed just because no one logged in to check the System Status. These alerts can be configured to be sent to an email address or the IP address of a syslog collector.

Follow these steps to configure the IMM2 to send alerts:

1. Log in to the IMM2 web interface by using the assigned IP address.
2. Go to **Events** → **Event Recipients** to add and modify email and syslog notifications.

Figure 4-16 shows the IMM2 Event Recipients function.



Figure 4-16 IMM2 Event Recipients

- Use the Create tab to choose to create an email Notification or Syslog Notification to define up to 12 unique recipients for alerts.

Options for configuring email notifications include defining the recipients name, email address, and the severity of the alerts sent. Figure 4-17 shows the Create email Notification window.

Create E-Mail Notification

Descriptive name:

E-Mail address:

Events to receive:

<input type="checkbox"/> Enable Critical	<input type="checkbox"/> Enable Attention	<input type="checkbox"/> Enable System
<input type="checkbox"/> Critical Temperature Threshold Exceeded	<input type="checkbox"/> Power redundancy warning	<input type="checkbox"/> Successful Remote Login
<input type="checkbox"/> Critical Voltage Threshold Exceeded	<input type="checkbox"/> Warning Temperature Threshold Exceeded	<input type="checkbox"/> Operating System Timeout
<input type="checkbox"/> Critical Power Failure	<input type="checkbox"/> Warning Voltage Threshold Exceeded	<input type="checkbox"/> All other informational/system events
<input type="checkbox"/> Hard Disk Drive Failure	<input type="checkbox"/> Warning Power Threshold Exceeded	<input type="checkbox"/> System Power On/Off
<input type="checkbox"/> Fan Failure	<input type="checkbox"/> Non-critical Fan events	<input type="checkbox"/> Operating System boot failure
<input type="checkbox"/> CPU Failure	<input type="checkbox"/> CPU in degraded state	<input type="checkbox"/> Operating System Loader watchdog timeout
<input type="checkbox"/> Memory Failure	<input type="checkbox"/> Memory Warning	<input type="checkbox"/> Predicted failure (PFA)
<input type="checkbox"/> Hardware Incompatibility	<input type="checkbox"/> All other attention events	<input type="checkbox"/> Event log 75% full
<input type="checkbox"/> Power redundancy failure		<input type="checkbox"/> Network change
<input type="checkbox"/> All other critical events		

Include the event log contents in the e-mail body

Disable this notification

▶ Advanced Settings

OK Cancel

Figure 4-17 Create email Notification window

Options for configuring Syslog Notifications include defining a descriptive name, a Host name, or IP address of the syslog collector, and the severity of the alerts sent.

Figure 4-18 shows the Create SysLog Notification window.

Create SysLog Notification

Descriptive name:

Host name or IP address of the SysLog collector: Port:

Events to receive:

<input type="checkbox"/> Enable Critical	<input type="checkbox"/> Enable Attention	<input type="checkbox"/> Enable System
<input type="checkbox"/> Critical Temperature Threshold Exceeded	<input type="checkbox"/> Power redundancy warning	<input type="checkbox"/> Successful Remote Login
<input type="checkbox"/> Critical Voltage Threshold Exceeded	<input type="checkbox"/> Warning Temperature Threshold Exceeded	<input type="checkbox"/> Operating System Timeout
<input type="checkbox"/> Critical Power Failure	<input type="checkbox"/> Warning Voltage Threshold Exceeded	<input type="checkbox"/> All other informational/system events
<input type="checkbox"/> Hard Disk Drive Failure	<input type="checkbox"/> Warning Power Threshold Exceeded	<input type="checkbox"/> System Power On/Off
<input type="checkbox"/> Fan Failure	<input type="checkbox"/> Non-critical Fan events	<input type="checkbox"/> Operating System boot failure
<input type="checkbox"/> CPU Failure	<input type="checkbox"/> CPU in degraded state	<input type="checkbox"/> Operating System Loader watchdog timeout
<input type="checkbox"/> Memory Failure	<input type="checkbox"/> Memory Warning	<input type="checkbox"/> Predicted failure (PFA)
<input type="checkbox"/> Hardware Incompatibility	<input type="checkbox"/> All other attention events	<input type="checkbox"/> Event log 75% full
<input type="checkbox"/> Power redundancy failure		<input type="checkbox"/> Network change
<input type="checkbox"/> All other critical events		

Disable this notification

▶ Advanced Settings

OK Cancel

Figure 4-18 Create SysLog Notification window

4. After the parameters are defined to create notifications, click **OK** to add the recipients to the list of Event Recipients in the IMM2.

Abbreviations and acronyms

AC	alternating current	DSA	Dynamic System Analysis
AES	Advanced Encryption Standard	DVD	Digital Video Disc
AFT	adapter fault tolerance	ECC	error checking and correcting
ALB	adaptive load balancing	ETS	Enhanced Technical Support
AMM	Advanced Management Module	FC	Fibre Channel
ASHRAE	American Society of Heating, Refrigerating and Air Conditioning Engineers	FCC	Federal Communications Commission
ASU	Advanced Settings Utility	FDX	full duplex
ATS	Advanced Technical Support	FTP	File Transfer Protocol
AVX	Advanced Vector Extensions	GB	gigabyte
BBU	battery backup unit	GUI	graphical user interface
BGI	background initialization	HBA	host bus adapter
BIOS	basic input output system	HDD	hard disk drive
BIST	Built-in Self Test	HE	high efficiency
BSMI	Bureau of Standards, Metrology and Inspection	HPC	high performance computing
BTU	British Thermal Unit	HS	hot swap
CB	Certification Body	HTTP	Hypertext Transfer Protocol
CC	consistency check	HW	hardware
CD	compact disk	I/O	input/output
CD-ROM	compact disc read only memory	IBM	International Business Machines
CE	Conformité Européene	IDE	integrated drive electronics
CIM	Common Information Model	IEC	International Electrotechnical Commission
CISPR	International Special Committee on Radio Interference	IEEE	Institute of Electrical and Electronics Engineers
CLI	command-line interface	IMM	integrated management module
CMA	Cable Management Arm	IOPS	I/O operations per second
CMOS	complementary metal oxide semiconductor	IP	Internet Protocol
CNA	Converged Network Adapter	IPMI	Intelligent Platform Management Interface
CPU	central processing unit	ISO	International Organization for Standards
CSA	Canadian Standards Association	IT	information technology
CTO	configure-to-order	ITSO	International Technical Support Organization
DC	domain controller	KVM	keyboard video mouse
DCBX	Data Center Bridging eXchange	LAN	local area network
DHCP	Dynamic Host Configuration Protocol	LDAP	Lightweight Directory Access Protocol
DIMM	dual inline memory module	LED	light emitting diode
DL	dual-layer	LOM	LAN on motherboard
DNS	Domain Name System	LP	low profile
DRAM	dynamic random access memory		

LR-DIMM	load-reduced DIMM	SNMP	Simple Network Management Protocol
LRO	Large receive offload	SS	simple swap
LSO	Large send offload	SSD	solid state drive
MAC	media access control	SSH	Secure Shell
MB	megabyte	TB	terabyte
MLC	multi-level cell	TCP	Transmission Control Protocol
MS	Microsoft	TDP	thermal design power
MSI-X	Message Signal Interrupt	TFTP	Trivial File Transfer Protocol
MTU	maximum transmission unit	TPM	Trusted Platform Module
NBD	next business day	TTY	teletypewriter
NFS	network file system	TUV-GS	Technischer Überwachungs-Verein Geprüfte Sicherheit (TUV tested safety)
NIC	network interface card	UDIMM	unbuffered DIMM
NL	nearline	UDP	user datagram protocol
NVRAM	non-volatile random access memory	UEFI	Unified Extensible Firmware Interface
OCE	Online Capacity Expansion	URL	Uniform Resource Locator
OS	operating system	USB	universal serial bus
PCI	Peripheral Component Interconnect	UXSPI	UpdateXpress System Packs Installer
PFA	Predictive Failure Analysis	VCCI	Voluntary Control Council for Interference
PFC	Priority-based Flow Control	VFA	Virtual Fabric Adapter
POST	power-on self test	VLAN	virtual LAN
PXE	Preboot eXecution Environment	VPD	vital product data
QPI	QuickPath Interconnect	WOL	Wake on LAN
RAID	redundant array of independent disks		
RDIMM	registered DIMM		
RHEL	Red Hat Enterprise Linux		
RLM	RAID Level Migration		
ROC	RAID-on-card		
ROM	read-only memory		
RSS	Receive-side scaling		
SAN	storage area network		
SAS	Serial Attached SCSI		
SATA	Serial ATA		
SCSI	Small Computer System Interface		
SD	sales and distribution		
SED	self-encrypting drive		
SFF	Small Form Factor		
SFP	small form-factor pluggable		
SFT	switch fault tolerance		
SLC	Single Level Cell		
SLES	SUSE Linux Enterprise Server		
SMART	Self-Monitoring Analysis and Report Technology		

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this paper.

Lenovo Press publications

The following Lenovo Press publications provide additional information about the topic in this document.

- ▶ *System x3750 M4*, TIPS0881
- ▶ *Integrated Management Module (IMM) Support on System x and BladeCenter Servers*, TIPS0849
- ▶ *ServeRAID M5110 and M5110e SAS/SATA Controllers for System x*, TIPS0857
- ▶ *ServeRAID M5120 SAS/SATA Controller for System x*, TIPS0858

Other publications and online resources

These publications and websites are also relevant as further information sources:

These publications are also relevant as further information sources:

- ▶ *System x3750 M4 Installation and Service Guide*
<http://ibm.com/support/entry/portal/docdisplay?lnocid=MIGR-5090828>
- ▶ *Integrated Management Module II User's Guide*
<http://ibm.com/support/entry/portal/docdisplay?lnocid=MIGR-5086346>
- ▶ *ServeRAID M5110 Quick Installation Guide and User's Guide*
<http://ibm.com/support/entry/portal/docdisplay?lnocid=MIGR-5089524>
- ▶ *ServeRAID M5100 Series Quick Installation Guides for hardware upgrades*
<http://ibm.com/support/entry/portal/docdisplay?lnocid=MIGR-5089482>
- ▶ *ServeRAID-M Software User's Guide*
<http://ibm.com/support/entry/portal/docdisplay?lnocid=MIGR-5086126>
- ▶ *Features on Demand User's Guide:*
<http://ibm.com/support/entry/portal/docdisplay?lnocid=migr-5089568>
- ▶ *Advanced Settings Utility (ASU) User's Guide*
<http://ibm.com/support/entry/portal/docdisplay?lnocid=MIGR-5085890>
- ▶ *System x Update Tools Transition*
<http://ibm.com/support/entry/portal/docdisplay?lnocid=MIGR-5082923>
- ▶ *Firmware Update Best Practices (for IMMv2 Systems)*
<http://ibm.com/support/entry/portal/docdisplay?lnocid=MIGR-5082923>

- ▶ *UpdateXpress System Pack Installer (UXSPI) User's Guide*

<http://ibm.com/support/entry/portal/docdisplay?lnocid=MIGR-5085892>