

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

# Lenovo System x3850 X6 and x3950 X6 Planning and Implementation Guide

Last Update: October 2018

---

**Covers the sixth generation  
Enterprise X-Architecture servers**

---

**Provides technical information about  
all server features**

---

**Explains what planning activities  
you need to do**

---

**Describes how to manage the server  
with Lenovo XClarity Administrator**

**David Watts**

**Rani Doughty**

**Ilya Solovyev**







# **Lenovo System x3850 X6 and x3950 X6 Planning and Implementation Guide**

October 2018

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

**Last update on October 2018**

This edition applies to Lenovo System x3850 X6 and x3950 X6, machine type 6241, with Intel Xeon Processor E7-4800/8800 v2, v3 and v4 processors.

© Copyright Lenovo 2018. All rights reserved.

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

# Contents

<b>Notices</b> .....	ix
Trademarks .....	x
<b>Summary of changes</b> .....	xi
October 2018 .....	xi
February 2017 .....	xi
October 2016 .....	xii
August 2016 .....	xii
June 2016 .....	xii
September 2015 .....	xiii
September 2014 .....	xiii
June 2014 .....	xiii
<b>Preface</b> .....	xv
The team who wrote this book .....	xv
Comments welcome .....	xvi
<b>Chapter 1. Introduction</b> .....	1
1.1 Target workloads .....	2
1.1.1 Databases .....	2
1.1.2 Business analytics .....	2
1.1.3 Virtualization .....	3
1.1.4 Enterprise applications: ERP and CRM .....	4
1.2 Key features .....	4
1.3 Positioning .....	5
1.4 Storage versus in-memory data .....	7
1.5 Flash storage .....	8
1.6 Energy efficiency .....	8
1.7 Services offerings .....	9
1.8 About this book .....	10
<b>Chapter 2. New technology in X6 servers</b> .....	11
2.1 Modular design .....	12
2.1.1 Compute Books .....	14
2.1.2 Storage Book .....	17
2.1.3 Primary I/O Book .....	17
2.1.4 More I/O Books .....	18
2.1.5 Power supplies .....	19
2.2 System architecture .....	20
2.2.1 x3850 X6 .....	20
2.2.2 x3950 X6 .....	22
2.3 Processors .....	24
2.3.1 Intel Xeon processor E7-4800/8800 v4 product family .....	25
2.3.2 Intel Xeon processor E7-4800/8800 v3 product family .....	26
2.3.3 Intel Xeon processor E7-4800/8800 v2 product family .....	27
2.3.4 Intel Xeon E7 features .....	27
2.3.5 Compute Books .....	31
2.4 Memory .....	33
2.4.1 Operational modes .....	34

2.4.2	Memory mirroring and rank sparing	38
2.4.3	Chipkill	42
2.4.4	Redundant bit steering	42
2.4.5	Advanced Page Retire	42
2.5	PCIe 3.0	43
2.6	Internal storage	43
2.6.1	Storage Book	44
2.6.2	SAS/SATA SSD technology	44
2.6.3	NVMe SSD technology	45
2.6.4	Flash Storage Adapters	47
2.7	Unified Extensible Firmware Interface	49
2.8	Integrated Management Module	49
2.9	Scalability	50
<b>Chapter 3. Product information</b>		<b>53</b>
3.1	Product features	54
3.1.1	Fast application performance	54
3.1.2	Agile system design	55
3.1.3	Resilient platform	55
3.2	Specifications	57
3.3	Standard models of X6 servers	60
3.3.1	x3850 X6 models	60
3.3.2	x3950 X6 models	62
3.4	SAP HANA Workload Optimized Solution models	64
3.4.1	x3850 X6 HANA models	64
3.4.2	x3950 X6 HANA models	66
3.5	Physical design	68
3.6	Ports and controls	70
3.6.1	Front operator panel	70
3.6.2	LCD system information panel	72
3.6.3	Rear ports and LEDs	74
3.7	Compute Books	75
3.7.1	Compute Book design	76
3.7.2	Compute Book population order	78
3.7.3	Upgrading Compute Books	79
3.8	Processor options	80
3.9	Memory	84
3.9.1	Memory options	85
3.9.2	Memory population order	89
3.10	Memory-channel storage	92
3.11	Storage subsystem	94
3.11.1	Storage Book	94
3.11.2	Backplanes	96
3.11.3	RAID controllers	97
3.11.4	Disk drive options	99
3.11.5	PCIe Flash Storage Adapters	99
3.12	I/O subsystem	100
3.13	Primary I/O Book	102
3.14	Half-length I/O Books and Full-length I/O Books	104
3.14.1	Half-length I/O Book	106
3.14.2	Full-length I/O Book	106
3.15	Hot-swap adapter support	108
3.16	Network adapters	109

3.17	SAS adapters for external storage	112
3.18	Fibre Channel host bus adapters	114
3.19	GPU adapters and co-processors	114
3.20	Partitioning	116
3.21	Standard onboard features	116
3.21.1	Integrated Management Module II	117
3.21.2	Unified Extensible Firmware Interface	117
3.21.3	Integrated Trusted Platform Module	117
3.21.4	Light path diagnostics	118
3.22	Integrated virtualization	120
3.23	Hot-swap capabilities	121
3.24	Power subsystem	122
3.25	Fans and cooling	124
3.26	Upgrading to an 8-socket X6 server	125
<b>Chapter 4. Infrastructure planning</b>		<b>129</b>
4.1	Physical and electrical specifications	130
4.2	Rack selection and rack options	131
4.3	Floor clearance	134
4.4	Using Rear Door Heat eXchanger	135
4.5	Power guidelines	136
4.5.1	Considerations	136
4.5.2	Power supply redundancy	137
4.5.3	Rules for achieving redundancy	139
4.5.4	Power supply installation order	140
4.5.5	Power policy	141
4.5.6	More power settings in the IMM2	143
4.5.7	Examples of power connections	145
4.6	Cooling considerations	145
4.7	Uninterruptible power supply units	146
4.8	PDU and line cord selection	147
4.8.1	Server-to-PDU power cord options	147
4.8.2	PDU and line cord options	148
<b>Chapter 5. Preparing the hardware</b>		<b>155</b>
5.1	Configuring the IMM2 settings	156
5.1.1	IMM2 virtual presence	156
5.1.2	IMM2 network access	157
5.1.3	Configuring the IMM2 network interface	158
5.1.4	IMM2 dedicated versus shared ML2 Ethernet port	159
5.1.5	x3950 X6 IMM2 communication	160
5.1.6	Troubleshooting IMM2 communications	161
5.1.7	IMM2 functions to diagnose and manage the server	161
5.2	Unified Extensible Firmware Interface settings	165
5.2.1	Operating modes	168
5.3	UEFI common settings	169
5.3.1	System power settings	170
5.3.2	Processor settings	171
5.3.3	Memory settings	173
5.3.4	UEFI settings for different workload types	174
5.3.5	ServeRAID M5210 RAID controller configuration	176
5.3.6	Speeding up the boot process	181
5.4	Use of NVMe drives and adapters	184

5.4.1	NVMe drive placement	184
5.4.2	NVMe PCIe SSD adapter placement	188
5.4.3	Using NVMe drives with Linux	188
5.4.4	Using NVMe drives with Microsoft Windows Server 2012 R2	192
5.4.5	Using NVMe drives with VMware ESXi server	196
5.4.6	Ongoing NVMe drive management	198
5.5	PCIe adapter placement advice	198
5.6	Hot-swap procedures	201
5.6.1	Hot-swapping a power supply	201
5.6.2	Hot-swapping an I/O Book	202
5.7	Partitioning the x3950 X6	205
5.7.1	Partitioning an x3950 X6 via the IMM2 web interface	206
5.8	Updating firmware	211
5.8.1	Firmware tools	211
5.8.2	Updating firmware	212
5.9	Troubleshooting	221
5.9.1	Integrated Management Module	221
5.9.2	LCD system information panel	222
5.9.3	System event log	222
5.9.4	POST event log	223
5.9.5	Installation and Service Guide	223
<b>Chapter 6. Operating system installation</b>		<b>225</b>
6.1	Installing an OS without a local optical drive	226
6.1.1	Integrated Management Module	226
6.1.2	Local USB port	229
6.1.3	Preboot eXecution Environment	229
6.1.4	Lenovo XClarity Administrator	229
6.2	ServerGuide	233
6.3	ServerGuide Scripting Toolkit	237
6.4	Use of embedded VMware ESXi	239
6.5	Booting from SAN	240
<b>Chapter 7. Managing X6 hardware</b>		<b>243</b>
7.1	Introduction	244
7.2	Integrated Management Module II	244
7.2.1	Configuring IMM2 for out-of-band-management	245
7.2.2	Configuring IMM2 in-band configuration	248
7.3	Remote control	248
7.3.1	Accessing the remote control feature in the IMM2	248
7.4	Lenovo XClarity Administrator	249
7.4.1	X6 considerations	250
7.4.2	Discovering the IMM2 of an x3850 X6	250
7.4.3	Remote control	253
7.4.4	Hardware monitoring	254
7.4.5	Firmware updates	256
7.4.6	Operating system deployment	257
7.4.7	Call home service	258
7.5	Lenovo XClarity integrators	258
7.6	Lenovo XClarity Energy Manager	265
7.6.1	Lenovo XClarity Energy Manager licensing	265
7.6.2	System requirements	266
7.6.3	Download and install Lenovo XClarity Energy Manager	267



7.6.4	Launching Lenovo XClarity Energy Manager . . . . .	271
7.6.5	Setting up Lenovo XClarity Energy Manager . . . . .	274
7.6.6	Energy optimization with Energy Manager . . . . .	284
7.7	Upward Integration Module for Zenoss . . . . .	288
7.8	Advanced Settings Utility . . . . .	289
7.8.1	Using ASU to configure settings in IMM2-based servers . . . . .	289
7.8.2	Command examples . . . . .	290
7.9	MegaRAID Storage Manager . . . . .	292
7.9.1	MegaRAID Storage Manager installation . . . . .	292
7.9.2	Drive states . . . . .	292
7.9.3	Virtual drive states . . . . .	293
7.9.4	StorCLI utility for storage management . . . . .	293
7.10	Managing NVMe drives and software RAID recovery . . . . .	295
7.10.1	Software RAID initialization in Linux . . . . .	295
7.10.2	NVMe drive hot-replacement in Linux . . . . .	299
7.10.3	NVMe drive hot-plug and software RAID recovery in Linux . . . . .	301
7.10.4	NVMe drive hot-replacement and software RAID recovery in Windows . . . . .	303
7.11	Serial over LAN . . . . .	312
7.11.1	Enabling SoL in UEFI . . . . .	313
7.11.2	Enabling SoL in the operating system . . . . .	313
7.11.3	Starting an SoL connection . . . . .	317
	<b>Abbreviations and acronyms . . . . .</b>	<b>319</b>
	<b>Related publications . . . . .</b>	<b>323</b>
	Lenovo Press publications . . . . .	323
	Other publications . . . . .	324
	Online resources . . . . .	324



# 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®	ThinkServer®
BladeCenter®	Lenovo XClarity™	ToolsCenter™
Bootable Media Creator™	MAX5™	TopSeller™
Dynamic System Analysis™	Lenovo (logo)®	TruDDR4™
eX5™	ServeRAID™	UltraNav®
eXFlash™	ServerGuide™	UpdateXpress System Packs™
Flex System™	Storage Configuration Manager™	vNIC™
Intelligent Cluster™	System x®	X5™

The following terms are trademarks of other companies:

Intel, Intel Core, Intel Xeon Phi, 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.

BitLocker, Excel, Hyper-V, Internet Explorer, Microsoft, 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.

# Summary of changes

This section describes the changes made in this update and in previous updates. These updates might also include minor corrections and editorial changes that are not identified.

## October 2018

### Additions:

- ▶ Added the capacities of the blank USB keys, Table 3-25 on page 120

### Updates:

- ▶ Models with E7 v3 processors are now withdrawn from marketing, however compute books are still available for field upgrades.
- ▶ Updated the list of supported operating systems
- ▶ Indicated options that are now withdrawn from marketing

### Corrections:

- ▶ Corrected which bays compute books should be installed in for a 6-processor configuration, 3.7.2, "Compute Book population order" on page 78

## February 2017

- ▶ Added new compute book option, 00YG935, based on the Intel Xeon Processor E7-8894 v4, Table 3-6 on page 81
- ▶ Added S3520 and PM863a solid-state drives
- ▶ Added 32GB LRDIMM memory option, 00KH391, for certain E7 v3 processors Table 3-7 on page 85
- ▶ Compute Books with E7 v2 processors now all withdrawn from marketing
- ▶ Removed withdrawn options:
  - Emulex VFA5 ML2 Dual Port 10GbE SFP+ Adapter, 00D1996
  - 300GB 15K 12Gbps SAS 2.5" G3HS 512e HDD, 00NA221
  - 500GB 7.2K 6Gbps NL SAS 2.5" G3HS HDD, 00AJ121
  - 600GB 15K 12Gbps SAS 2.5" G3HS 512e HDD, 00NA231
  - 500GB 7.2K 6Gbps NL SATA 2.5" G3HS HDD, 00AJ136
  - 146GB 15K 6Gbps SAS 2.5" G3HS HDD, 00AJ111
  - 1.2TB 10K 12Gbps SAS 2.5" G3HS 512e HDD, 00NA261
  - USB Memory Key for VMware ESXi 5.1 Update 2, 00ML233
  - USB Memory Key for VMware ESXi 5.1 U1, 41Y8382
  - NVIDIA Quadro M6000, 00KH377
  - NVIDIA Grid K1 PCIe x16 for System x3850/x3950 X6, 00FP671
  - NVIDIA Grid K2 Actively Cooled PCIe x16 for System x3850/x3950 X6, 00FP674
  - NVIDIA Tesla M60 GPU, PCIe (active), 00YL377

## October 2016

- ▶ Windows Server 2016 is now supported
- ▶ Added Intel S3610 SSDs
- ▶ Added Mellanox ConnectX-4 Lx 2x25GbE SFP28 Adapter, Table 3-20 on page 110
- ▶ Added NVIDIA Quadro M6000 24GB GPU, PCIe (active), Table 3-24 on page 115
- ▶ Corrected the part number for the NVIDIA Grid K1 PCIe x16, Table 3-24 on page 115
- ▶ The following recently added options are not supported with E7 v2 compute books:
  - Network adapters: Intel X710-DA2, Intel X550-T2, and Mellanox ConnectX-4
  - FC host bus adapters: QLogic 16Gb FC and Emulex 16Gb FC
  - NVIDIA GPUs: Tesla M60 and Quadro M5000
  - USB memory keys preloaded with VMware ESXi 5.5 U3B and 6.0 U2
- ▶ Corrected Table 5-3 on page 185 and Figure 5-40 on page 187 regarding the connections between the NVMe PCIe Extender Adapter and the NVMe drive backplane.
- ▶ Corrected the numbering of the x3950 X6 slots, Figure 5-63 on page 199
- ▶ Withdrawn from marketing:
  - Models with E7 v2 processors (the compute books are still available as options for existing customers)
  - 600GB 10K 12Gbps SAS 2.5" G3HS 512e SED, 00NA291
  - Solarflare SFN5162F 2x10GbE SFP+ Performant Adapter, 47C9952
  - Intel Xeon Phi 3120A PCIe x16, 00FP670

## August 2016

- ▶ Added additional options:
  - Network adapters: Intel X710-DA2, Intel X550-T2, and Mellanox ConnectX-4
  - FC host bus adapters: QLogic 16Gb FC and Emulex 16Gb FC
  - NVIDIA GPUs: Tesla M60 and Quadro M5000
  - USB memory keys preloaded with VMware ESXi 5.5 U3B and 6.0 U2
- ▶ Grammar and style corrections

## June 2016

### New information

- ▶ New models based on the Intel Xeon E7 v4 processors
- ▶ New SAP HANA models
- ▶ New compute book options (E7-4800 v4 and E7-8800 v4 processors)
- ▶ New TruDDR4™ memory options
- ▶ New adapter and drive options
- ▶ Integrated Management Module is now IMM 2.1
- ▶ New PDU options
- ▶ Added console keyboard options
- ▶ New Instructions on how to perform hot-remove and hot-add operations for NVMe drives
- ▶ Lenovo® XClarity Energy Manager
- ▶ Updated XClarity Integrator part number options

#### Updated information

- ▶ eXFlash™ DIMMs now withdrawn from marketing
- ▶ io3 Enterprise Value Flash Storage Adapters now withdrawn from marketing
- ▶ Additional information about the included rack rail kit
- ▶ New IMM and XClarity screenshots
- ▶ New links to Lenovo support pages

## September 2015

#### New information:

- ▶ Models of the x3850 X6 and x3950 X6, machine type 6241
- ▶ Compute Books that are based on the Intel Xeon Processor E7-4800/8800 v3 family
- ▶ Compute Books with DDR3 and TruDDR4 memory
- ▶ PCIe NVMe solid-state drive technology and drive options
- ▶ Support for more PCIe SSD adapters
- ▶ Support for more network adapters
- ▶ Support for more HDDs and SSDs
- ▶ Information about upgrading from E7 v2-based Compute Books to E7 v3-based
- ▶ Use of Lenovo XClarity™ Administrator

## September 2014

#### Changed information:

- ▶ All processors support eXFlash DIMMs
- ▶ The RAID 1 feature of eXFlash DIMMs is currently not supported

## June 2014

#### New information:

- ▶ NVIDIA GPUs support a maximum of 1 TB of system memory, page 115
- ▶ Information about the cable management kit shipped with the server, page 60
- ▶ Added Intel I350 Ethernet adapters

#### Changed information:

- ▶ Corrected the depth dimensions, page 60 and page 130
- ▶ Certain processors do not support eXFlash DIMMs
- ▶ The eXFlash DIMM driver does not support RAID
- ▶ VMware vSphere 5.1 supports a maximum of 160 concurrent threads





# Preface

The increasing demand for cloud computing and business analytical workloads by enterprises to meet business needs drives innovation to find new ways to build informational systems. Clients are looking for cost-optimized fit-for-purpose IT solutions that manage large amounts of data, easily scale performance, and provide reliable real-time access to actionable information.

This book describes the four-socket Lenovo System x3850 X6 and eight-socket Lenovo System x3950 X6. These servers provide the computing power to drive mission-critical scalable databases, business analytics, virtualization, enterprise applications, and cloud applications.

This Lenovo Press book covers product information as well as planning and implementation information. In the first few chapters, we provide detailed technical information about the X6 servers. This information is most useful in designing, configuring, and planning to order a server solution. In the later chapters of the book, we provide detailed configuration and setup information to get your server operational.

This book is for clients, Lenovo Business Partners, and Lenovo employees that want to understand the features and capabilities of the X6 portfolio of servers and want to learn how to install and configure the servers for use in production.

## The team who wrote this book

This document is produced by the following subject matter experts working in the Lenovo offices in Morrisville, NC, USA.



**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).



**Rani Doughty** is a hardware specialist and data center consultant with a background in System x, BladeCenter, and Flex System. She currently works with the worldwide Lenovo Data Center Services (DCS) team as a developer of the Power Configurator. She has 12 years of technical experience in the x86 field. She holds an MBA and Honors degree in IT from the University of Ballarat (Australia). She has written and presented worldwide extensively on IBM Systems Director, pre- and post-sale tools, and infrastructure planning.



**Ilya Solovyev** is a Technical Consultant in Lenovo Enterprise Solution Services, based in Moscow. He currently provides technical consulting services for System x, Flex System, and BladeCenter hardware and software. His areas of expertise also include Linux systems, and virtualization, cloud, and SAP HANA solutions. Ilya is a certified Red Hat Engineer and has a Bachelor degree in Mathematics.

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

Lenovo data center marketing:

- ▶ Mary Beth Daughtry
- ▶ John Encizo
- ▶ Kyle Hampton
- ▶ Tisha Jefferson
- ▶ Randy Lundin
- ▶ Tony Meli
- ▶ Mark T. Chapman

Lenovo data center development:

- ▶ Ron Birtles
- ▶ Mike French
- ▶ Clifton Kerr
- ▶ Sumeet Kochar
- ▶ Randy Kolvick
- ▶ Tim Schlude
- ▶ Mehul Shah

## Comments welcome

Your comments are important to us!

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

- ▶ Use the online feedback form found at the following web page:  
<http://lenovopress.com/sg248208>
- ▶ Send your comments in an email to:  
[comments@lenovopress.com](mailto:comments@lenovopress.com)

# Introduction

The X6 family of scalable rack servers consists of the following servers:

- ▶ Lenovo System x3850 x6, a four-socket 4U rack-mount server
- ▶ Lenovo System x3950 X6, an eight-socket 8U rack-mount server

These servers are the sixth generation of servers that are built upon Lenovo Enterprise X-Architecture. Enterprise X-Architecture is the culmination of generations of Lenovo technology and innovation that is derived from our experience in high-end enterprise servers.

The X6 servers deliver innovation with enhanced scalability, reliability, availability, and serviceability (RAS) features to enable optimal break-through performance that is ideal for mission-critical scalable databases, business analytics, virtualization, enterprise applications, and cloud applications.

The X6 generation servers pack numerous fault-tolerant and high-availability features into a high-density, rack-optimized, chassis-like package where all serviceable components are front and rear accessible. This design significantly reduces the space that is needed to support massive network computing operations and simplifies servicing.

These servers can be expanded on demand, offering flexible modular scalability in processing, I/O, and memory dimensions so that you can provision what you need now and expand the system to meet future requirements. X6 is fast, agile, and resilient.

This chapter includes the following topics:

- ▶ 1.1, “Target workloads” on page 2
- ▶ 1.2, “Key features” on page 4
- ▶ 1.3, “Positioning” on page 5
- ▶ 1.4, “Storage versus in-memory data” on page 7
- ▶ 1.5, “Flash storage” on page 8
- ▶ 1.6, “Energy efficiency” on page 8
- ▶ 1.7, “Services offerings” on page 9
- ▶ 1.8, “About this book” on page 10

## 1.1 Target workloads

The X6 servers introduce new levels of fault tolerance with advanced RAS features that are implemented in hardware and software. It also introduces simplified servicing and upgrades with a bookshelf concept and lid-less design, which results in dramatic improvements in response time with stretched memory speeds and innovative flash storage offerings. These servers also use proven technologies of the previous generations of Enterprise X-Architecture.

These servers provide users who are looking for the highest level of scalable performance, the maximum memory capacity, and the richest set of RAS features for maximum productivity. The servers are designed for mission-critical, scalable workloads, including large databases, and ERP/CRM systems to support online transaction processing, business analytics, virtualization, and enterprise applications.

This section describes how X6 technology helps to address challenges clients are facing in these mission-critical enterprise environments.

### 1.1.1 Databases

Leadership performance, scalability, and large memory support means that X6 systems can be highly used, which yields the best return for the following database applications:

- ▶ SAP Business Suite on X6
- ▶ Microsoft SQL Data Warehouse on X6
- ▶ SAP HANA on X6
- ▶ IBM DB2 BLU on X6

X6 is well-suited for Online transaction processing (OLTP) workloads. OLTP workloads are characterized by small, interactive transactions that generally require subsecond response times. For most OLTP systems, the processor, memory, and I/O subsystem in a server are well-balanced and are not considered performance bottlenecks.

The major source of performance issues in OLTP environments often is related to the storage I/O. The speed of traditional hard disk drive (HDD)-based storage systems does not match the processing capabilities of the server. As a result, a situation often occurs where a powerful processor sits idle, waiting for storage I/O requests to complete, which negatively affects the user and business productivity. This wait is not an issue with X6.

The OLTP workload optimization goal for X6 systems is to address storage I/O bottlenecks through technologies, such as a large capacity memory subsystem to enable in-memory data, and high-performance/low-latency storage subsystem that uses flash storage technologies. For more information, see 1.4, “Storage versus in-memory data” on page 7 and 1.5, “Flash storage” on page 8.

### 1.1.2 Business analytics

Data warehouses are commonly used with online analytical processing (OLAP) workloads in decision support systems, such as financial analysis. Unlike OLTP, where transactions often are relatively simple and feature small amounts of data, OLAP queries are more complex and process larger volumes of data.

For OLAP workloads, transactional delays can significantly increase business and financial risks. Usually, decision making is stalled or delayed because of lack of accurate, real-time operational data for analytics, which can mean missed opportunities.

These transactional delays primarily come from batch data loads and performance issues because of handling heavy complex queries and massive amounts of data (frequently referred to as *big data*) that use I/O resources. For OLAP workloads, a fast response time is critical to ensure that strategic business decisions can be made quickly in dynamic market conditions.

In general, clients might experience the following challenges with OLAP environments:

- ▶ Slow query execution and response times, which delay business decision making.
- ▶ Dramatic growth in data, which requires deeper analysis.

Lenovo X6 systems can help to make businesses more agile and analytics-driven by providing up-to-the-minute analytics that are based on real-time data. As with OLTP workloads, in-memory databases or flash storage are used for workload optimization (see 1.4, “Storage versus in-memory data” on page 7 and 1.5, “Flash storage” on page 8).

By using X6 technology, we help address challenges in OLAP environments in the following ways:

- ▶ Dramatically boosting the performance of OLAP workloads with a scale-up architecture.
- ▶ Significantly improving response time for better and timely decision making.

### 1.1.3 Virtualization

Virtualization commonly increases effectiveness in the use of resources and reduces capital expenses, software licensing fees, and operational and management costs.

The first wave of server consolidation focused on lightly loaded servers that easily tapped into a hypervisor’s ability to share processor and memory resources across applications. Hypervisors struggle to manage and share the heavy I/O loads that are typical of performance-intensive workloads. As a result, performance-intensive databases that are used for core enterprise workloads, such as customer relationship management (CRM), enterprise resource planning (ERP), and supply chain management (SCM), are left to run on physical, non-virtual servers.

The next wave of server virtualization with X6 expands the virtualization footprint to the workhorse applications of enterprise IT, namely those performance-intensive databases.

X6 makes virtualization of mission-critical, performance-intensive workloads possible in the following ways:

- ▶ X6 support for an integrated hypervisor: All x3850 X6 and x3950 X6 models support the addition of an internal USB key with VMware ESXi that is installed.
- ▶ Workload optimized models help you choose the best components for a specific workload.
- ▶ Processor support: The Intel Xeon Processor E7-4800 and E7-8800 series support Intel Virtualization Technology (Intel VT) Flex Priority and Intel VT Flex migration.
- ▶ Large VM support for enterprise applications:
  - Virtualized SAP HANA on X6
  - VMware vCloud Suite on X6
  - Microsoft Hyper-V Private Cloud on X6

## 1.1.4 Enterprise applications: ERP and CRM

Enterprise applications, such as ERP or CRM represent a mixed workload in which transaction processing and a certain level of real-time reporting are available. In a 2-tier implementation, database server and application modules are on the same server. The key performance metric is response time, as with OLTP and OLAP workloads.

X6 offerings provide low latency, extreme performance, and efficient transaction management to accommodate mixed workload requirements. X6 in-memory and flash storage offerings can help to deliver the following benefits for enterprise applications:

- ▶ Dramatically boost the performance of applications and lower cost per IOPS ratio without redesigning the application architecture.
- ▶ Increase user productivity with better response times, which improves business efficiency.
- ▶ Increase data availability by using advanced system-level high availability and reliability technologies, which reduces the number of solution components and shortens batch processing and backup times.
- ▶ Increase storage performance and capacity while decreasing power, cooling, and space requirements.

## 1.2 Key features

The Lenovo X6 servers are fast, agile, and resilient and make meeting the business needs of your enterprise easier. Consider the following points:

- ▶ Fast performance means immediate access to information.

This business-critical, enterprise-class server delivers fast performance with the Lenovo flash storage offerings for the X6 servers. Flash storage combines extreme IOPS performance and low response time for transactional database workloads. The flash technologies that are used in the X6 servers include PCIe NVMe drives and adapters (higher bandwidth and shorter latency), Flash Storage Adapters (high performance NAND technology), and SAS/SATA SSDs (capacity, plus 12Gb host connection for performance).

- ▶ Agile system design helps to reduce costs and provides simple access to server components.

X6 delivers a unique and adaptive modular design that enables you to grow on-demand with the bookshelf design. Scale from 2-way to 4-way to 8-way and grow your memory and I/O to meet your needs.

With X6 technology, you can realize the following benefits:

- Support multiple generations of Intel processor technology that allows the choice of deploying Intel Xeon Processor E7-4800/8800 v2, v3, and v4 processors.
  - Offer an easy upgrade path from E7 v2 processors to E7 v3 and E7 v4 processors.
  - Add I/O capability, such as extra network adapters or storage devices, while the server is still running, which enables upgrades with minimal application downtime.
  - Most components that are used in a four-socket x3850 X6 can be reused when you are upgrading to an eight-socket x3950 X6.
- ▶ A resilient platform maximizes application uptime and promotes easy integration in virtual environments.

This new server is designed to continue operating if there is a component failure and helps you to reduce planned and unplanned downtime.

The RAS features of the new X6 servers include the following capabilities:

- Predict failures before they occur  
Predictive Failure Analysis (PFA) allows the server to monitor the status of critical subsystems and notify the system administrator when components appear to be degrading. Because of this information, replacing failing parts in most cases can be performed as part of planned maintenance activity. This ability reduces the need for unscheduled outages and so your system continues to run.
- Find failed components fast  
Light path diagnostics allows systems engineers and administrators to easily and quickly diagnose hardware problems. The LCD display on the front of the server gives you more information about the problem than LEDs, so failures can now be evaluated in seconds and costly downtime can be reduced or avoided.
- Survive a processor failure  
The server recovers from a failed processor and restart automatically. Even if the primary processor (the one used for starting the operating system) fails, the X6 system can start from another processor by using redundant links to key resources.
- Survive memory failures  
The combination of Chipkill and Redundant Bit Steering (RBS, also known as Double Device Data Correction or DDDC) allows the server to tolerate two sequential DRAM memory chip failures without affecting overall system performance.
- Survive an adapter failure and replace it while the server is running  
The new servers feature up to six adapter slots that support hot-swapping. This configuration means that the I/O Books can be removed and any failed adapters can be replaced without any server downtime.
- Swap components easily with the server's lidless design  
There is no need to pull this server in or out of the rack to service it because all components can be accessed from the front or rear. This design allows for faster maintenance by simplifying service procedures. This concept is similar to the design that is used in BladeCenter and Flex System.

These built-in technologies drive the outstanding system availability and uninterrupted application performance that is needed to host mission-critical applications.

## 1.3 Positioning

The Lenovo System x3850 X6 and x3950 X6 servers are the next generation of X-Architecture following the highly successful eX5 server. X6 servers include various new features when compared to the previous generation of eX5, including support for more memory and I/O, and model-dependent support for v2, v3, and v4 Intel Xeon Processors in a modular design.

When compared to the 4-socket x3750 M4 server, the X6 servers fill the demand for enterprise workloads that require 4-socket and 8-socket performance, high availability, and advanced RAS features.

Table 1-1 shows a high-level comparison between the 4-socket x3750 M4, the eX5-based x3850 and x3950 X5™, and the X6-based x3850 and x3950 X6.

Table 1-1 Maximum configurations for the X6 systems

Maximum configurations		x3750 M4	x3850/x3950 X5	x3850/x3950 X6
Form factor	4 -socket	2U	4U	4U
	8-socket	Not available	8U	8U
Processors	1-node	4	4	4
	2-node	Not available	8	8
Cores	1-node	32	40	96
	2-node	Not available	80	192
Memory	1-node	48 DIMM slots <sup>a</sup>	64 DIMM slots <sup>a</sup>	96 DIMM slots <sup>a</sup>
	1-node with MAX5™	Not available	96 DIMM slots <sup>a</sup>	Not available
	2-node	Not available	128 DIMM slots <sup>a</sup>	192 DIMM slots <sup>a</sup>
2.5-inch drive bays	1-node	16	8	8
	2-node	Not available	16	16
1.8-inch SSDs drive bays	1-node	32	16	16
	2-node	Not available	32	32
2.5-inch PCIe NVMe SSDs	1-node	Not available	Not available	4
	2-node	Not available	Not available	8
Standard 1 Gb Ethernet interfaces	1-node	2	2 <sup>b</sup>	4 (Optional)
	2-node	Not available	4	8 (Optional)
Standard 10 Gb Ethernet interface	1-node	2 (Optional)	2	2 (Optional)
	2-node	Not available	4	4 (Optional)
USB ports	1-node	4 USB 2.0	8 USB 2.0	6 USB 2.0, 2 USB 3.0
	2-node	Not available	16 USB 2.0	12 USB 2.0, 4 USB 3.0
Power supplies	1-node	2 x 1400 W	2 x 1975 W	4 x 900 W or 1400 W <sup>c</sup>
	2-node	Not available	4 x 1975 W	8 x 900 W or 1400W <sup>c</sup>

a. Requires all processors to be installed to use all memory slots.

b. Model dependent.

c. Mixing of power supplies in pairs is supported.



## 1.4 Storage versus in-memory data

Main memory (RAM) is the fastest storage type that can hold a significant amount of data. Data in main memory can be accessed more than a 100,000 times faster than data on a spinning hard disk, and even flash technology storage is about a thousand times slower than main memory.

Main memory is connected directly to the processors through a high-speed bus, whereas hard disks are connected through a chain of buses (QPI, PCIe, and SAN) and controllers (I/O hub, RAID controller or SAN adapter, and storage controller).

Compared to keeping data on disk, keeping the data in main memory can dramatically improve database performance because of the improved access time. However, there is one potential drawback. In a database transaction that was committed, the transaction cannot stay committed.

In database technology, *atomicity*, *consistency*, *isolation*, and *durability* (ACID) is a set of requirements that ensures that database transactions are processed reliably. Consider the following points:

- ▶ A transaction must be atomic. That is, if part of a transaction fails, the entire transaction must fail and leave the database state unchanged.
- ▶ The consistency of a database must be preserved by the transactions that it performs.
- ▶ Isolation ensures that no transaction interferes with another transaction.
- ▶ Durability means that after a transaction is committed, it will remain committed.

Although the first three requirements are not affected by the in-memory concept, durability is a requirement that cannot be met by storing data in main memory alone because main memory is volatile storage. That is, it loses its content when no electrical power is present. To make data persistent, it must be on non-volatile storage. Therefore, some sort of permanent storage is still needed, such as hard disk drives (HDDs) or solid-state drives (SSDs) to form a hybrid solution that uses in-memory and disk technology together.

The advantage of a hybrid solution can mean flexibility by balancing the performance, cost, and persistence and form factor in the following ways:

- ▶ Performance: Use in-memory technology to enhance performance of sorting, storing, and retrieving specified data rather than going to disk.
- ▶ Persistence and form factor: Memory cannot approach the density of a small HDD.
- ▶ Cost: Less costly HDDs can be substituted for more memory.

In the next section, we describe Lenovo technologies for flash storage.

## 1.5 Flash storage

Lenovo flash storage offerings for X6 servers combine extreme IOPS performance and low response time for transactional database workloads. The flash technologies that are used in the X6 servers include PCIe NVMe drives, Flash Storage Adapters, and SAS/SATA SSDs.

The following flash storage options are available:

- ▶ PCIe NVMe drives

Non-Volatile Memory Express (NVMe) is new PCIe 3.0 high-performance SSD technology that provides higher bandwidth and shorter latency, which enables systems to process large amounts of data. These SSDs connect directly to the PCIe host interface and do not require a SAS or SATA controller.

Each NVMe drive has a direct PCIe 3.0 x4 connection that provides more than 4x random read IOPs versus 12 Gb SAS, and more than 2x sequential read/write MBps versus 12 Gb SAS. NVMe drives are optimized for heavy multithreaded workloads by using internal parallelism and many other improvements, such as enlarged I/O queues.

For more information, see 2.6.3, “NVMe SSD technology” on page 45.

- ▶ Flash Storage Adapters

These storage devices offer high-performance NAND technology in a standard PCIe adapter form factor. Up to 5.2 TB capacities are available. These adapters have very high sustained sequential read/write throughput rates and very high random read/write I/O operations per second.

The adapters contain no moving parts and do not have the issues that are associated with HDDs, such as vibration, noise, and mechanical failure. The adapters are built as block devices on a PCIe bus with advanced wear-leveling, ECC protection, and chip-level fault tolerance, which provides exceptional reliability and efficiency.

For more information, see 2.6.4, “Flash Storage Adapters” on page 47.

- ▶ SSDs with SAS/SATA interfaces

The X6 servers support 2.5-inch and 1.8-inch SSDs with 12 Gb SAS/SATA host connections. These drives include the 3.84 TB Enterprise Capacity SAS SSD. These drives provide high performance and capacity. Supported RAID controllers offer various RAID and performance options.

For more information, see 2.6.2, “SAS/SATA SSD technology” on page 44.

These technologies allow X6 servers to deliver break-through performance for targeted workloads by offering lower latency and higher performance compared to traditional SSDs.

## 1.6 Energy efficiency

The x3850 X6 and x3950 X6 offer the following energy-efficiency features to save energy, reduce operational costs, increase energy availability, and contribute to the green environment:

- ▶ Energy-efficient electronic components that help lower operational costs.
- ▶ Highly efficient 900 W AC and 1400 W AC power supplies have 80 PLUS Platinum certification.
- ▶ Intel Xeon processor E7-4800/8800 v4 product families offer better performance over the previous generation while fitting into the same thermal design power (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 use 10% less energy compared to 1.5 V DDR3 RDIMMs.
- ▶ Low-voltage 1.2 V DDR4 memory RDIMMs use 20% less energy compared to 1.5 V DDR3 RDIMMs.
- ▶ SSDs use as much as 80% less power than traditional spinning 2.5-inch HDDs.
- ▶ The server uses hexagonal ventilation holes that can be grouped more densely than round holes, which provide more efficient airflow through the system.
- ▶ Lenovo XClarity Energy Manager provides data center energy management by monitoring power and temperature at the server level and the group level, to help data center administrators improve business continuity and energy efficiency.

## 1.7 Services offerings

The Lenovo Professional Services and Data Center Services teams help clients worldwide with deployment of System x solutions and technologies and data center assessments. The Professional Services team offers the following benefits in support of System x servers and solutions:

- ▶ Asset management
- ▶ Information infrastructure
- ▶ Service management
- ▶ Security
- ▶ Virtualization
- ▶ Consolidation
- ▶ Business and collaborative solutions
- ▶ Installation
- ▶ Configuration
- ▶ Integration
- ▶ Migration

In addition to these offerings for System x, the Professional Services team has the following offerings specifically for X6:

- ▶ Virtualization Enablement
- ▶ Database Enablement
- ▶ Enterprise Application Enablement
- ▶ Migration Study
- ▶ Virtualization Health Check

The Data Center Services team offers in-depth data center power and cooling assessments, including the following areas:

- ▶ Planning for high-density systems and cloud for the data center
- ▶ Data center baseline cooling assessment
- ▶ Data center power and cooling resiliency assessment
- ▶ Retail and campus data closet power and cooling planning

For more information about these services, see this website:

<http://shop.lenovo.com/us/en/systems/services/>

The services offerings are designed around having the flexibility to be customized to meet your needs and can provide preconfigured services, custom services, expert skills transfer, off-the-shelf training, and online or classroom courses for X6.

For more information, contact the appropriate team that is listed in Table 1-2.

*Table 1-2 Lab Services and Data Center Services contact details*

Contact	Geography
<b>Lab Services team</b>	
x86svcs@lenovo.com	North America (US and Canada)
x86svcAP@lenovo.com	Asia Pacific (GCG, ANZ, ASEAN, Japan, Korea, and ISA)
x86svcLA@lenovo.com	Latin America (Brazil, Mexico, and SSA)
x86svcEP@lenovo.com	Europe, Middle East, and Africa
<b>Data Center Services team</b>	
power@lenovo.com	Worldwide

## 1.8 About this book

In this book, readers get a general understanding of X6 technology, which sets it apart from previous models, and the architecture that makes up this product line. This book is broken down into the following chapters:

- ▶ The first three chapters give an in-depth look at the X6 hardware and architecture.
- ▶ The remainder of the book describes preparing and implementing the X6 server. We describe power and cooling considerations, operating installations, systems management, and firmware update tools.

## New technology in X6 servers

In this chapter, we give an overview of the technologies that are included in the Lenovo System x3850 X6 and x3950 X6 servers. We describe system architecture, chassis design with modular structure, and Intel Xeon E7-4800 and E7-880 series processors, including the latest v4 generation with DDR4 memory support.

Then we describe the current memory options and features of the storage subsystem, including innovative memory-channel storage technology and PCIe NVMe solid-state drives (SSDs). We also describe other advanced technology in the servers, including X6 scaling and partitioning capabilities.

This chapter includes the following topics:

- ▶ 2.1, “Modular design” on page 12
- ▶ 2.2, “System architecture” on page 20
- ▶ 2.3, “Processors” on page 24
- ▶ 2.4, “Memory” on page 33
- ▶ 2.5, “PCIe 3.0” on page 43
- ▶ 2.6, “Internal storage” on page 43
- ▶ 2.7, “Unified Extensible Firmware Interface” on page 49
- ▶ 2.8, “Integrated Management Module” on page 49
- ▶ 2.9, “Scalability” on page 50

## 2.1 Modular design

X6 rack family consists of the following new flagship servers of the Lenovo x86 server family:

- ▶ Lenovo System x3850 X6: 4U rack-optimized server; scalable to four sockets
- ▶ Lenovo System x3950 X6: 8U rack-optimized server; scalable to eight sockets

Figure 2-1 shows the System x3850 X6.



Figure 2-1 Lenovo System x3850 X6

The x3950 X6 resembles two x3850 X6 servers, with one server placed on top of the other server. However, unlike the earlier x3950 eX5 servers (which connected two x3850 servers via external cables), x3950 X6 uses a single chassis with a single-midplane design, which eliminates the need for external connectors and cables.

Figure 2-2 shows the System x3950 X6.



Figure 2-2 *Lenovo System x3950 X6*

The X6 systems offer a new “bookshelf” design concept that is based on a fixed chassis mounted in a standard rack cabinet. There is no need to pull the chassis out of the rack to access components because all components can be accessed from the front or rear, just as books are pulled from a bookshelf.

Figure 2-3 shows the x3850 X6 server with one of the four Compute Books partially removed.



Figure 2-3 *x3850 X6 server with a Compute Book partially removed*

The modular component that can be installed in a chassis is called a *book*. The following types of books are available:

► **Compute Books**

A Compute Book contains 1 processor, 24 DIMM slots, and 2 hot-swap fan modules. It is accessible from the *front* of the server.

The x3850 X6 supports up to four Compute Books. The x3950 X6 supports up to eight Compute Books.

► **Storage Books**

The Storage Book contains standard 2.5-inch hot-swap SSDs/HDDs or 1.8-inch hot-swap SSDs. It also provides front USB and video ports, and it has two PCIe slots that are reserved for internal storage adapters. The Storage Book is accessible from the *front* of the server.

The x3850 X6 includes one Storage Book standard (maximum one); the x3950 X6 includes two Storage Books standard (maximum two).

► **I/O Books**

An I/O Book is a container that provides PCIe expansion capabilities. I/O Books are accessible from the *rear* of the server.

The following types of I/O Books are available:

- **Primary I/O Book.** This book provides core I/O connectivity, including the ML2 unique slot for an onboard network, three standard PCIe 3.0 slots, Integrated Management Module II, hot-swap fan modules and USB, video, serial, and systems management ports.
- **Full-length I/O Book.** This hot-swap book provides three optional full-length PCIe slots, and two of the slots can host a Graphics Processing Unit (GPU) or coprocessor adapters with a total power consumption of 300 W.
- **Half-length I/O Book.** This hot-swap book provides three optional half-length PCIe slots.

The x3850 X6 includes one Primary I/O Book and supports one or two of the full or half-length I/O Books (one of each book or two of either book). The x3950 X6 includes *two* Primary I/O Books and supports up to four of the full or half-length I/O Books (any combination).

## 2.1.1 Compute Books

The x3850 X6 server supports up to four Compute Books, one Storage Book, one Primary I/O Book, and up to two optional I/O Books. In addition, the 4U chassis supports up to four power supplies and up to 10 hot-swap dual-motor fans (eight fans on the front and two fans on the rear of the server).



Figure 2-4 shows the front of the x3850 X6 server in which the four Compute Books and the Storage Book are highlighted.




Figure 2-4 x3850 X6 front view

Figure 2-5 shows the rear view of the x3850 X6 server in which the Primary I/O Book, other I/O Books, and power supplies are highlighted.




Figure 2-5 x3850 X6 rear view

In the middle of the server, a passive midplane connects all modular components. Figure 2-6 shows the x3850 X6 midplane.




Figure 2-6 x3850 X6 midplane (front side showing Compute Book connections)

Each Compute Book contains 1 Intel Xeon processor, 24 memory DIMMs, and 2 dual-motor fans. The 24 DIMM slots are on both sides of the Compute Book's processor board, with 12 memory modules on each side.

Figure 2-7 shows the Compute Book with the clear side cover removed. The front of the Compute Book includes two hot-swap fans.




Figure 2-7 Compute Book

The x3850 X6 server supports up to four Compute Books; the x3950 X6 server supports up to eight . For more information about the Compute Books, see 3.7, “Compute Books” on page 75.

## 2.1.2 Storage Book

The Storage Book is accessible from the front of the server. It has bays for 1.8-inch (up to 16) or 2.5-inch (up to 8) drives, two slots for RAID controllers to connect to those drive bays, a front operator panel, and USB and video ports. Figure 2-8 shows the Storage Book.



Figure 2-8 Storage Book

The x3850 X6 server includes one Storage Book (maximum one); the x3950 X6 includes two Storage Books (maximum two). For more information about the Storage Book, see 3.11.1, “Storage Book” on page 94.

## 2.1.3 Primary I/O Book

The Primary I/O Book is the third core module, which provides base I/O connectivity, including the following four PCIe slots:

- ▶ Three standard PCIe 3.0 half-length slots
- ▶ A single slot for an ML2 form-factor network adapter

The Primary I/O Book also contains core logic, such as the Integrated Management Module II (IMM2) and Unified Extensible Firmware Interface (UEFI), fan modules, and peripheral ports. The Primary I/O Book installs in the rear of the server.

Figure 2-9 shows the internal components of the Primary I/O Book, as viewed from the front of the server. The black air baffle is raised to show other internal components.




Figure 2-9 Primary I/O Book internal components

The x3850 X6 includes one Primary I/O Book (maximum one); the x3950 X6 includes two Primary I/O Books (maximum two). For more information about the Primary I/O Book, see 3.13, “Primary I/O Book” on page 102.

## 2.1.4 More I/O Books

More I/O Books provide extra I/O capabilities, each with *three* PCIe slots. The following types of other I/O Books are available:

- ▶ Half-length I/O Book, which supports half-length adapters
- ▶ Full-length I/O Book, which supports full-length or half-length adapters

The Full-length I/O Book accepts GPU adapters and coprocessors, including double-wide adapters that require up to 300 W of power. The Full-length I/O Book includes *two* auxiliary power connectors, 150 W and 75 W, and power cables.

These I/O Books are accessible from the *rear* of the server. Figure 2-10 shows the two I/O Books.



Figure 2-10 Other I/O Books

The x3850 X6 server supports up to *two* extra I/O Books of any type; the x3950 X6 server supports up to *four* extra I/O Books.

For more information about the extra I/O Books, see 3.14.1, “Half-length I/O Book” on page 106 and 3.14.2, “Full-length I/O Book” on page 106.

## 2.1.5 Power supplies

The x3850 X6 supports up to *four* hot-swap redundant power supplies and the x3950 X6 supports up to *eight* power supplies. Available power supplies are the 900 W AC, 1400 W AC and 750W -48V DC power supplies. Figure 2-11 shows a 1400 W AC power supply.



Figure 2-11 1400 W power supply

For more information about the power supplies, see 3.24, “Power subsystem” on page 122.

## 2.2 System architecture

This section describes the overall architecture of the x3850 X6 and x3950 X6.

### 2.2.1 x3850 X6

Figure 2-12 shows the system architecture of the x3850 X6 server.

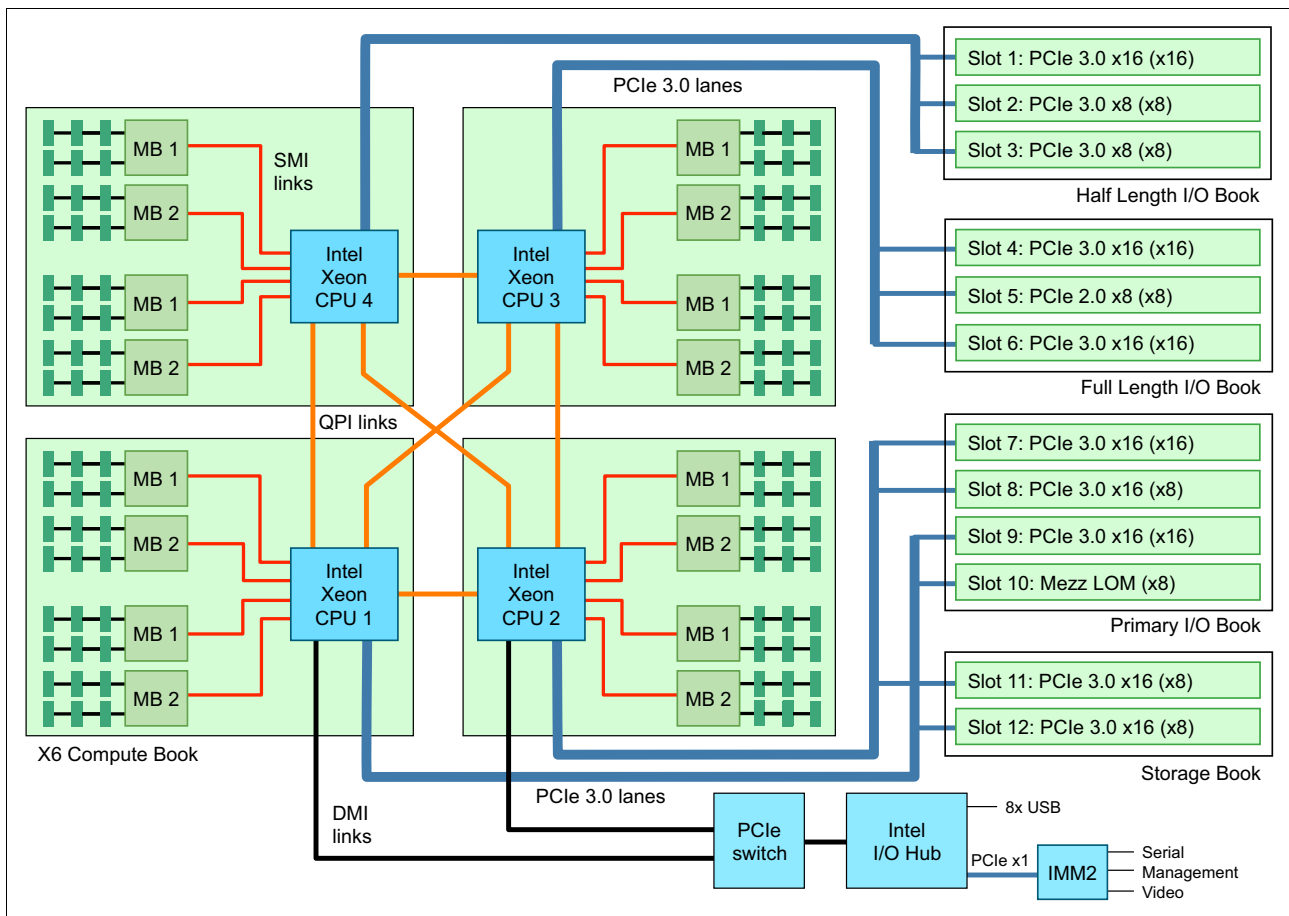


Figure 2-12 x3850 X6 system architecture

Processor-to-processor communication is carried over shared-clock or coherent quick path interconnect (QPI) links. Each processor has *three* QPI links to connect to other processors.

Figure 2-13 shows how the four processors of the x3850 X6 are connected via QPI links. In this configuration, each processor is directly connected to all other processors.




Figure 2-13 QPI links between processors

In the Compute Book, each processor has four Scalable Memory Interconnect Generation 2 (SMI2) channels (two memory controllers per processor, each with two SMI channels) that are connected to four scalable memory buffers. Each memory buffer has six DIMM slots (two channels with three DIMMs per channel) for a total of 24 DIMMs (eight channels with three DIMMs per channel) per processor. Compute Books are connected to each other via QPI links.

The Primary I/O Book has three PCIe 3.0 slots, a Mezzanine LOM slot, an I/O Controller Hub, IMM2, and peripheral ports (such as USB, video, serial) on the board. Extra I/O Books (Full-length and Half-length) have three PCIe 3.0 slots each and support hot-swap PCIe adapters.

**Extra I/O Books:** For illustration purposes, Half-length and Full-length I/O Books are shown in Figure 2-12 on page 20, where the Half-length I/O Book supplies slots 1, 2, and 3, and the Full-length I/O Book supplies slots 4, 5, and 6. Their order can be reversed, or two of either type can be used.

The Primary I/O Book is connected to the Compute Books 1 (CPU 1) and 2 (CPU 2) directly via PCIe links from those processors: PCIe slots 9 and (for dedicated mezzanine NIC) 10 are connected to CPU 1, and PCIe slots 7 and 8 are connected to CPU 2. Also, CPU 1 and CPU 2 are connected to the Intel I/O Hub via DMI switched links for redundancy purposes.

The Storage Book also is connected to Compute Books 1 and 2; however, PCIe slots 11 and 12 (both used for dedicated HBA/RAID cards) are connected to different processors (CPU 2 and CPU 1, respectively). In addition, certain peripheral ports are routed from the Intel I/O Hub and IMM2 to the Storage Book.

Extra I/O Books are connected to Compute Books 3 and 4 and use PCIe links from CPU 3 and CPU 4. If you need to install more I/O Books, you should install the Compute Book in an appropriate slot first.

## 2.2.2 x3950 X6

The x3950 X6 server uses the same architecture, but adapted for an 8-socket configuration. Figure 2-14 and Figure 2-15 on page 23 show the system architecture of the x3950 X6 server. (Figure 2-14 shows the lower half and Figure 2-15 on page 23 shows the upper half.)

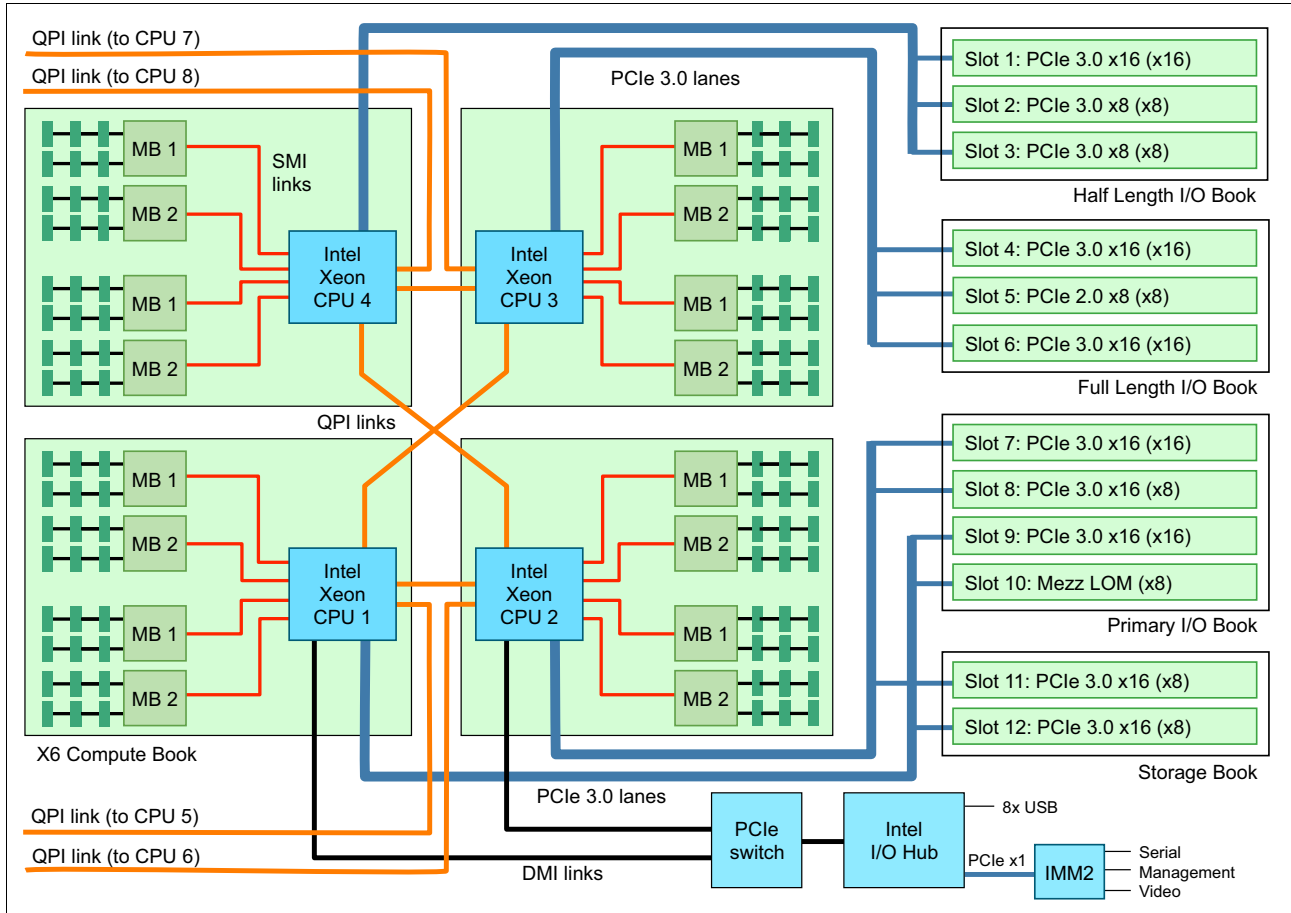


Figure 2-14 x3950 X6 (bottom half)



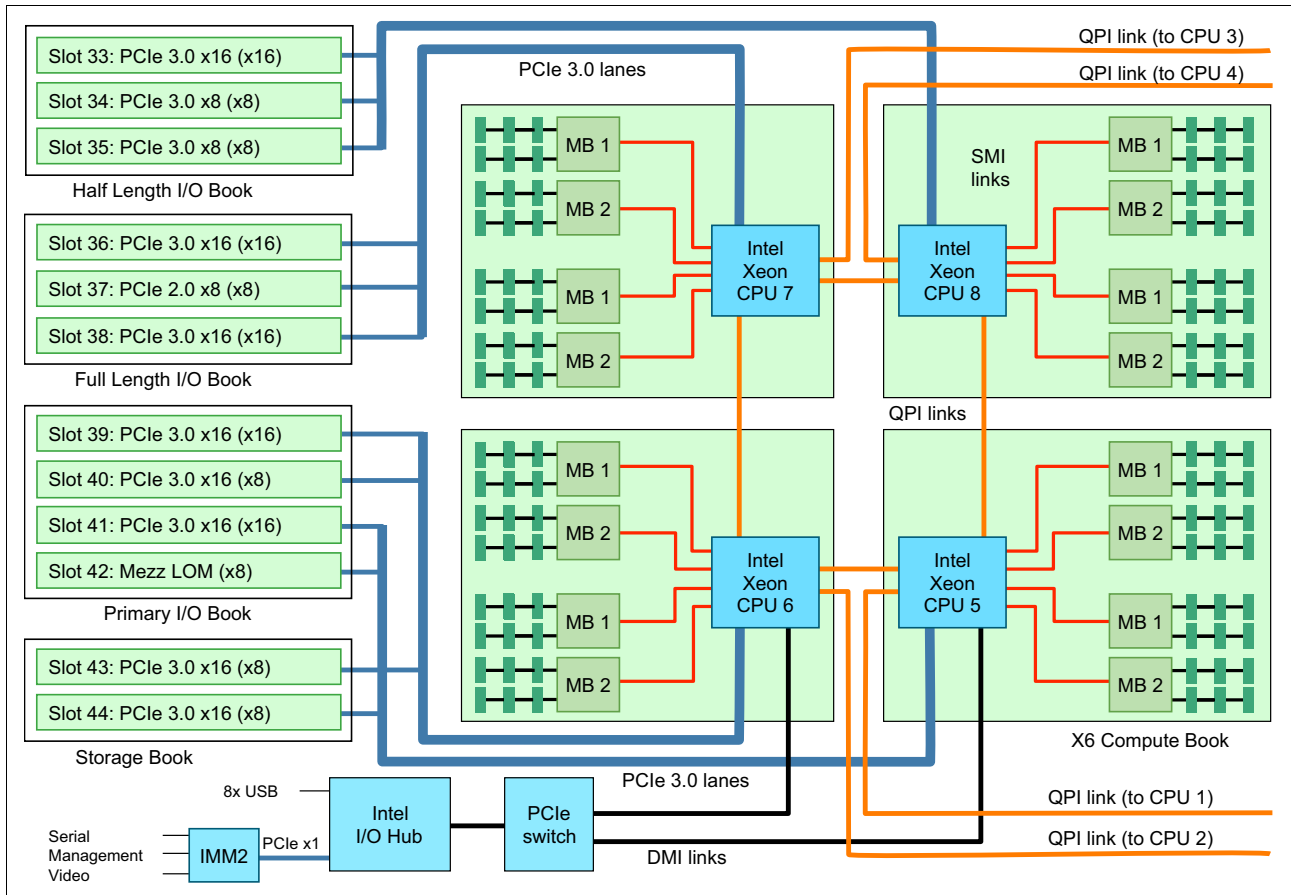


Figure 2-15 x3950 X6 system architecture (top half)

The 8-socket configuration is formed by using the native QPI scalability of the Intel Xeon processor E7 family.

Figure 2-16 shows how the processors are connected via QPI links.




Figure 2-16 QPI connectivity: x3950 X6 with eight processors installed

The x3950 X6 is an eight-socket server but also supports four and six processors installed, with upgradability to eight processors. Figure 2-17 shows the x3950 X6 with only *four* processors installed. The effective connectivity is shown on the right side of the figure where each processor is either connected directly to a processor or one hop away.

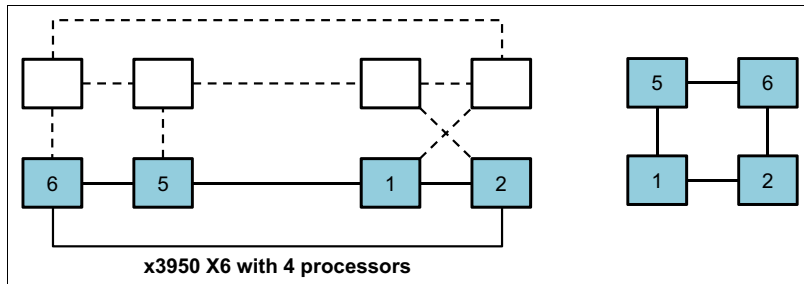


Figure 2-17 QPI connectivity: x3950 X6 with 4 processors installed

Figure 2-18 shows the x3950 X6 with only *six* processors installed. The connectivity is shown on the right of the figure where each processor is either connected directly to another processor or one hop away.




Figure 2-18 QPI connectivity: x3950 X6 with 6 processors installed

In addition, the 8-socket server can form *two independent systems* that contain four sockets in each node, as if two independent 4U x3850 X6 servers are housed in one 8U chassis. This partitioning feature is enabled via the IMM2 interface. When partitioning is enabled, each partition can deploy its own operating system and applications. Each partition uses its own resources and can no longer access the other partition's resources.

## 2.3 Processors

The current models of the X6 systems use the Intel Xeon E7 processor family. The new Intel Xeon E7 v4 processors feature the new Intel microarchitecture (formerly codenamed "Broadwell-EX") that provides higher core count, larger cache sizes, and DDR4 memory support. Intel Xeon E7 v2, v3 and v4 families support up to 24 DIMMs per processor and provide fast low-latency I/O with integrated PCIe 3.0 controllers.

### 2.3.1 Intel Xeon processor E7-4800/8800 v4 product family

The following groups of the Intel Xeon processor E7 family are used in the X6 servers:

- ▶ The Intel Xeon processor E7-4800/8800 v4 product family is supported in the x3850 X6. This family supports four-processor configurations. (Using E7-8800 v4 processors enables the swapping of Compute Books between x3850 X6 and x3950 X6 servers.)
- ▶ The Intel Xeon processor E7-8800 v4 product family is used in the x3950 X6 to scale to eight-socket configurations.

The X6 systems support the latest generation of the Intel Xeon processor E7-4800 v4 and E7-8800 v4 product family, which offers the following key features:

- ▶ Up to 24 cores and 48 threads (by using Hyper-Threading feature) per processor
- ▶ Up to 60 MB of shared last-level cache
- ▶ Up to 3.2 GHz core frequencies
- ▶ Up to 9.6 GTps bandwidth of QPI links
- ▶ DDR4 memory interface support, which brings greater performance and power efficiency
- ▶ Integrated memory controller with four SMI2 Gen2 channels that support up to 24 DDR4 DIMMs
- ▶ Memory channel (SMI2) speeds up to 1866 MHz in RAS (lockstep) mode and up to 3200 MHz in performance mode.
- ▶ Integrated PCIe 3.0 controller with 32 lanes per processor
- ▶ Intel Virtualization Technology (VT-x and VT-d)
- ▶ Intel Turbo Boost Technology 2.0
- ▶ Improved performance for integer and floating point operations
- ▶ Virtualization improvements with regards to posted interrupts, page modification logging, and VM enter/exit latency reduction
- ▶ New Intel Transactional Synchronization eXtensions (TSX)
- ▶ Intel Advanced Vector Extensions 2 (AVX2.0) with new optimized turbo behavior
- ▶ Intel AES-NI instructions for accelerating of encryption
- ▶ Advanced QPI and memory reliability, availability, and serviceability (RAS) features
- ▶ Machine Check Architecture recovery (non-running and running paths)
- ▶ Enhanced Machine Check Architecture Gen2 (eMCA2)
- ▶ Machine Check Architecture I/O
- ▶ Resource director technology: Cache monitoring technology, cache allocation technology, memory bandwidth monitoring
- ▶ Security technologies: OS Guard, Secure Key, Intel TXT, Crypto performance (ADOX/ADCX), Malicious Software (SMAP), Key generation (RDSEED)

Table 2-1 compares the Intel Xeon E7-4800/8800 processors that are supported in X6 systems.

Table 2-1 X6 processors comparisons

Feature	X6 family, Xeon E7 v2	X6 family, Xeon E7 v3	X6 family, Xeon E7 v4
Processor family	Intel Xeon E7-8800 v2 Intel Xeon E7-4800 v2	Intel Xeon E7-8800 v3 Intel Xeon E7-4800 v3	Intel Xeon E7-8800 v4 Intel Xeon E7-4800 v4
Processor codenames	Ivy Bridge EX	Haswell EX	Broadwell EX

Feature	X6 family, Xeon E7 v2	X6 family, Xeon E7 v3	X6 family, Xeon E7 v4
Cores per CPU	Up to 15 cores	Up to 18 cores	Up to 24 cores
Last level cache	Up to 37.5 MB L3 cache	Up to 45 MB L3 cache	Up to 60 MB L3 cache
QPI	QPI 1.1 at 8.0 GT/s max	QPI 1.1 at 9.6 GT/s max	QPI 1.1 at 9.6 GT/s max
CPU TDP rating	Up to 155 W	Up to 165 W	Up to 165 W
DIMM sockets	24 DDR3 DIMMs per CPU	24 DDR3 DIMMs per CPU 24 DDR4 DIMMs per CPU	24 DDR4 DIMMs per CPU
Maximum memory speeds	2667 MHz SMI2	3200 MHz SMI2	3200 MHz SMI2
PCIe technology	PCIe 3.0 (8 GTps)	PCIe 3.0 (8 GTps)	PCIe 3.0 (8 GTps)

### 2.3.2 Intel Xeon processor E7-4800/8800 v3 product family

Similar to the latest Intel Xeon E7 v4 processors, the Intel Xeon E7 v3 processor family provides two types of processors: E7-4800 v3 and E7-8800 v3 for four- and eight-socket configurations.

The Intel Xeon processor E7-4800 v3 and E7-8800 v3 product family offers the following key features:

- ▶ Up to 18 cores and 36 threads (by using Hyper-Threading feature) per processor
- ▶ Up to 45 MB of shared last-level cache
- ▶ Up to 3.2 GHz core frequencies
- ▶ Up to 9.6 GTps bandwidth of QPI links
- ▶ Integrated memory controller with four SMI2 channels that support up to 24 DDR3/DDR4 DIMMs
- ▶ Up to 1600 MHz DDR3 or 1866 MHz DDR4 memory speeds
- ▶ DDR4 memory channel (SMI2) speeds up to 1866 MHz in RAS (lockstep) mode and up to 3200 MHz in performance mode.
- ▶ Integrated PCIe 3.0 controller with 32 lanes per processor
- ▶ Intel Virtualization Technology (VT-x and VT-d)
- ▶ Intel Turbo Boost Technology 2.0
- ▶ Intel Advanced Vector Extensions 2 (AVX2)
- ▶ Intel AES-NI instructions for accelerating of encryption
- ▶ Advanced QPI and memory RAS features
- ▶ Machine Check Architecture recovery (non-running and running paths)
- ▶ Enhanced Machine Check Architecture Gen2 (eMCA2)
- ▶ Machine Check Architecture I/O
- ▶ Security technologies: OS Guard, Secure Key, Intel TXT

### 2.3.3 Intel Xeon processor E7-4800/8800 v2 product family

The Intel Xeon E7 v2 processor family provides two types of processors as well: E7-4800 v2 and E7-8800 v2 for four- and eight-socket configurations.

The Intel Xeon processor E7-4800 v2 and E7-8800 v2 product family offers the following key features:

- ▶ Up to 15 cores and 30 threads (by using Hyper-Threading feature) per processor
- ▶ Up to 37.5 MB of L3 cache
- ▶ Up to 3.4 GHz core frequencies
- ▶ Up to 8 GTps bandwidth of QPI links
- ▶ Integrated memory controller with four SMI2 channels that support up to 24 DDR3 DIMMs
- ▶ Up to 1600 MHz DDR3 memory speeds
- ▶ Integrated PCIe 3.0 controller with 32 lanes per processor
- ▶ Intel Virtualization Technology (VT-x and VT-d)
- ▶ Intel Turbo Boost Technology 2.0
- ▶ Intel Advanced Vector Extensions (AVX)
- ▶ Intel AES-NI instructions for accelerating of encryption
- ▶ Advanced QPI and memory RAS features
- ▶ Machine Check Architecture recovery (non-running and running paths)
- ▶ Enhanced Machine Check Architecture Gen1 (eMCA1)
- ▶ Machine Check Architecture I/O
- ▶ Security technologies: OS Guard, Secure Key, Intel TXT

### 2.3.4 Intel Xeon E7 features

Intel Xeon E7 processors include a broad set of features and extensions. Many of these technologies are common for all Intel Xeon E7 generations; some technologies are unique to the latest Intel Xeon E7 v4 family.

#### Intel Transactional Synchronization eXtensions

Intel Transactional Synchronization eXtensions (TSX) feature the latest v4 and v3 generations of Intel Xeon E7 processor families and brings hardware transactional memory support. Intel TSX implements a memory-locking approach that is called Hardware Lock Elision (HLE), which facilitates running multithreaded applications more efficiently.

Much TSX-aware software gained great performance boosts by running on Intel Xeon E7 v4 processors. For example, SAP HANA SPS 09 in-memory database showed twice as many transactions per minute with Intel TSX enabled versus TSX disabled on E7 v3 processors and three times more transactions per minute compared to Intel Xeon E7 v2 processors.

For more information about Intel TSX, see the Solution Brief, *Ask for More from Your Data*, which is available here:

<http://www.intel.com/content/dam/www/public/us/en/documents/solution-briefs/sap-hana-real-time-analytics-solution-brief.pdf>

#### Intel Advanced Encryption Standard

Advanced Encryption Standard (AES) is an encryption standard that is widely used to protect network traffic and sensitive data. Advanced Encryption Standard - New Instructions

(AES-NI), which is available with the E7 processors, implements certain complex and performance intensive steps of the AES algorithm by using processor hardware. AES-NI can accelerate the performance and improve the security of an implementation of AES versus an implementation that is performed by software.

For more information about Intel AES-NI, see this website:

<http://software.intel.com/en-us/articles/intel-advanced-encryption-standard-instructions-aes-ni>

## Intel Virtualization Technology

Intel Virtualization Technology (Intel VT) is a suite of processor and I/O hardware enhancements that assists virtualization software to deliver more efficient virtualization solutions and greater capabilities.

Intel Virtualization Technology for x86 (Intel VT-x) allows the software hypervisors to better manage memory and processing resources for virtual machines (VMs) and their guest operating systems.

Intel Virtualization Technology for Directed I/O (Intel VT-d) helps improve I/O performance and security for VMs by enabling hardware-assisted direct assignment and isolation of I/O devices.

For more information about Intel Virtualization Technology, see this website:

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

## Hyper-Threading Technology

Intel Hyper-Threading Technology enables a single physical processor to run two separate code streams (threads) concurrently. To the operating system, a processor core with Hyper-Threading is seen as two logical processors. Each processor has its own architectural state; that is, its own data, segment, and control registers, and its own advanced programmable interrupt controller (APIC).

Each logical processor can be individually halted, interrupted, or directed to run a specified thread independently from the other logical processor on the chip. The logical processors share the running resources of the processor core, which include the running engine, caches, system interface, and firmware.

Hyper-Threading Technology improves server performance. This process is done by using the multithreading capability of operating systems and server applications in such a way as to increase the use of the on-chip running resources that are available on these processors. Application types that make the best use of Hyper-Threading include virtualization, databases, email, Java, and web servers.

For more information about Hyper-Threading Technology, see this website:

<http://www.intel.com/technology/platform-technology/hyper-threading>

**vSphere 5.1 and 8-socket systems:** VMware vSphere 5.1 has a fixed upper limit of 160 concurrent threads. Therefore, if you use an 8-socket system with more than 10 cores per processor, you should disable Hyper-Threading.

## Turbo Boost Technology 2.0

The Intel Xeon E7-8800/4800 family of processors brings enhanced capabilities for changing processor speed with Intel Turbo Boost 2.0 technology.

Intel Turbo Boost Technology dynamically saves power on unused processor cores and increases the clock speed of the cores in use. Depending on current workload, Intel Turbo Boost Technology allows a dynamic increase in the clock speed of the active cores to gain a performance boost. For example, a 3.4 GHz 15-core processor can overclock the cores up to 3.7 GHz.

Turbo Boost Technology is available on a per-processor basis for the X6 systems. For ACPI-aware operating systems and hypervisors, such as Microsoft Windows 2008/2012, RHEL 5/6, SLES 11, VMware ESXi 4.1, and later, no changes are required to use it. Turbo Boost Technology can be used with any number of enabled and active cores, which results in increased performance of multithreaded and single-threaded workloads.

Turbo Boost Technology dynamically saves power on unused processor cores and increases the clock speed of the cores in use. In addition, it can temporarily increase the speed of all cores by intelligently managing power and thermal headroom. For example, a 2.5 GHz 15-core processor can temporarily run all 15 active cores at 2.9 GHz. With only two cores active, the same processor can run those active cores at 3.0 GHz. When the other cores are needed again, they are turned back on dynamically and the processor frequency is adjusted.

When temperature, power, or current exceeds factory-configured limits and the processor is running above the base operating frequency, the processor automatically steps the core frequency back down to reduce temperature, power, and current. The processor then monitors these variables, and reevaluates whether the current frequency is sustainable or if it must reduce the core frequency further. At any time, all active cores run at the same frequency.

For more information about Turbo Boost Technology, see this website:

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

### QuickPath Interconnect

The Intel Xeon E7 processors implemented in X6 servers include two integrated memory controllers in each processor. Processor-to-processor communication is carried over shared-clock or coherent QPI links. Each processor has three QPI links to connect to other processors.

Figure 2-19 shows the QPI configurations. On the left is how the four sockets of the x3850 X6 are connected. On the right is how all eight sockets of the x3950 X6 are connected.




Figure 2-19 QPI links between processors

Each processor has some memory that is connected directly to that processor. To access memory that is connected to *another* processor, each processor uses QPI links through the other processor. This design creates a non-uniform memory access (NUMA) system. Similarly, I/O can be local to a processor or remote through another processor.

For QPI use, Intel modified the MESI cache coherence protocol to include a forwarding state. Therefore, when a processor asks to copy a shared cache line, only one other processor responds.

For more information about QPI, see this website:

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

## Intel Data Direct I/O

For I/O, Intel no longer has a separate I/O hub. Instead, 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.

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>

## RAS features

The Intel Xeon processor E7 family of processors has the following RAS features on their interconnect links (SMI and QPI):

- ▶ Cyclic redundancy checking (CRC) on the QPI links  
The data on the QPI link is checked for errors.
- ▶ QPI packet retry  
If a data packet on the QPI link has errors or cannot be read, the receiving processor can request that the sending processor try sending the packet again.
- ▶ QPI clock failover  
If there is a clock failure on a coherent QPI link, the processor on the other end of the link can become the clock. This action is not required on the QPI links from processors to I/O hubs because these links are asynchronous.
- ▶ QPI self-healing  
If persistent errors are detected on a QPI link, the link width can be reduced dynamically to allow the system to run in a degraded mode until repair can be performed. QPI link can reduce its width to a half width or a quarter width, and slowdown its speed.
- ▶ Scalable memory interconnect (SMI) packet retry  
If a memory packet has errors or cannot be read, the processor can request that the packet be resent from the memory buffer.

## Machine Check Architecture recovery

The Intel Xeon processor E7 family also includes Machine Check Architecture (MCA) recovery, a RAS feature that enables the handling of system errors that otherwise require that the operating system be halted. For example, if a dead or corrupted memory location is discovered but it cannot be recovered at the memory subsystem level and it is not in use by the system or an application, an error can be logged and the operation of the server can continue. If it is in use by a process, the application to which the process belongs can be stopped or informed about the situation.

Implementation of the MCA recovery requires hardware support, firmware support (such as found in the UEFI), and operating system support. Microsoft, SUSE, Red Hat, VMware, and other operating system vendors include or plan to include support for the Intel MCA recovery feature on the Intel Xeon processors in their latest operating system versions.



The new MCA recovery features of the Intel Xeon processor E7-4800/8800 product family include:

- ▶ Execution path recovery: The ability to work with hardware and software to recognize and isolate the errors that were delivered to the execution engine (core).
- ▶ Enhanced MCA (eMCA) Generation 1: Provides enhanced error log information to the operating system, hypervisor, or application that can be used to provide better diagnostic and predictive failure analysis for the system. This feature enables higher levels of uptime and reduced service costs.
- ▶ Enhanced MCA (eMCA) Generation 2: Provides more capabilities for error handling (E7 v3 and v4 processors only).

## Security improvements

The Intel Xeon E7-4800/8800 processor families feature the following important security improvements that help to protect systems from different types of security threats:

- ▶ Intel OS Guard: Evolution of Intel Execute Disable Bit technology, which helps to protect against escalation of privilege attacks by preventing code execution from user space memory pages while in kernel mode. It helps to protect against certain types of malware attacks.
- ▶ #VE2 (Beacon Pass 2 Technology): #VE utilizes ISA-level CPU assists to allow memory-monitoring of antimalicious software performance to scale on virtualized and non-virtualized servers, making deep malicious software detection possible on server platforms.
- ▶ Intel Trusted Execution Technology (Intel TXT), Intel VT-x, and Intel VT-d: New hardware-based techniques, with which you can isolate VMs and start VMs in a trusted environment only. In addition, malware-infected VMs cannot affect nother VMs on the same host.
- ▶ Intel Secure Key: Provides hardware random number generation without storing any data in system memory. It keeps generated random numbers out of sight of malware, which enhances encryption protection.

For more information, see *Crimeware Protection: 3rd Generation Intel Core vPro Processors*, which is available at this website:

<http://www.intel.com/content/dam/www/public/us/en/documents/white-papers/3rd-gen-core-vpro-security-paper.pdf>

## 2.3.5 Compute Books

The core modular element of the X6 design is a Compute Book. It contains the following components:

- ▶ One Intel Xeon E7 v2, v3, or v4 processor
- ▶ 24 memory DDR3 or DDR4 DIMM slots, depending on the CPU used
- ▶ Two dual-motor fans

**x3850 X6 Essential:** x3850 X6 Essential models with compute books containing only 12 DIMM sockets are not covered in this document.

Figure 2-20 shows the Compute Book.



Figure 2-20 The Compute Book

The system board of the Compute Book has two sides, on which all components are installed.

Figure 2-21 and Figure 2-22 show the left and right sides of the Compute Book, respectively. On the left side, there is 1 processor and 12 DIMM slots. On the right, there are 12 DIMM slots, for a total of 24 DIMMs per Compute Book.




Figure 2-21 Left side of Compute Book




Figure 2-22 Right side of Compute Book

The x3850 X6 supports up to four Compute Books, and the x3950 X6 supports up to eight Compute Books. The following configurations are supported:

- ▶ 1, 2, or 4 Compute Books for the x3850 X6
- ▶ 4, 6, or 8 Compute Books for the x3950 X6

**Note:** *All Compute Books in a server must be the same type; that is, they must all have the same processor and either all DDR3 or all DDR4 memory (processor-dependent) at the same frequency/speed.*

## 2.4 Memory

The System x3850 X6 and x3950 X6 servers support three generations of Intel Xeon E7 processors:

- ▶ E7 v4 processors support DDR4 memory only
- ▶ E7 v3 processors can use either DDR3 or DDR4 memory
- ▶ E7 v2 processors support DDR3 memory only

DDR4 is a new memory standard that is supported by the Intel Xeon E7 v3 and v4 processor families. DDR4 memory modules can run at greater speeds than DDR3 DIMMs, operate at lower voltage, and are more energy-efficient than DDR3 modules.

X6 Compute Books with E7 v3 or v4 processors and DDR4 memory interface support Lenovo TruDDR4 memory modules, which are tested and tuned to maximize performance and reliability. Lenovo TruDDR4 DIMMs can operate at greater speeds and have higher performance than DIMMs that only meet industry standards.

DDR3 and TruDDR4 memory types have ECC protection and support Chipkill and Redundant Bit Steering technologies.

Each processor has two integrated memory controllers, and each memory controller has two Scalable Memory Interconnect generation 2 (SMI2) links that are connected to two scalable memory buffers. Each memory buffer has two memory channels, and each channel supports three DIMMs, for a total of 24 DIMMs per processor.

Figure 2-23 shows the processor's memory architecture.

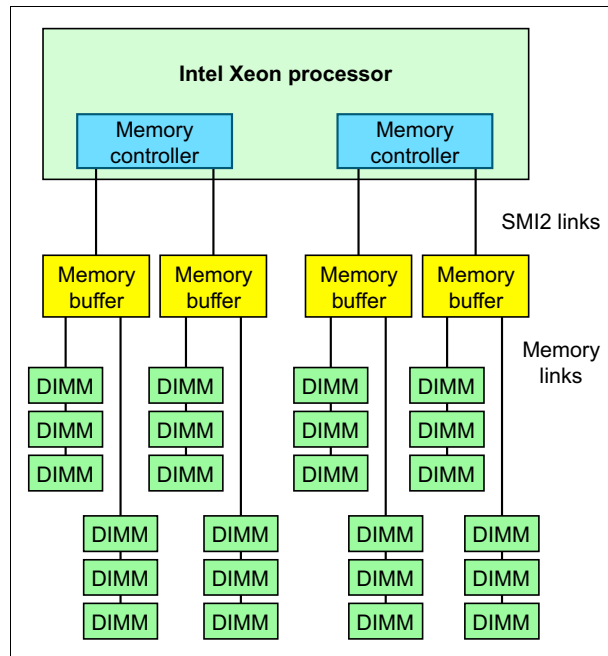


Figure 2-23 Intel Xeon processor E7-4800/8800 memory architecture

E7 **v2**-based Compute Books each support up to 24 DDR3 memory modules running at speeds up to 1600 MHz. Compute Books with E7 **v3** processors can be DDR3-based with up to 24 DDR3 DIMMs operating at speeds up to 1600 MHz, or DDR4-based with up to 24 TruDDR4 memory modules operating at speeds up to 1866MHz. *DDR3 and TruDDR4 DIMMs cannot be mixed*; all memory across all books within a server must be of the same type (either all DDR3 or all TruDDR4). Compute Books with E7 **v4** processors can support up to 24 TruDDR4 DIMMs operating at speeds up to 1866 MHz. DDR3 memory is *not* supported by E7 **v4** processors.

The x3850 X6 supports up to 96 DIMMs when all processors are installed (24 DIMMs per processor), and the x3950 X6 supports up to 192 DIMMs. The processor and the corresponding memory DIMM slots are on the Compute Book (for more information, see 2.3.5, “Compute Books” on page 31).

## 2.4.1 Operational modes

The following memory modes are supported by the Intel Xeon processor E7 product families:

- ▶ Performance mode

In this operation mode, each memory channel works independently and it is addressed individually via burst lengths of 8 bytes (64 bits). The Intel SMI2 channel operates at twice the memory speed. All channels can be populated in any order and modules have no matching requirements.

Chipkill (Single Device Data Correction, or SDDC) is supported in Performance mode. Redundant Bit Steering (RBS) is *not* supported in Performance mode.

Although in this mode DIMMs can be populated in any order, memory modules should be placed based by round-robin algorithm between SMI2 channels and alternating between DDR channels for best performance. For more information about DIMMs population order, see 3.9.2, “Memory population order” on page 89.

► RAS (Lockstep) mode

In RAS operation mode (also known as Lockstep mode), the memory controller operates two memory channels behind one memory buffer as single channel.

In RAS mode, the SMI2 channel operates at the memory transfer rate. DIMMs must be installed in pairs.

Because data is moved by using both channels at once, more advanced memory protection schemes can be implemented to provide protection against single-bit and multibit errors:

- Chipkill, also known as SDDC
- RBS (multibit correction)

The combination of these two RAS features is also known as Double Device Data Correction (DDDC).

Mirroring and sparing are also supported in both modes, as described in 2.4.2, “Memory mirroring and rank sparing” on page 38.

Figure 2-24 shows the two modes. In RAS mode, both channels of one memory buffer are in lockstep with each other.

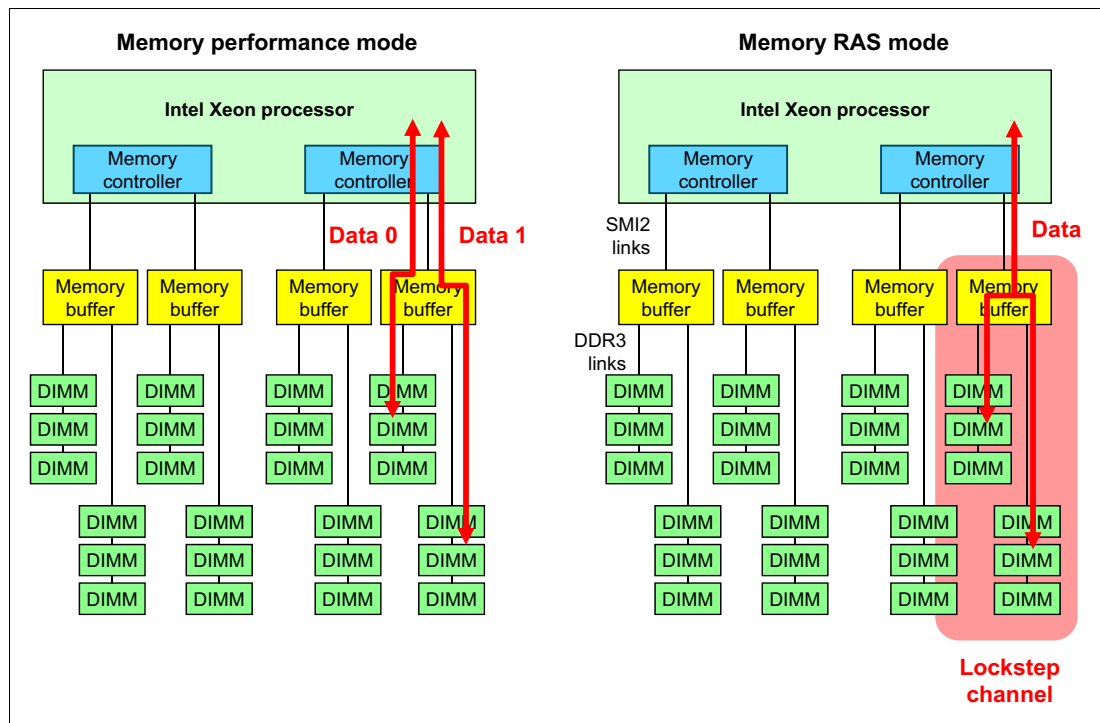


Figure 2-24 Memory modes: Performance mode (left) and RAS mode (right)

The following tables show the maximum speed and bandwidth for the SMI2 and memory channels in both modes for DDR3 and TruDDR4 memory modules, as well as their operating voltages.

Table 2-2 DDR3 memory speeds in Performance mode

DIMM type / capacity	1 DPC 1.35 V	1 DPC 1.5 V	2 DPC 1.35 V	2 DPC 1.5 V	3 DPC 1.35 V	3 DPC 1.5 V
<b>DDR3, Performance mode (2:1): Maximum performance</b>						

DIMM type / capacity	1 DPC 1.35 V	1 DPC 1.5 V	2 DPC 1.35 V	2 DPC 1.5 V	3 DPC 1.35 V	3 DPC 1.5 V
1RX4, 2Gb, 1600 MHz / 4Gb	1333 MHz		1333 MHz			1333 MHz
1RX4, 4Gb, 1600 MHz / 8Gb	1333 MHz		1333 MHz			1333 MHz
2RX4, 4Gb, 1600 MHz / 16Gb	1333 MHz		1333 MHz			1333 MHz
4RX4, 4Gb, 1600 MHz / 32Gb	1333 MHz		1333 MHz		1333 MHz	
8Rx4, 4Gb, 1333 MHz / 64Gb	1333 MHz		1333 MHz		1333 MHz	
<b>DDR3, Performance mode (2:1): Balanced</b>						
1RX4, 2Gb, 1600 MHz / 4Gb	1067 MHz		1067 MHz		1067 MHz	
1RX4, 4Gb, 1600 MHz / 8Gb	1067 MHz		1067 MHz		1067 MHz	
2RX4, 4Gb, 1600 MHz / 16Gb	1067 MHz		1067 MHz		1067 MHz	
4RX4, 4Gb, 1600 MHz / 32Gb	1067 MHz		1067 MHz		1067 MHz	
8Rx4, 4Gb, 1333 MHz / 64Gb	1067 MHz		1067 MHz		1067 MHz	
<b>DDR3, Performance mode (2:1): Minimal</b>						
1RX4, 2Gb, 1600 MHz / 4Gb	1067 MHz		1067 MHz		1067 MHz	
1RX4, 4Gb, 1600 MHz / 8Gb	1067 MHz		1067 MHz		1067 MHz	
2RX4, 4Gb, 1600 MHz / 16Gb	1067 MHz		1067 MHz		1067 MHz	
4RX4, 4Gb, 1600 MHz / 32Gb	1067 MHz		1067 MHz		1067 MHz	
8Rx4, 4Gb, 1333 MHz / 64Gb	1067 MHz		1067 MHz		1067 MHz	

Table 2-3 DDR3 memory speeds in RAS mode

DIMM type / capacity	1 DPC 1.35 V	1 DPC 1.5 V	2 DPC 1.35 V	2 DPC 1.5 V	3 DPC 1.35 V	3 DPC 1.5 V
<b>DDR3, RAS mode (1:1): Maximum performance</b>						
1RX4, 2Gb, 1600 MHz / 4Gb		1600 MHz		1600MHz		1333MHz
1RX4, 4Gb, 1600 MHz / 8Gb		1600 MHz		1600 MHz		1333 MHz
2RX4, 4Gb, 1600 MHz / 16Gb		1600 MHz		1600 MHz		1333 MHz
4RX4, 4Gb, 1600 MHz / 32Gb		1600 MHz		1600 MHz	1333 MHz	
8Rx4, 4Gb, 1333 MHz / 64Gb	1333 MHz		1333 MHz		1333 MHz	
<b>DDR3, RAS mode (1:1): Balanced</b>						
1RX4, 2Gb, 1600 MHz / 4Gb	1333 MHz		1333 MHz		1067 MHz	
1RX4, 4Gb, 1600 MHz / 8Gb	1333 MHz		1333 MHz		1067 MHz	
2RX4, 4Gb, 1600 MHz / 16Gb	1333 MHz		1333 MHz		1067 MHz	
4RX4, 4Gb, 1600 MHz / 32Gb	1333 MHz		1333 MHz		1333 MHz	
8Rx4, 4Gb, 1333 MHz / 64Gb	1067 MHz		1067 MHz		1067 MHz	
<b>DDR3, RAS mode (1:1): Minimal</b>						

DIMM type / capacity	1 DPC 1.35 V	1 DPC 1.5 V	2 DPC 1.35 V	2 DPC 1.5 V	3 DPC 1.35 V	3 DPC 1.5 V
1RX4, 2Gb, 1600 MHz / 4Gb	1067 MHz		1067 MHz		1067 MHz	
1RX4, 4Gb, 1600 MHz / 8Gb	1067 MHz		1067 MHz		1067 MHz	
2RX4, 4Gb, 1600 MHz / 16Gb	1067 MHz		1067 MHz		1067 MHz	
4RX4, 4Gb, 1600 MHz / 32Gb	1067 MHz		1067 MHz		1067 MHz	
8Rx4, 4Gb, 1333 MHz / 64Gb	1067 MHz		1067 MHz		1067 MHz	

Table 2-4 DDR4 memory speeds in Performance mode

DIMM type / capacity	1 DPC 1.2 V	2 DPC 1.2 V	3 DPC 1.2 V
<b>TruDDR4, Performance mode (2:1): Maximum performance</b>			
1RX4, 4Gb, 2133 MHz / 8Gb	1600 MHz	1600 MHz	1600 MHz
2RX4, 4Gb, 2133 MHz / 16Gb	1600 MHz	1600 MHz	1600 MHz
2RX4, 8Gb, 2133 MHz / 32Gb	1600 MHz	1600 MHz	1600 MHz
4RX4, 16Gb, 2133 MHz / 64Gb	1600 MHz	1600 MHz	1600 MHz
<b>TruDDR4, Performance mode (2:1): Balanced</b>			
1RX4, 4Gb, 2133 MHz / 8Gb	1333 MHz	1333 MHz	1333 MHz
2RX4, 4Gb, 2133 MHz / 16Gb	1333 MHz	1333 MHz	1333 MHz
2RX4, 8Gb, 2133 MHz / 32Gb	1333 MHz	1333 MHz	1333 MHz
4RX4, 16Gb, 2133 MHz / 64Gb	1333 MHz	1333 MHz	1333 MHz
<b>TruDDR4, Performance mode (2:1): Minimal</b>			
1RX4, 4Gb, 2133 MHz / 8Gb	1333 MHz	1333 MHz	1333 MHz
2RX4, 4Gb, 2133 MHz / 16Gb	1333 MHz	1333 MHz	1333 MHz
2RX4, 8Gb, 2133 MHz / 32Gb	1333 MHz	1333 MHz	1333 MHz
4RX4, 16Gb, 2133 MHz / 64Gb	1333 MHz	1333 MHz	1333 MHz

Table 2-5 DDR4 memory speeds in RAS mode

DIMM type / capacity	1 DPC 1.2 V	2 DPC 1.2 V	3 DPC 1.2 V
<b>TruDDR4, RAS mode (1:1): Maximum performance</b>			
1RX4, 4Gb, 2133 MHz / 8Gb	1867 MHz	1867 MHz	1600 MHz
2RX4, 4Gb, 2133 MHz / 16Gb	1867 MHz	1867 MHz	1600 MHz
2RX4, 8Gb, 2133 MHz / 32Gb	1867 MHz	1867 MHz	1600 MHz
4RX4, 16Gb, 2133 MHz / 64Gb	1867 MHz	1867 MHz	1600 MHz
<b>TruDDR4, RAS mode (1:1): Balanced</b>			
1RX4, 4Gb, 2133 MHz / 8Gb	1600 MHz	1600 MHz	1333 MHz

DIMM type / capacity	1 DPC 1.2 V	2 DPC 1.2 V	3 DPC 1.2 V
2RX4, 4Gb, 2133 MHz / 16Gb	1600 MHz	1600 MHz	1333 MHz
2RX4, 8Gb, 2133 MHz / 32Gb	1600 MHz	1600 MHz	1333 MHz
4RX4, 16Gb, 2133 MHz / 64Gb	1600 MHz	1600 MHz	1333 MHz
<b>TruDDR4, RAS mode (1:1): Minimal</b>			
1RX4, 4Gb, 2133 MHz / 8Gb	1333 MHz	1333 MHz	1333 MHz
2RX4, 4Gb, 2133 MHz / 16Gb	1333 MHz	1333 MHz	1333 MHz
2RX4, 8Gb, 2133 MHz / 32Gb	1333 MHz	1333 MHz	1333 MHz
4RX4, 16Gb, 2133 MHz / 64Gb	1333 MHz	1333 MHz	1333 MHz

## 2.4.2 Memory mirroring and rank sparing

In addition to Performance and RAS modes, the memory subsystem has the following RAS features that can be enabled from the UEFI:

- ▶ Memory mirroring
- ▶ Rank sparing

### Memory mirroring

To improve memory reliability and availability, the memory controller can mirror memory data across two memory channels. To enable the mirroring feature, you must have both memory channels of a processor populated with the same DIMM type and amount of memory.

Memory mirroring provides the user with a redundant copy of all code and data addressable in the configured memory map. Two copies of the data are kept, similar to the way RAID-1 writes to disk. Reads are interleaved between memory channels. The system automatically uses the most reliable memory channel as determined by error logging and monitoring.

If errors occur, only the alternative memory channel is used until bad memory is replaced. Because a redundant copy is kept, mirroring results in only half the installed memory being available to the operating system. Memory mirroring does not support asymmetrical memory configurations and requires that each channel be populated in identical fashion. For example, you must install two identical 4 GB 2133MHz DIMMs equally and symmetrically across the two memory channels to achieve 4 GB of mirrored memory.

Memory mirroring is a hardware feature that operates independent of the operating system. There is a slight memory performance trade-off when memory mirroring is enabled.

The memory mirroring feature can be used with performance or RAS modes:

- ▶ When Performance mode is used, memory mirroring duplicates data between memory *channels* of the two memory buffers connected to one memory controller.
- ▶ In RAS (Lockstep) mode, memory mirroring duplicates data between memory *buffers* that are connected to the same memory controller.



Figure 2-25 shows how memory mirroring is implemented in Performance mode (left) and RAS mode (right).

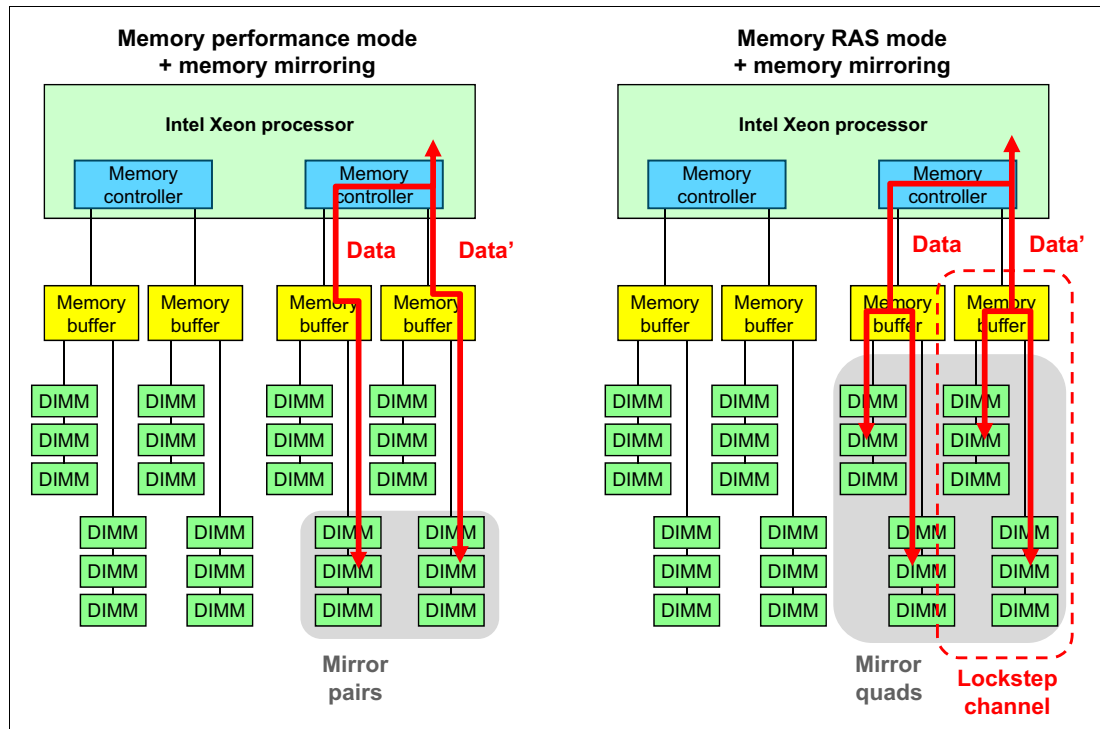


Figure 2-25 Memory mirroring with used with Performance mode (left) and RAS mode (right)

The following memory mirroring rules apply:

- ▶ The server supports single-socket memory mirroring. The Compute Book memory channel 0 mirrors memory channel 1, and memory channel 2 mirrors memory channel 3. This mirroring provides redundancy in memory but reduces the total memory capacity by half.
- ▶ DIMMs must be installed in pairs for each Compute Book when the memory mirroring feature is used.
- ▶ The DIMM population must be identical (capacity, organization, and so on) for memory channel 0 and memory channel 1, and identical for memory channel 2 and memory channel 3.
- ▶ Memory mirroring reduces the maximum available memory by half of the installed memory. For example, if the server has 64 GB of installed memory, only 32 GB of addressable memory is available when memory mirroring is enabled.

### Memory Address Range Mirroring

Intel Xeon E7 v3 and v4 processors provide an advanced version of the memory mirroring feature named Memory Address Range Mirroring. Unlike basic memory mirroring, when whole memory channels or memory buffers are mirrored, memory address range mirroring designates a memory *region* to mirror. You can choose a range of memory addresses to be mirrored, which allows you to keep critical data secured and save large amounts of memory at the same time. Each processor supports up to two mirror ranges, each mirror range size uses 64 MB granularity. For additional information on Memory Address Range Mirroring, refer to the following paper:

<http://software.intel.com/en-us/articles/reliability-availability-and-serviceability-integration-and-validation-guide-for-the-intel>

## Rank sparing

In rank-sparing mode, one rank is held in reserve as a spare for the other ranks in the same memory channel. There are eight memory channels per processor.

Memory rank sparing provides a degree of redundancy in the memory subsystem, but not to the extent of mirroring. In contrast to mirroring, sparing leaves more memory for the operating system. In sparing mode, the trigger for failover is a preset threshold of correctable errors. When this threshold is reached, the content is copied to its spare rank. The failed rank is then taken offline, and the spare counterpart is activated for use.

In rank sparing mode, one rank per memory channel is configured as a spare. The spare rank must have identical or larger memory capacity than all the other ranks (sparing source ranks) on the same channel.

For example, if dual-rank DIMMs are installed and are all of the same capacity, there are six ranks total for each memory channel (three DIMMs per channel). This configuration means that one of the six ranks are reserved and five of the six ranks can be used for the operating system.

Memory sparing is a hardware feature that operates independent of the operating system. There is a slight memory performance trade-off when memory sparing is enabled.

The rank sparing feature can be used in addition to performance or RAS modes. Consider the following points:

- ▶ When Performance mode is used, rank sparing duplicates data between memory modules of the *same channel* of one memory buffer. If there is an imminent failure (as indicated by a red X in Figure 2-26 on page 41), that rank is taken offline and the data is copied to the spare rank.
- ▶ When RAS (Lockstep) mode is used, rank sparing duplicates data between memory channels of one memory buffer. If there is an imminent failure (as indicated by a red X in Figure 2-26 on page 41), that rank is taken offline and the data is copied to the spare rank. In addition, the partner rank on the other channel that is connected to the same memory buffer also is copied over.

Figure 2-26 shows the rank sparing usage with Performance mode (left) and RAS mode (right).




Figure 2-26 Rank sparing: Performance mode (left) and RAS mode (right)

The following configuration rules apply for rank sparing:

- ▶ Memory rank sparing is not supported if memory mirroring is enabled.
- ▶ The spare rank must have identical or larger memory capacity than all the other ranks on the same memory channel.
- ▶ When single-rank DIMMs are used, a minimum of *two* DIMMs must be installed per memory channel to support memory sparing.
- ▶ When multirank DIMMs are used, *one* multirank DIMM can be installed per memory channel to support memory sparing.
- ▶ The total memory available in the system is reduced by the amount of memory that is allocated for the spare ranks.

### Multiple rank sparing

Intel Xeon E7 v3 and v4 processors provide an advanced version of rank sparing called multiple rank sparing. It is now possible to specify more than one rank as a hot spare. When the memory errors threshold is reached, the content of the failed rank is copied to its spare rank. The failed rank is then taken offline and the spare counterpart is activated for use. If another memory rank fails, the same procedure triggers again to take over to the next standby rank, and so on.

Note that the spare rank(s) must have memory capacity identical to, or larger than, all the other ranks (sparing source ranks). The total memory available in the system is reduced by the amount of memory allocated for the spare ranks.

### 2.4.3 Chipkill

Chipkill memory technology, which is an advanced form of error correction code (ECC), is available on X6 servers. Chipkill protects the memory in the system from any single memory chip failure. It also protects against multibit errors from any portion of a single memory chip.

Chipkill on its own can provide 99.94% memory availability to the applications without sacrificing performance and with standard ECC DIMMs.

Chipkill is used in Performance and RAS modes.

### 2.4.4 Redundant bit steering

Redundant bit steering provides the equivalent of a hot-spare drive in a RAID array. It is based in the memory controller and senses when a chip on a DIMM fails and when to route the data around the failed chip.

Redundant bit steering is used in *RAS mode only*.

X6 servers support the Intel implementation of Chipkill plus redundant bit steering, which Intel refers to as DDDC.

Redundant bit steering uses the ECC coding scheme that provides Chipkill coverage for x4 DRAMs. This coding scheme leaves the equivalent of one x4 DRAM spare in every pair of DIMMs. If a chip failure on the DIMM is detected, the memory controller can copy data from the failed chip through the spare x4.

Redundant bit steering operates automatically without issuing a Predictive Failure Analysis (PFA) or light path diagnostics alert to the administrator, although an event is logged to the service processor log. After the second DRAM chip failure on the DIMM in RAS (Lockstep) mode, more single bit errors result in PFA and light path diagnostics alerts.

### 2.4.5 Advanced Page Retire

Advanced Page Retire is an algorithm to handle memory errors. It is built-in sophisticated error-handling firmware that uses and coordinates memory recovery features, which balances the goals of maximum up time and minimum repair actions.

The algorithm uses short- and long-term thresholds per memory rank with leaky bucket and automatic sorting of memory pages with the highest correctable error counts. First, it uses hardware recovery features, followed by software recovery features, to optimize recovery results for newer and older operating systems and hypervisors.

When recovery features are exhausted, the firmware issues a Predictive Failure Alert. Memory that failed completely is held offline during starts until it is repaired. Failed DIMMs are indicated by light path diagnostics LEDs that are physically at the socket location.

## 2.5 PCIe 3.0

The X6 family servers supports the latest generation of PCI Express (PCIe) protocol, Version 3.0. PCIe 3.0 is the evolution of PCI Express I/O standard that brings doubled bandwidth over PCIe 2.0 and preserves compatibility with previous generations of PCIe protocol. PCIe 1.x and 2.x cards work in PCIe 3.0-capable slots, and PCIe 3.0 cards work when they are plugged into PCIe slots of previous generations.

PCIe 3.0 uses a 128b/130b encoding scheme, which is more efficient than the 8b/10b encoding that is used in PCIe 2.0 protocol. This approach reduces overhead to less than 2% when compared to 20% of PCIe 2.0. The bandwidth also can be doubled at 8 GTps speed.

Each processor contains an Integrated I/O (IIO) module that provides 32 lanes of PCIe 3.0. These 32 lanes can be split into any combination of x4, x8, and x16.

Table 2-6 shows a comparison of the PCIe capabilities of the eX5 and X6 families.

Table 2-6 X5 and X6 PCIe capabilities comparison

Specification	x3850 eX5	x3850 X6
PCIe version	PCIe 2.0	PCIe 3.0
Raw bit rate	5 GTps	8 GTps
Encoding	8b/10b	128b/130b
Bandwidth per lane	~500 MBps	~1 GBps
Available PCIe slots	8	12
Number of x16 slots	1	Up to 6 <sup>a</sup>
Hot-swap support	No	Yes <sup>b</sup>

a. Six PCIe x16 slots are available for configurations with two full-length I/O Books.

b. Hot-swap capabilities are available for additional I/O Books only.

For more information about PCIe 3.0, see the PCI Express Base 3.0 specification by PCI-SIG, which is available at this website:

<http://www.pcisig.com/specifications/pciexpress/base3/>

For more information about the implementation of the I/O subsystem, see Chapter 3, “Product information” on page 53

## 2.6 Internal storage

In this section, we give an overview of the following technologies that are used in the storage subsystem of the X6 servers family:

- ▶ 2.6.1, “Storage Book” on page 44
- ▶ 2.6.2, “SAS/SATA SSD technology” on page 44
- ▶ 2.6.3, “NVMe SSD technology” on page 45
- ▶ 2.6.4, “Flash Storage Adapters” on page 47

## 2.6.1 Storage Book

The x3850 X6 and x3950 X6 servers support 1.8-inch SSDs or 2.5-inch SSDs and HDDs. The drives are installed in the Storage Book, as shown in Figure 2-27:




Figure 2-27 The X6 Storage Book

In addition to the drive bays, a Storage Book contains two PCIe 3.0 x8 slots for internal RAID controllers or host bus adapters (HBAs).

The X6 servers bring a new generation of SAS protocol: 12 Gb SAS. It doubles the data transfer rate of 6 Gb SAS solutions, to fully unlock the potential of the PCIe 3.0 interface and to maximize performance for storage I/O-intensive applications, including databases, business analytics, and virtualization and cloud environment.

In addition to or as an alternative to the use of SAS or SATA drives, the X6 servers support the use of PCIe NVMe SSDs. These drives connect directly to the PCIe bus of the processors and provide the ultimate in storage bandwidth and latency while still in a drive form factor.

## 2.6.2 SAS/SATA SSD technology

Each SSD unit can accommodate up to four 2.5-inch or eight 1.8-inch hot-swap front-accessible SSDs, and the unit occupies the space of four 2.5-inch SAS HDD bays. You can install the following number of SSD units:

- ▶ The x3850 X6 can have up to 8x 2.5-inch SSDs with up to two SSD units (up to four SSDs per unit).
- ▶ The x3850 X6 can have up to 16x 1.8-inch SSDs with up to two SSD units (up to eight SSDs per unit).
- ▶ The x3950 X6 can have up to 16x 2.5-inch SSDs with up to four SSD units.
- ▶ The x3950 X6 can have up to 32x 1.8-inch SSDs with up to four SSD units.

SSD solutions can provide the following benefits:

- ▶ Significantly lower implementation cost of high performance I/O-intensive storage systems.

- ▶ Significantly higher performance and better response time of storage-intensive applications with up to 10x less response time comparing with solutions that are based on spinning HDDs.
- ▶ Significant savings in power and cooling with high performance-per-watt ratio.

SSDs are optimized for a heavy mix of random read and write operations, such as transaction processing, data mining, business intelligence, and decision support, and other random I/O-intensive applications. Built on enterprise-grade MLC NAND flash memory, the SSD drives used in the X6 systems deliver up to 30,000 IOPS per single drive. Combined into an SSD unit, these drives can deliver up to 240,000 IOPS and up to 2 GBps of sustained read throughput *per SSD unit*. In addition to its superior performance, SSD offers superior uptime with three times the reliability of mechanical disk drives because SSDs have no moving parts to fail.

SSD units require the following components:

- ▶ Hot-swap SAS SSD backplane
- ▶ 2.5-inch or 1.8-inch SSDs
- ▶ High-performance disk controllers

### 2.6.3 NVMe SSD technology

Non-Volatile Memory Express (NVMe) is new PCIe 3.0 high performance SSD technology that provides high I/O throughput and low latency. NVMe PCI SSDs are similar to the PCIe Flash Adapter (discussed in 2.6.4, “Flash Storage Adapters” on page 47), both of which share the same NVMe technology, NVMe interfaces remove SAS/SATA bottlenecks and enable all of the capabilities of contemporary NAND flash memory. Each NVMe PCI SSD has a direct PCIe 3.0 x4 connection, which provides at least 2x the bandwidth with *half* the latency of SATA/SAS-based SSD solutions. NVMe drives are also optimized for heavy multithreaded workloads by using internal parallelism and many other improvements, such as enlarged I/O queues.

NVMe technology has the following key characteristics:

- ▶ Direct PCIe 3.0 connection. There is a PCIe 3.0 x4 connection for each NVMe drive with up to 4 GBps overall throughput.
- ▶ Low I/O latency. The average read/write latency for the P3700 drives is 20  $\mu$ s.
- ▶ Up to 2800 MBps sequential read speed with 128 KB blocks, and up to 2000 MBps sequential write speed with 128 KB blocks per drive.
- ▶ Up to 460,000 IOPS of random *reads* with 4 KB blocks, and up to 175,000 IOPS of random writes with 4 KB blocks.
- ▶ A total of 65,536 I/O queues supported and 65,536 commands per queue supported, which provides great performance for heavily multithreaded workloads with combined sequential and random access.
- ▶ High endurance: The P3700 drives include High Endurance Technology (HET) which combines NAND silicon enhancements and advanced SSD NAND management techniques to extend SSD write endurance up to 17 full drive writes per day (DWPD) for 5 years.
- ▶ Available drive capacities of 400 GB, 800 GB, 1.6 TB, and 2.0 TB.
- ▶ Support for software RAID under operating system management.
- ▶ Hot-add and hot-remove features are available with supported operating systems.

- ▶ Most operating systems have native support of NVMe drives or provide support through software drivers, including RHEL 6.5 and later, RHEL 7.0 and later, SLES 11 SP3, SLES 12, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, and VMware vSphere (ESXi) 5.5 and 6.0.
- ▶ NVMe drives can be used as boot drives.

NVMe drives are available in 2.5-inch drive form-factor compatible with the X6 Storage Book, but require a special NVMe backplane and PCIe extender adapters that are installed in Storage Book PCIe 3.0 x8 slots.

Figure 2-28 shows the PCIe NVMe SSD.



Figure 2-28 NVMe PCIe SSD

Each NVMe PCIe extender supports one or two NVMe drives. You can install up to two NVMe PCI extenders in each Storage Book, with which up to four NVMe PCIe drives can be used in one Storage Book. As a result, the drive maximums per server are possible:

- ▶ x3850 X6: four NVMe drives per server (one Storage Book)
- ▶ x3950 X6: eight NVMe drives per server (two Storage Books)

The extender adapter is shown in Figure 2-29.





Figure 2-29 NVMe PCIe SSD Extender Adapter

If you use four NVMe drives in one Storage Book and two NVMe PCI extenders accordingly, no more PCIe slots are available for RAID-adapters in that Storage Book. In this case, only NVMe drives can be installed in the Storage Book and the other four 2.5-inch drive bays must be left empty.

For more information about the available combinations of drives in the Storage Book, see 3.11.1, “Storage Book” on page 94.

**Note:** Each pair of NVMe PCIe drives require one NVMe PCIe extender, which is installed in the Storage Book. Each NVMe PCIe SSD uses four PCIe 3.0 lanes; therefore, a pair of NVMe PCIe drives completely uses the bandwidth of one available PCIe 3.0 x8 slot in the Storage Book. One NVMe PCIe extender is needed for every two NVMe drives.

## 2.6.4 Flash Storage Adapters

Flash Storage Adapters (also known as High IOPS adapters or Flash Adapters) provide high-performance storage based on SSD technology, but in a regular PCIe adapter form factor. Designed for high-performance servers and computing appliances, these adapters deliver read/write throughput of up to 2.8 GBps or 2.2 GBps and the added benefits of lower power, cooling, and management overhead and a smaller storage footprint. The maximum capacity of Flash Storage Adapters is up to 6.4 TB.

The Flash Storage Adapter options include the following:

- ▶ NVMe Enterprise Performance Flash Adapters
- ▶ io3 Enterprise Performance Flash Adapters
- ▶ io3 Enterprise Mainstream Flash Adapters

The Flash Storage Adapters combine high IOPS performance with low latency. As an example, with 4 KB block random reads, the 3.2 TB io3 Enterprise Mainstream Flash Adapter can read 345,000 IOPS, and the 2.0 TB P3700 NVMe Enterprise Performance Flash Adapter can read 400,000 IOPS, compared with 420 IOPS for a 15 K RPM 146 GB disk drive. The

Enterprise Mainstream adapter can also write 385,000 IOPS with 4 KB random writes while the NVMe Enterprise Performance Flash Adapter can write 150,000 IOPS with 4 KB random writes. The read access latency is approximately 79 microseconds for the Enterprise Mainstream Flash Adapter, which is about 1% of the latency of a 15K RPM 146 GB disk drive. The write access latency is even lower, at about 15 microseconds.

Reliability features include the use of enterprise-grade MLC (eMLC), advanced wear-leveling, ECC protection, and Adaptive Flashback redundancy for RAID-like chip protection with self-healing capabilities, which provides unparalleled reliability and efficiency.

Advanced bad-block management algorithms enable taking blocks out of service when their failure rate becomes unacceptable. These reliability features provide a predictable lifetime and up to 25 years of data retention.

Figure 2-30 shows the 6.4TB io3 Enterprise Mainstream Flash Storage Adapter:




Figure 2-30 6.4 TB io3 Enterprise Mainstream Flash Storage Adapter

Figure 2-31 on page 48 shows the P3700 NVMe Enterprise Performance Flash Adapter:



Figure 2-31 P3700 Enterprise Performance Flash Adapter

For more information about SSD Adapters that are supported in the X6 servers, see the Lenovo Press Product Guides in the PCIe SSD Adapters category:

<http://lenovopress.com/servers/options/ssdadapter>

## 2.7 Unified Extensible Firmware Interface

The UEFI is a prestart environment that provides an interface between server firmware and the operating system. It replaces BIOS as the software that manages the interface between server firmware, operating system and hardware initialization, and eliminates the 16-bit, real-mode limitation of the BIOS.

The UEFI includes the following main advantages:

- ▶ It uses 64-bit code architecture, with no restrictions for code size.
- ▶ It offers a modular design; adapter vendors can add any functions in ROM.
- ▶ An adapter configuration can move into F1 Setup.
- ▶ There are no beep codes; all errors are covered by light path diagnostics.
- ▶ There are human-readable Event logs in F1 Setup.
- ▶ UEFI Settings can be accessed Out of Band via the Advanced Setup Utility (ASU) through the IMM web interface.
- ▶ UEFI Code Updates can be performed in band through the operating system, and out of band through the IMM web interface.

For more information about UEFI, see this website:

<http://www.uefi.org/home/>

## 2.8 Integrated Management Module

The x3850 X6 and x3950 X6 have an IMM2 service processor. The specific version of the IMM2 used in the X6 servers is IMM Version 2.1. The IMM2 provides advanced control and monitoring features to manage an X6 server. IMM V2.1 (referred to in this book as IMM2) enables easy console redirection with text and graphics, and keyboard and mouse support over the system management LAN connections.

By using this feature, the user can display server activities from power-on to full operation remotely, with remote user interaction, at virtually any time.

IMM2 monitors the following components:

- ▶ System voltages
- ▶ System temperatures
- ▶ Fan speed control
- ▶ Fan tachometer monitor
- ▶ Good Power signal monitor
- ▶ System ID and planar version detection
- ▶ System power and reset control
- ▶ Non-maskable interrupt (NMI) detection (system interrupts)
- ▶ SMI detection and generation (system interrupts)
- ▶ Serial port text console redirection
- ▶ System LED control (power, HDD, activity, alerts, and heartbeat)

IMM2 provides the following features:

- ▶ An embedded web server, which gives you remote control from any standard web browser with the Oracle Java plug-in installed.

- ▶ Shared or dedicated Ethernet port. Support for LAN over USB for in-band communications to the IMM.
- ▶ A command-line interface (CLI), which the administrator can use from a Telnet or SSH session.
- ▶ Secure Sockets Layer (SSL) and Lightweight Directory Access Protocol (LDAP).
- ▶ Built-in LAN and serial connectivity that support virtually any network infrastructure.
- ▶ Multiple alerting functions with access to VPD/PFA to warn systems administrators of potential problems through email, IPMI platform event traps (PETs), and Simple Network Management Protocol (SNMP).
- ▶ Event logs that are time stamped, saved on the IMM, and that can be attached to e-mail alerts
- ▶ User authentication using a secure connection to a Lightweight Directory Access Protocol (LDAP) server

IMM2 has many enhancements compared to the previous generation of IMM:

- ▶ Faster processor and more memory, which delivers more responsive user interface
- ▶ Improved system power-on and start time
- ▶ New web user interface that provides a common design across all System x products
- ▶ Increased color depth and resolution for remote presence function
- ▶ IMM2 configuration changes requires no restart and becomes effective immediately
- ▶ Hardware management of non-volatile storage
- ▶ 1 Gb Ethernet management capability
- ▶ More detailed information for UEFI detected events
- ▶ Simplified update and flashing mechanism
- ▶ Security enhancements, including authentication policies

The following table represents the main differences between the IMM2 and the IMM2.1.

Table 2-7 IMM2 vs IMM2.1

	IMM2	IMM2.1
Management controller	Renesas SH7757	Renesas SH7758
Intel Chipset family	Romley, Brickland	Grantley
System power management	RTMM (H8S-2117), IBM AME	FPGA, Intel ME
DCMI 1.5 support	No	Yes
Intel SPS 3.0 support	No	Yes
Lenovo efficiency mode support	No	Yes

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

<https://support.lenovo.com/docs/UM103336>

## 2.9 Scalability

The X6 servers have a flexible modular design with which you can increase the server's compute power and I/O capabilities by adding Compute Books and I/O Books. You can build your initial configuration of an x3850 X6 server with one Compute Book, one Primary I/O

Book, and one Storage Book, and ultimately reuse these components in an 8U x3950 X6 server that consists of eight Compute Books, two primary I/O Books, four other I/O Books, and two Storage Books.

Table 2-8 shows a comparison of minimum and maximum configurations of X6 servers, from an entry level x3850 X6 configuration to a fully populated x3950 X6.

*Table 2-8 Scalability capabilities of X6 servers*

	<b>Minimum x3850 X6 configuration</b>	<b>Maximum x3950 X6 configuration</b>
Number of processors	1 E7-4800 v2	8 E7-8800 v4
Memory modules	2 DDR3	192 TruDDR4
Total memory capacity	8 GB	12 TB
PCIe slots	6	24
Internal storage drives	4 2.5-inch HDDs	32 1.8-inch SSDs

The x3850 X6 and x3950 X6 servers use native QPI scaling capabilities to achieve 4-socket and 8-socket configurations. Unlike eX5 systems, there are no external connectors and cables for X6 systems; all interconnects are integrated in the midplane.

Lenovo supports upgrading a 4-socket X6 server to an 8-socket X6 server. The following methods are recommended:

- ▶ Start with a 4-socket x3950 X6 (8U chassis) and add processors when needed.
- ▶ Start with a 4-socket x3850 X6 (4U chassis) and upgrade when needed by using an RPQ process.

For more information about the upgrading process, see 3.26, “Upgrading to an 8-socket X6 server” on page 125.



# Product information

In this chapter, we describe the components of the Lenovo System x3850 X6 and x3950 X6.

This chapter includes the following topics:

- ▶ 3.1, “Product features” on page 54
- ▶ 3.2, “Specifications” on page 57
- ▶ 3.3, “Standard models of X6 servers” on page 60
- ▶ 3.4, “SAP HANA Workload Optimized Solution models” on page 64
- ▶ 3.5, “Physical design” on page 68
- ▶ 3.6, “Ports and controls” on page 70
- ▶ 3.7, “Compute Books” on page 75
- ▶ 3.8, “Processor options” on page 80
- ▶ 3.9, “Memory” on page 84
- ▶ 3.10, “Memory-channel storage” on page 92
- ▶ 3.11, “Storage subsystem” on page 94
- ▶ 3.12, “I/O subsystem” on page 100
- ▶ 3.13, “Primary I/O Book” on page 102
- ▶ 3.14, “Half-length I/O Books and Full-length I/O Books” on page 104
- ▶ 3.15, “Hot-swap adapter support” on page 108
- ▶ 3.16, “Network adapters” on page 109
- ▶ 3.18, “Fibre Channel host bus adapters” on page 114
- ▶ 3.19, “GPU adapters and co-processors” on page 114
- ▶ 3.20, “Partitioning” on page 116
- ▶ 3.21, “Standard onboard features” on page 116
- ▶ 3.22, “Integrated virtualization” on page 120
- ▶ 3.23, “Hot-swap capabilities” on page 121
- ▶ 3.24, “Power subsystem” on page 122
- ▶ 3.25, “Fans and cooling” on page 124
- ▶ 3.26, “Upgrading to an 8-socket X6 server” on page 125

## 3.1 Product features

The increasing demand for cloud computing and analytics workloads by enterprises to meet social, mobile, and big data requirements drives innovation to find new ways to build informational systems. Clients are looking for cost-optimized fit-for-purpose IT solutions that manage large amounts of data, easily scale performance, and provide enterprise class reliability.

Built on decades of innovation, Lenovo introduces its sixth generation of Enterprise X-Architecture technology, Lenovo X6 servers. X6 servers are fast, agile, and resilient. Consider the following points:

- ▶ *Fast* application performance means immediate access to actionable information.
- ▶ *Agile* system design helps to reduce acquisition costs and upgrades processor and memory technology at each refresh within the same chassis.
- ▶ *Resilient* platforms maximize application uptime and promote easy integration in virtual environments.

X6 servers continue to lead the way as the shift toward mission-critical scalable databases, business analytics, virtualization, enterprise applications, and cloud applications accelerates.

### 3.1.1 Fast application performance

The server offers the following features to boost performance:

- ▶ Based on the Intel Xeon E7 family of processors:
  - Models based on E7-4800 v4 and E7-8800 v4 processors (Also available are models based on E7 v3 processors)
  - Supports up to four E7 v4 processors with a total of 96 cores and 192 threads to maximize the concurrent running of multi-threaded applications.
  - Supports processors with up to 24 cores, core frequencies up to 3.2 GHz, L3 cache of up to 60 MB, and three QPI interconnect links at up to 9.6 GTps.
- ▶ It supports memory speeds up to 1866 MHz for DDR4 memory and up to 1600 MHz for DDR3 memory (model dependent).
- ▶ It supports up to 96 DIMM sockets (x3850 X6) or 192 DIMM sockets (x3950 X6), 24 DIMMs per processor.
- ▶ Intelligent and adaptive system performance with Intel Turbo Boost Technology 2.0 allows CPU cores to run at maximum speeds during peak workloads by temporarily going beyond processor TDP.
- ▶ Intel Hyper-Threading Technology boosts performance for multi-threaded applications by enabling simultaneous multi-threading 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.
- ▶ Intel Advanced Vector Extensions (AVX) significantly improve floating-point performance for compute-intensive technical and scientific applications.
- ▶ Supports 12 Gbps SAS RAID portfolio.
- ▶ The use of solid-state drives (SSDs) instead of, or along with, traditional spinning hard disk drives (HDDs) can significantly improve I/O performance. An SSD can support up to 100 times more I/O operations per second (IOPS) than a typical HDD.



- ▶ Support for PCIe Non-Volatile Memory Express (NVMe) storage. These disks are SSDs that use the high-speed PCIe 3.0 interface to help increase IOPS and lower latency.
- ▶ PCI Express 3.0 I/O adapter slots that improve the theoretical maximum bandwidth by almost 100% (8 GTps per link by using 128b/130b encoding) compared to the previous generation of PCI Express 2.0 (5 GTps per link by using 8b/10b encoding).
- ▶ With Intel Integrated I/O Technology, the PCI Express 3.0 controller is integrated into the Intel Xeon processors. This integration helps to dramatically reduce I/O latency and increase overall system performance.
- ▶ Support for up to two graphics processing units (GPUs) and co-processors to maximize computing power.
- ▶ Energy-efficient electronic components help lower operational costs, including highly efficient 900 W AC and 1400 W AC power supplies with 80 PLUS Platinum certification.

### 3.1.2 Agile system design

The server provides the following scalability and flexibility features:

- ▶ Innovative module “book” design for each of the three subsystems: Compute Books, Storage Book, and I/O Books. Front and rear access means that you can easily scale the system by adding components without removing the entire server from the rack.
- ▶ The modular book design also allows clients to create the configuration that fits their application and environment needs, which reduces acquisition costs while giving them the flexibility to grow and modify their configuration later.
- ▶ The book design also means that subsystem upgrades are simpler, quicker to perform, and have less of an effect on the rest of the server.
- ▶ By using 64 GB LRDIMMs, it supports a large amount of internal memory:
  - x3850 X6: Supports up to 6 TB
  - x3950 X6: Supports up to 12 TB
- ▶ Internal storage bays provide a flexible and scalable all-in-one platform to meet your increasing demands:
  - x3850 X6: Up to 8x 2.5-inch HDD/SSD bays, or up to 16x 1.8-inch SSD bays
  - x3950 X6: Up to 16x 2.5-inch HDD/SSD bays, or up to 32x 1.8-inch SSD bays
- ▶ Offers up to 11 PCIe slots plus a dedicated Mezzanine LOM (ML2) adapter slot; most slots are PCIe 3.0 to maximize I/O scalability.
- ▶ PCIe slots are implemented in I/O Books to maximize modularity. Choose from Half-length I/O Books or Full-length I/O Books, depending on the adapters you need to deploy.
- ▶ Most components are common between the four-socket x3850 X6 and eight-socket x3950 X6, which makes for a simple upgrade path with minimal parts on the floor.

### 3.1.3 Resilient platform

The server provides the following features to simplify serviceability and increase system uptime:

- ▶ Upward Integration Modules for standard hypervisors enable the creation and management of policies to maintain high availability of virtual machines and concurrent updating of the system firmware, with no effect on application performance or availability.

- ▶ Advanced Processor Recovery allows the system to automatically switch access and control of networking, management, and storage if there is a processor 1 failure, which provides higher availability and productivity.
- ▶ Advanced Page Retire proactively protects applications from corrupted pages in memory, which is crucial for scaling memory to terabytes.
- ▶ Redundant bit steering, memory mirroring, and memory rank sparing provide redundancy if there is a non-correctable memory failure.
- ▶ Intel Execute Disable Bit functionality can help 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, which allows an application to run in its own isolated space that is protected from all other software that is running on a system.
- ▶ Redundant Intel Platform Controller Hub (PCH) connections to the processors allow the platform to maintain access to networking, storage, and server management during a processor failure.
- ▶ Hot-swap drives support RAID redundancy for data protection and greater system uptime.
- ▶ By using hot-swap I/O Books, you can install or replace adapters while the server is still running.
- ▶ Tool-less, lid-less design provides front and rear access for easy upgrades and serviceability; no need to pull the server out of the rack to access internal components.
- ▶ Hot-swap power supplies and hot-swap dual-motor redundant fans to provide availability for mission-critical applications.
- ▶ A new LCD diagnostics panel combined with individual light path diagnostic LEDs quickly lead the technician to failed (or failing) components, which simplifies servicing, speeds up problem resolution, and helps improve system availability.
- ▶ Predictive Failure Analysis (PFA) detects when system components (processors, memory, HDDs, SSDs, fans, and power supplies) operate outside of standard thresholds and generates proactive alerts in advance of a possible failure, which increases uptime.
- ▶ Built-in Integrated Management Module Version II (IMM2) continuously monitors system parameters, triggers alerts, and performs recovering actions if there are failures to minimize downtime.
- ▶ Includes a special Mezzanine LOM (ML2) adapter slot with support for adapters with two 10 Gb ports or 4 Gigabit ports; supports direct connectivity to the IMM2 service processor for out-of-band systems management.
- ▶ Integrated industry-standard Unified Extensible Firmware Interface (UEFI) enables improved setup, configuration, and updates, and simplifies error handling.
- ▶ Two integrated Trusted Platform Modules (TPMs) 1.2 support enables advanced cryptographic functionality, such as digital signatures and remote attestation.
- ▶ Industry-standard Advanced Encryption Standard (AES) NI supports faster and stronger encryption.
- ▶ Lenovo XClarity provides proactive systems management. It offers comprehensive systems management that helps to increase uptime, reduce costs, and improve productivity through advanced server management capabilities.
- ▶ Lenovo XClarity Energy Manager (LXEM) provides data center energy management by monitoring power and temperature at the server level and the group level to help data center administrators improve business continuity and energy efficiency.
- ▶ SSDs offer better reliability than HDDs for greater uptime.

- ▶ Built-in diagnostics that use Dynamic Systems Analysis (DSA) Preboot speed up troubleshooting tasks to reduce service time.
- ▶ A three-year customer-replaceable unit and onsite limited warranty, 9x5 next business day, is included. Optional service upgrades are available.

## 3.2 Specifications

Table 3-1 lists the standard specifications.

Table 3-1 Specifications

Components	Specification
Form factor	<ul style="list-style-type: none"> <li>▶ x3850 X6: 4U rack</li> <li>▶ x3950 X6: 8U rack</li> </ul>
Processor	<ul style="list-style-type: none"> <li>▶ x3850 X6: One, two, or four Intel Xeon E7-4800 v4 or E7-8800 v4 processors, each in a Compute Book. Processor options between 4 cores (3.2 GHz) and 24 cores (up to 2.7 GHz). Three QPI links up to 9.6 GTps each. Compute Books have TruDDR4 memory that operate up to 1866 MHz. Up to 45 MB L3 cache. Intel C602J chipset. Models with E7 v3 processors are now withdrawn from marketing, however compute books are still available as field upgrades.</li> <li>▶ x3950 X6: Four, six, or eight Intel E7-8800 v4 processors, each in a Compute Book. Processor options from 4 cores (3.2 GHz) to 24 cores (up to 2.7 GHz). Three QPI links up to 9.6 GTps each. Compute Books feature TruDDR4 memory that operates at up to 1866 MHz. Up to 45 MB L3 cache. Intel C602J chipset. Models with E7 v3 processors are now withdrawn from marketing, however compute books are still available as field upgrades.</li> </ul>
Chipset	Intel C602J chipset.
Memory	<ul style="list-style-type: none"> <li>▶ x3850 X6: Up to 96 DIMM sockets</li> <li>▶ x3950 X6: Up to 192 DIMM sockets</li> <li>▶ v4 processors: 24 DIMMs per processor (installed in the Compute Book). RDIMMs and LRDIMMs (Load Reduced DIMMs) are supported, but memory types cannot be intermixed. Memory is TruDDR4 memory (up to 1866 MHz). DDR3 memory is not supported with v4 processors.</li> <li>▶ v3 processors: 24 DIMMs per processor (installed in the Compute Book). RDIMMs and LRDIMMs (Load Reduced DIMMs) are supported, but memory types cannot be intermixed. Memory is TruDDR4 memory (up to 1866 MHz) or DDR3 memory (up to 1600 MHz), depending on the Compute Book that is selected. DDR3 memory and TruDDR4 memory cannot be mixed in the same Compute Book. Compute Books with DDR3 memory and Compute Books with TruDDR4 memory cannot be mixed in the same system.</li> </ul>
Memory maximums	<ul style="list-style-type: none"> <li>▶ x3850 X6 with TruDDR4 Compute Books: <ul style="list-style-type: none"> <li>With RDIMMs: Up to 3 TB with 96x 32 GB RDIMMs and four processors</li> <li>With LRDIMMs: Up to 6 TB with 96x 64 GB LRDIMMs and four processors</li> </ul> </li> <li>▶ x3850 X6 with DDR3 Compute Books (now withdrawn): <ul style="list-style-type: none"> <li>With RDIMMs: Up to 1.5 TB with 96x 16 GB RDIMMs and four processors (v3 only)</li> <li>With LRDIMMs: Up to 6 TB with 96x 64 GB LRDIMMs and four processors (v3 only)</li> </ul> </li> <li>▶ x3950 X6 with TruDDR4 Compute Books: <ul style="list-style-type: none"> <li>With RDIMMs: Up to 6 TB with 192x 32 GB RDIMMs and eight processors</li> <li>With LRDIMMs: Up to 12 TB with 192x 64 GB LRDIMMs and eight processors</li> </ul> </li> <li>▶ x3950 X6 with DDR3 Compute Books (now withdrawn): <ul style="list-style-type: none"> <li>With RDIMMs: Up to 3 TB with 192x 16 GB RDIMMs and eight processors (v3 only)</li> <li>With LRDIMMs: Up to 12 TB with 192x 64 GB LRDIMMs and eight processors (v3 only)</li> </ul> </li> </ul>

Components	Specification
Memory protection	ECC, Chipkill, RBS, memory mirroring, and memory rank sparing.
Disk drive bays	<ul style="list-style-type: none"> <li>▶ x3850 X6: Up to eight 2.5-inch hot-swap SAS/SATA bays, or up to 16 1.8-inch SSD bays.</li> <li>▶ x3950 X6: Up to 16x eight 2.5-inch hot-swap SAS/SATA bays, or up to 32 1.8-inch SSD bays.</li> </ul>
Maximum internal storage	<ul style="list-style-type: none"> <li>▶ x3850 X6: Up to 30.7 TB with 3.84 TB 2.5-inch SAS SSDs Up to 14.4 TB with 1.8 TB 2.5-inch SAS HDDs</li> <li>▶ x3950 X6: Up to 61.4 TB with 3.84 TB 2.5-inch SAS SSDs Up to 28.8 TB with 1.8 TB 2.5-inch SAS HDDs</li> </ul>
RAID support	12 Gb SAS/SATA RAID 0, 1, or 10 with ServeRAID™ M5210; optional upgrades to RAID 5 or 50 are available (zero-cache; 1 GB non-backed cache; 1 GB or 2 GB flash-backed cache). Upgrades to RAID 6 or 60 available for M5210 with 1 GB or 2 GB upgrades.
Optical and tape bays	No internal bays; use an external USB drive. For more information, see this website: <a href="http://support.lenovo.com/en/documents/pd011281">http://support.lenovo.com/en/documents/pd011281</a>
Network interfaces	Mezzanine LOM (ML2) slot for dual-port 10 GbE cards with SFP+ or RJ-45 connectors or quad-port GbE cards with RJ-45 connectors. x3950 X6 has two ML2 slots. Dedicated 1 GbE port for systems management.
PCI Expansion slots	<p>x3850 X6: Up to 11 PCIe slots plus one dedicated Mezzanine LOM slot (12 total). The following slots are available:</p> <ul style="list-style-type: none"> <li>▶ Two PCIe 3.0 x8 slots for internal RAID controllers (Storage Book)</li> <li>▶ Two PCIe 3.0 x16 slots (x16-wired), half-length, full-height (Primary I/O Book)</li> <li>▶ One PCIe 3.0 x16 (x8-wired), half-length, full-height (Primary I/O Book)</li> <li>▶ One ML2 slot for network adapter (PCIe 3.0 x8) (Primary I/O Book)</li> <li>▶ Two optional I/O Books, each with three slots, all full height (use of these I/O Books requires four processors). Optional books are hot-swap capable.</li> </ul> <p>x3950 X6: Up to 22 PCIe slots plus two dedicated Mezzanine LOM slots (24 total). The following slots are available:</p> <ul style="list-style-type: none"> <li>▶ Four PCIe 3.0 x8 slots for internal RAID controllers (Storage Book)</li> <li>▶ Four PCIe 3.0 x16 slots (x16-wired), half-length, full height (Primary I/O Book)</li> <li>▶ Two PCIe 3.0 x16 (x8-wired), half-length, full-height (Primary I/O Book)</li> <li>▶ Two ML2 slots for network adapter (PCIe 3.0 x8) (Primary I/O Book)</li> <li>▶ Four optional I/O Books, each with three slots, all full height (use of these I/O Books requires four processors). Optional books are hot-swap capable.</li> </ul> <p>The following optional I/O Books are available:</p> <ul style="list-style-type: none"> <li>▶ Half-length I/O Book: Two PCIe 3.0 x8 slots, one PCIe 3.0 x16 slot</li> <li>▶ Full-length I/O Book: Two PCIe 3.0 x16, one PCIe 2.0 x4 slot; two auxiliary power connectors: 150 W and 75 W; supports one double-wide GPU up to 300 W</li> </ul>

Components	Specification
Ports	<ul style="list-style-type: none"> <li>▶ x3850 X6: Front: Two USB 3.0, one USB 2.0, and one DB-15 video ports. Rear: Four USB 2.0, one DB-15 video, one DB-9 serial, and one 1 GbE RJ-45 systems management. Internal: USB 2.0 port for embedded hypervisor.</li> <li>▶ x3950 X6: Front: Four USB 3.0, two USB 2.0, and two DB-15 video ports. Rear: Eight USB 2.0, two DB-15 video, two DB-9 serial, and two 1 GbE RJ-45 IMM2 systems management. Internal: two USB 2.0 port for embedded hypervisor. Note: The second video, IMM2, and internal USB hypervisor ports are used only when the server is partitioned into two four-socket servers.</li> </ul>
Cooling	<p>Each fan pack includes two counter-rotated dual-motor fans.</p> <ul style="list-style-type: none"> <li>▶ x3850 X6: Up to 10 redundant hot-swap fan packs; five fan zones with N+1 fan redundancy.</li> <li>▶ x3950 X6: Up to 20 redundant hot-swap fan packs; 10 fan zones with N+1 fan redundancy.</li> </ul>
Power supply	<ul style="list-style-type: none"> <li>▶ x3850 X6: Up to four redundant hot-swap power supplies.</li> <li>▶ x3950 X6: Up to eight redundant hot-swap power supplies.</li> </ul> <p>Power supplies can be 900 W AC or 1400 W AC power supplies (all 80 PLUS Platinum certified); -48V 750 W DC power supplies available via configure-to-order (CTO).</p> <p>Note: In a four or eight power supply configuration, mixing 900 W and 1400 W power supplies is allowed; however, they must be balanced and form pairs or a set of four of the same type.</p>
Hot-swap parts	Drives, power supplies, fans, optional I/O Books.
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.
Security features	<ul style="list-style-type: none"> <li>▶ x3850 X6: Power-on password, admin password, two TPMs</li> <li>▶ x3950 X6: Power-on password, admin password, four TPMs</li> </ul>
Systems management	UEFI, IMM2 (Version 2.1) with remote presence feature, Predictive Failure Analysis, Light Path Diagnostics, Automatic Server Restart, Lenovo XClarity Administrator, Lenovo XClarity Energy Manager, and ServerGuide™.
Supported operating systems	Microsoft Windows Server, Red Hat Enterprise Linux, SUSE Linux Enterprise Server, and VMware vSphere ESXi. See Chapter 6, "Operating system installation" on page 225 for specific versions.
Limited warranty	Three-year customer-replaceable unit (CRU) and onsite limited warranty with 9x5 next business day (NBD).
Service and support	Optional service upgrades are available through Lenovo Services offerings: 4-hour or 2-hour response time, 8-hour fix time, 1-year or 2-year warranty extension, and remote technical support for Lenovo hardware and some Lenovo/OEM applications.
Dimensions	<ul style="list-style-type: none"> <li>▶ x3850 X6: Height: 173 mm (6.8 in)</li> <li>▶ x3950 X6: Height: 351 mm (13.7 in)</li> <li>▶ Both: Width: 482 mm (19.0 in), depth: 804 mm (31.6 in), depth with cable management brackets installed: 836 mm (32.9 in), depth with Full-length I/O Book installed: 921 mm (36.2 in).</li> </ul>

Components	Specification
Weight	<ul style="list-style-type: none"> <li>▶ x3850 X6: Minimum configuration: 35.9 kg (79.2 lb), typical: 46.4 kg (102.3 lb), maximum: 54.7 kg (120 lb).</li> <li>▶ x3950 X6: Minimum configuration: 84.5 kg (186.3 lb), typical: 88.2 kg (194.5 lb), maximum: 110.0 kg (242.6 lb).</li> </ul>

The x3850 X6 and x3950 X6 servers are shipped with the following items:

- ▶ Rail kit
- ▶ Cable management bracket kit (2 brackets for x3850 X6, 4 brackets for the x3950 X6)
- ▶ 2.8 m (9.18 in) C13-C14 power cord (one for each power supply)
- ▶ Statement of Limited Warranty
- ▶ Important Notices
- ▶ Rack Installation Instructions
- ▶ Documentation CD that contains the Installation and User's Guide

Figure 3-1 shows the components of the cable management bracket kit that ship with the x3850 X6. A second set (that is, two left and two right) is also shipped with the x3950 X6.




Figure 3-1 Cable management bracket kit

### 3.3 Standard models of X6 servers

This section lists the standard models of the x3850 X6 and x3950 X6, machine type 6241.

**Machine type 3837:** Machine type 3837 (withdrawn from marketing) is not covered in this book. For information about these systems, see the following Lenovo Press product guides:

- ▶ System x3850 X6 (3837):  
<http://lenovopress.com/tips1084>
- ▶ System x3950 X6 (3837):  
<http://lenovopress.com/tips1132>

#### 3.3.1 x3850 X6 models

Table 3-2 lists the standard models of x3850 X6, machine type 6241.

Table 3-2 x3850 X6 models

Model <sup>a</sup>	Intel Xeon Processor <sup>b</sup>	Memory	eXFlash DIMMs	RAID	Drive bays Drives	Ethernet <sup>c</sup>	I/O slots std/max <sup>d</sup>	Power supplies
x3850 X6 models with E7 v4 processors (TruDDR4 memory)								

Model <sup>a</sup>	Intel Xeon Processor <sup>b</sup>	Memory	eXFlash DIMMs	RAID	Drive bays Drives	Ethernet <sup>c</sup>	I/O slots std/max <sup>d</sup>	Power supplies
6241-14x	1x E7-4809 v4 8C 2.1GHz 20M 115W	2x 8GB DDR4 (1866MHz) <sup>e</sup>	No support	Optional	Optional	4x 1 GbE	6 / 12	1x 900 W 4 maximum
6241-24x <sup>f</sup>	2x E7-4820 v4 10C 2.0GHz 25M 115W	4x 16GB DDR4 (1866MHz) <sup>e</sup>	No support	1x M5210	4x 2.5" / 8 Open bay	4x 1 GbE	6 / 12	2x 900 W 4 maximum
6241-44x <sup>f</sup>	2x E7-4850 v4 16C 2.1GHz 40M 115W	4x 16GB DDR4 (1866MHz) <sup>e</sup>	No support	1x M5210	4x 2.5" / 8 Open bay	4x 1 GbE	6 / 12	2x 900 W 4 maximum
6241-54x <sup>f</sup>	2x E7-8860 v4 18C 2.2GHz 45M 140W	4x 16GB DDR4 (1866MHz) <sup>e</sup>	No support	1x M5210	4x 2.5" / 8 Open bay	4x 1 GbE	6 / 12	2x 900 W 4 maximum
6241-94x <sup>f</sup>	2x E7-8890 v4 24C 2.2GHz 60M 165W	4x 16GB DDR4 (1866MHz) <sup>e</sup>	No support	1x M5210	4x 2.5" / 8 Open bay	4x 1 GbE	6 / 12	2x 900 W 4 maximum
<b>x3850 X6 TopSeller™ models with E7 v4 processors (TruDDR4 memory) - North America only</b>								
6241-E1U	2x E7-4820 v4 10C 2.0GHz 25M 115W	4x 16GB DDR4 (1866MHz) <sup>e</sup>	No support	1x M5210	4x 2.5" / 8 Open bay	4x 1 GbE	6 / 12	2x 900 W 4 maximum
6241-E2U	2x E7-4850 v4 16C 2.1GHz 40M 115W	4x 16GB DDR4 (1866MHz) <sup>e</sup>	No support	1x M5210	4x 2.5" / 8 Open bay	4x 1 GbE	6 / 12	2x 900 W 4 maximum
6241-E3U	2x E7-8890 v4 24C 2.2GHz 60M 165W	4x 16GB DDR4 (1866MHz) <sup>e</sup>	No support	1x M5210	4x 2.5" / 8 Open bay	4x 1 GbE	6 / 12	2x 900 W 4 maximum
<b>x3850 X6 models with E7 v3 processors and TruDDR4 memory (all withdrawn from marketing)</b>								
6241-D5x	1x E7-4809 v3 8C 2.0GHz 20MB 115W	2x 8 GB DDR4 (1866MHz) <sup>e</sup>	No support	Optional	Optional	4x 1 GbE	6 / 12	1x 900 W 4 maximum
6241-F2x	2x E7-4820 v3 10C 1.9GHz 25MB 115W	4x 16 GB DDR4 (1866 MHz) <sup>e</sup>	No support	1x M5210	4x 2.5-inch HS 8 maximum Open bay	4x 1 GbE	6 / 12	2x 900 W 4 maximum
6241-F4x	2x E7-4850 v3 14C 2.2GHz 35MB 115W	4x 16 GB DDR4 (1866 MHz) <sup>e</sup>	No support	1x M5210	4x 2.5-inch HS 8 maximum Open bay	2x 10 GbE	6 / 12	2x 900 W 4 maximum
6241-G2x	2x E7-8860 v3 16C 2.2GHz 40MB 140W	4x 16 GB DDR4 (1866 MHz) <sup>e</sup>	No support	1x M5210	4x 2.5-inch HS 8 maximum Open bay	4x 1 GbE	6 / 12	2x 900 W 4 maximum
6241-G4x	2x E7-8890 v3 18C 2.5GHz 45MB 165W	4x 16 GB DDR4 (1866 MHz) <sup>e</sup>	No support	1x M5210	4x 2.5-inch HS 8 maximum Open bay	4x 1 GbE	6 / 12	2x 900 W 4 maximum
<b>x3850 X6 models with E7 v3 processors and DDR3 memory (all withdrawn from marketing)</b>								
6241-D4x	1x E7-4809 v3 8C 2.0GHz 20MB 115W	2x 8 GB DDR3 (1600 MHz)	Optional	Optional	Optional	4x 1 GbE	6 / 12	1x 900 W 4 maximum
6241-F1x	2x E7-4820 v3 10C 1.9GHz 25MB 115W	4x 16 GB DDR3 (1600 MHz)	Optional	1x M5210	4x 2.5-inch HS 8 maximum Open bay	4x 1 GbE	6 / 12	2x 900 W 4 maximum

Model <sup>a</sup>	Intel Xeon Processor <sup>b</sup>	Memory	eXFlash DIMMs	RAID	Drive bays Drives	Ethernet <sup>c</sup>	I/O slots std/max <sup>d</sup>	Power supplies
6241-F3x	2x E7-4850 v3 14C 2.2GHz 35MB 115W	4x 16 GB DDR3 (1600 MHz)	Optional	1x M5210	4x 2.5-inch HS 8 maximum Open bay	2x 10 GbE <sup>c</sup>	6 / 12	2x 900 W 4 maximum
6241-G1x	2x E7-8860 v3 16C 2.2GHz 40MB 140W	4x 16 GB DDR3 (1600 MHz)	Optional	1x M5210	4x 2.5-inch HS 8 maximum Open bay	4x 1 GbE	6 / 12	2x 900 W 4 maximum
6241-G3x	2x E7-8890 v3 18C 2.5GHz 45MB 165W	4x 16 GB DDR3 (1600 MHz)	Optional	1x M5210	4x 2.5-inch HS 8 maximum Open bay	4x 1 GbE	6 / 12	2x 900 W 4 maximum
<b>x3850 X6 models with E7 v2 processors and DDR3 memory (all withdrawn from marketing)</b>								
6241-A4x <sup>g</sup>	1x E7-4809 v2 6C 1.9GHz 12MB 105W	2x 8 GB DDR3 (1333 MHz) <sup>e</sup>	Optional	Optional	Optional	4x 1 GbE	6 / 12	1x 900 W 4 maximum
6241-B1x <sup>g</sup>	2x E7-4820 v2 8C 2.0GHz 16MB 105W	4x 8 GB DDR3 (1600 MHz)	Optional	1x M5210	4x 2.5-inch HS 8 maximum Open bay	4x 1 GbE	6 / 12	2x 900 W 4 maximum
6241-B3x <sup>g</sup>	2x E7-4850 v2 12C 2.3GHz 24MB 105W	4x 8 GB DDR3 (1600 MHz)	Optional	1x M5210	4x 2.5-inch HS 8 maximum Open bay	2x 10 GbE	6 / 12	2x 900 W 4 maximum
6241-C1x <sup>g</sup>	2x E7-4860 v2 12C 2.6GHz 30MB 130W	4x 8 GB DDR3 (1600 MHz)	Optional	1x M5210	4x 2.5-inch HS 8 maximum Open bay	4x 1 GbE	6 / 12	2x 900 W 4 maximum
6241-C4x <sup>g</sup>	2x E7-4890 v2 15C 2.8GHz 37.5MB 155W	4x 8 GB DDR3 (1600 MHz)	Optional	1x M5210	4x 2.5-inch HS 8 maximum Open bay	4x 1 GbE	6 / 12	2x 900 W 4 maximum

- The letter “x” in the Machine Type Model (MTM) represents a country-specific letter (for example, the EMEA MTM is 6241-A4G, and the US MTM is 6241-A4U). For more information, contact a Lenovo representative.
- Processor detail: Processor quantity and model, cores, core speed, L3 cache, and TDP.
- Models F3x, F4x, and B3x include the Broadcom NetXtreme II ML2 Dual Port 10GbE SFP+ adapter (BCM57810S based). All other models include Intel I350-T4 ML2 Quad Port GbE Adapter adapter (I350-AM4 based).
- Models with six slots have the Primary I/O Book (four slots) and Storage Book (two slots) standard.
- The DIMMs in these models operate at the speed indicated to match the memory bus speed of the processor.
- Models 24x, 44x, 54x and 94x not available in North America.
- Withdrawn from marketing

### 3.3.2 x3950 X6 modes

Table 3-3 lists the standard models of x3950 X6, machine type 6241.

Table 3-3 x3950 X6 models

Model <sup>a</sup>	Intel Xeon Processor <sup>b</sup>	Memory	eXFlash DIMMs	RAID	Drive bays (std/max)	Ethernet (ML2 slot) <sup>c</sup>	I/O slots std/max <sup>d</sup>	Power supplies
<b>x3950 X6 models with E7 v4 processors (TruDDR4 memory)</b>								
6241-58x	4x E7-8860 v4 18C 2.2GHz 45M 140W	8x 16 GB DDR4 1866 MHz	No support	1x M5210	4x 2.5-inch HS 16 maximum Open bay	One 4x1GbE	12 / 24	4x 900 W 8 maximum
6241-78x	4x E7-8870 v4 20C 2.1GHz 50M 140W	8x 16 GB DDR4 1866 MHz	No support	1x M5210	4x 2.5-inch HS 16 maximum Open bay	One 4x1GbE	12 / 24	4x 900 W 8 maximum



Model <sup>a</sup>	Intel Xeon Processor <sup>b</sup>	Memory	eXFlash DIMMs	RAID	Drive bays (std/max)	Ethernet (ML2 slot) <sup>c</sup>	I/O slots std/max <sup>d</sup>	Power supplies
6241-98x	4x E7-8890 v4 24C 2.2GHz 60M 165W	8x 16 GB DDR4 1866 MHz	No support	1x M5210	4x 2.5-inch HS 16 maximum Open bay	One 4x1GbE	12 / 24	4x 900 W 8 maximum
<b>x3950 X6 models with E7 v3 processors and TruDDR4 memory (all withdrawn from marketing)</b>								
6241-GAx	4x E7-8860 v3 16C 2.2GHz 140W	8x 16 GB DDR4 1866 MHz	No support	1x M5210	4x 2.5-inch HS 16 maximum Open bay	One 4x1GbE	12 / 24	4x 900 W 8 maximum
6241-JAx	4x E7-8870 v3 18C 2.1GHz 140W	8x 16 GB DDR4 1866 MHz	No support	1x M5210	4x 2.5-inch HS 16 maximum Open bay	One 4x1GbE	12 / 24	4x 900 W 8 maximum
6241-JCx	4x E7-8890 v3 18C 2.5GHz 165W	8x 16 GB DDR4 1866 MHz	No support	1x M5210	4x 2.5-inch HS 16 maximum Open bay	One 4x1GbE	12 / 24	4x 900 W 8 maximum
<b>x3950 X6 models with E7 v3 processors and DDR3 memory (withdrawn from marketing)</b>								
6241-DAx	4x E7-8860 v3 16C 2.2GHz 140W	8x 16 GB DDR3 1600 MHz	Optional	1x M5210	4x 2.5-inch HS 16 maximum Open bay	One 4x1GbE	12 / 24	4x 900 W 8 maximum
6241-FAx	4x E7-8870 v3 18C 2.1GHz 140W	8x 16 GB DDR3 1600 MHz	Optional	1x M5210	4x 2.5-inch HS 16 maximum Open bay	One 4x1GbE	12 / 24	4x 900 W 8 maximum
6241-FCx	4x E7-8890 v3 18C 2.5GHz 165W	8x 16 GB DDR3 1600 MHz	Optional	1x M5210	4x 2.5-inch HS 16 maximum Open bay	One 4x1GbE	12 / 24	4x 900 W 8 maximum
<b>x3950 X6 models with E7 v2 processors and DDR3 memory (withdrawn from marketing)</b>								
6241-BAx <sup>e</sup>	4x Xeon E7-8850 v2 12C 2.3GHz 24MB 105W	8x 8 GB 1600MHz	Optional	1x M5210	4x 2.5-inch HS 16 maximum Open bay	One 4x1GbE	12 / 24	4x 900 W 8 maximum
6241-CAx <sup>e</sup>	4x Xeon E7-8870 v2 15C 2.3GHz 30MB 130W	8x 8 GB 1600MHz	Optional	1x M5210	4x 2.5-inch HS 16 maximum Open bay	One 4x1GbE	12 / 24	4x 900 W 8 maximum
6241-CCx <sup>e</sup>	4x Xeon E7-8890 v2 15C 2.8GHz 37.5MB 155W	8x 8 GB 1600MHz	Optional	1x M5210	4x 2.5-inch HS 16 maximum Open bay	One 4x1GbE	12 / 24	4x 900 W 8 maximum

a. The letter “x” in the Machine Type Model (MTM) represents a country-specific letter (for example, the EMEA MTM is 6241-BAG, and the US MTM is 6241-BAU). For more information, contact a Lenovo representative.

b. Processor detail: Processor quantity and model, cores, core speed, L3 cache, and TDP.

c. All models include one Intel I350-T4 ML2 Quad Port GbE Adapter adapter (I350-AM4 based).

d. Models with 12 slots have two Primary I/O Books (each with four slots) and two Storage Books (each with two slots) standard.

e. Withdrawn from marketing

For more information about standard features of the server, see Table 3-1 on page 57.

## 3.4 SAP HANA Workload Optimized Solution models

This section lists the Workload Optimized Solution models for SAP HANA for x3850 X6.

### 3.4.1 x3850 X6 HANA models

Table 3-4 lists the SAP HANA models of x3850 X6, machine type 6241.

Table 3-4 SAP HANA Workload Optimized Solution models for x3850 X6

Model	Intel Xeon Processor	Memory	RAID	Drive bays Drives	Ethernet	I/O slots	Power supplies
<b>SAP HANA models with with E7 v4 processors (TruDDR4 memory)</b>							
6241-8Ax	2x E7-8880 v4 22C 2.2GHz 55M 150W	128 GB (16x 8GB) DDR4 RDIMM	1x M5210 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3710 SSD	1x 4x1GbE ML2; 2x Mellanox 10Gb <sup>a</sup>	6 std 12 max	4x 1400W HS / 4
6241-8Bx	2x E7-8880 v4 22C 2.2GHz 55M 150W	256 GB (16x 16GB) DDR4 RDIMM	1x M5210 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3710 SSD	1x 4x1GbE ML2; 2x Mellanox 10Gb <sup>a</sup>	6 std 12 max	4x 1400W HS / 4
6241-8Cx	2x E7-8880 v4 22C 2.2GHz 55M 150W	512 GB (16x 32GB) DDR4 RDIMM	1x M5210 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3710 SSD	1x 4x1GbE ML2; 2x Mellanox 10Gb <sup>a</sup>	6 std 12 max	4x 1400W HS / 4
6241-8Dx	4x E7-8880 v4 22C 2.2GHz 55M 150W	512 GB (32x 16GB) DDR4 RDIMM	1x M5210 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3710 SSD	1x 4x1GbE ML2; 2x Mellanox 10Gb <sup>a</sup>	12 std 12 max	4x 1400W HS / 4
6241-8Fx	4x E7-8880 v4 22C 2.2GHz 55M 150W	1024 GB (32x 32GB) DDR4 RDIMM	1x M5210 + upgrades <sup>a</sup> + 1x M5225 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3710 SSD	1x 4x1GbE ML2; 2x Mellanox 10Gb <sup>a</sup>	12 std 12 max	4x 1400W HS / 4
<b>SAP HANA models E7 v4 processors (TruDDR4 memory) - North America only</b>							
6241-EKU	2x E7-8880 v4 22C 2.2GHz 55M 150W	16x 16GB (1866MHz) <sup>□</sup>	1x M5210 + upgrades <sup>a</sup>	8x 2.5"; 6x 1.2 TB SAS HDD; 2x 400 GB S3710 SSD	1x 4x1GbE ML2; 2x Mellanox 10Gb <sup>a</sup>	6 std; 12 max	4x 1400W HS / 4
<b>SAP HANA models with E7 v3 processors and TruDDR4 memory (withdrawn from marketing)</b>							
6241-HQx	4x E7-8880 v3 18C 2.3GHz 45MB 150W	1024 GB 32x 32GB DDR4 RDIMM	1x M5210 + upgrades 1x M5120 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3700 SSD	1x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	12 std 12 max	4x 1400 W HS / 4
6241-HRx	4x E7-8880 v3 18C 2.3GHz 45MB 150W	512 GB 32x 16GB DDR4 RDIMM	1x M5210 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3700 SSD	1x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	12 std 12 max	4x 1400 W HS / 4
6241-HSx	2x E7-8880 v3 18C 2.3GHz 45MB 150W	512 GB 16x 32GB DDR4 RDIMM	1x M5210 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3700 SSD	1x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	6 std 12 max	4x 1400 W HS / 4
6241-HTx	2x E7-8880 v3 18C 2.3GHz 45MB 150W	256 GB 16x 16GB DDR4 RDIMM	1x M5210 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3700 SSD	1x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	6 std 12 max	4x 1400 W HS / 4
6241-HUx	2x E7-8880 v3 18C 2.3GHz 45MB 150W	128 GB 16x 8GB DDR4 RDIMM	1x M5210 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3700 SSD	1x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	6 std 12 max	4x 1400 W HS / 4

Model	Intel Xeon Processor	Memory	RAID	Drive bays Drives	Ethernet	I/O slots	Power supplies
<b>SAP HANA models with E7 v3 processors and DDR3 memory (withdrawn from marketing)</b>							
6241-HVx	4x E7-8880 v3 18C 2.3GHz 45MB 150W	1024 GB 32x 32 GB DDR3 LRDIMM	1x M5210 + upgrades <sup>a</sup> 1x M5120 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3700 SSD	1x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	12 std 12 max	4x 1400 W HS / 4
6241-HWx	4x E7-8880 v3 18C 2.3GHz 45MB 150W	512 GB 32x 16 GB DDR3 RDIMM	1x M5210 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3700 SSD	1x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	12 std 12 max	4x 1400 W HS / 4
6241-HXx	2x E7-8880 v3 18C 2.3GHz 45MB 150W	512 GB 16x 32 GB DDR3 RDIMM	1x M5210 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3700 SSD	1x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	6 std 12 max	4x 1400 W HS / 4
6241-HYx	2x E7-8880 v3 18C 2.3GHz 45MB 150W	256 GB 16x 16 GB DDR3 RDIMM	1x M5210 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3700 SSD	1x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	6 std 12 max	4x 1400 W HS / 4
6241-HZx	2x E7-8880 v3 18C 2.3GHz 45MB 150W	128 GB 16x 8 GB DDR3 RDIMM	1x M5210 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3700 SSD	1x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	6 std 12 max	4x 1400 W HS / 4
<b>SAP HANA models with E7 v2 processors and DDR3 memory (withdrawn from marketing)</b>							
6241-H2x <sup>b</sup>	2x E7-8880 v2 15C 2.5GHz 37.5MB 130W	128 GB 16x 8 GB DDR3 RDIMM	1x M5210 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3700 SSD	1x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	6 std 12 max	4x 1400 W HS / 4
6241-H3x <sup>b</sup>	2x E7-8880 v2 15C 2.5GHz 37.5MB 130W	256 GB 16x 16 GB DDR3 RDIMM	1x M5210 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3700 SSD	1x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	6 std 12 max	4x 1400 W HS / 4
6241-H4x <sup>b</sup>	2x E7-8880 v2 15C 2.5GHz 37.5MB 130W	512 GB 16x 32 GB DDR3 LRDIMM	1x M5210 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3700 SSD	1x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	6 std 12 max	4x 1400 W HS / 4
6241-H5x <sup>b</sup>	2x E7-8880 v2 15C 2.5GHz 37.5MB 130W	512 GB 32x 16 GB DDR3 RDIMM	1x M5210 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3700 SSD	1x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	12 std 12 max	4x 1400 W HS / 4
6241-H6x <sup>b</sup>	2x E7-8880 v2 15C 2.5GHz 37.5MB 130W	1024 GB 32x 32 GB DDR3 LRDIMM	1x M5210 + upgrades <sup>a</sup> 1x M5120 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3700 SSD	1x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	12 std 12 max	4x 1400 W HS / 4

a. See the list of specific components in the following section.

b. Withdrawn from marketing

Each of these HANA models include the following components

- ▶ ServeRAID M5210 controller (46C9110) with 2GB cache upgrade with flash backup (47C8664), plus Performance Accelerator (47C8710) and SSD Caching Enabler (47C8712)
- ▶ Models HQx, HVx, H6x: ServeRAID M5120 RAID controller for external connectivity (81Y4478), with 1GB cache upgrade with flash backup (47C8660) plus Performance Accelerator (47C8710) and SSD Caching Enabler (47C8712).
- ▶ Model 8Fx: ServeRAID M5225-2GB RAID controller for external connectivity (00AE938) plus Performance Accelerator (47C8710) and SSD Caching Enabler (47C8712).
- ▶ Six 1.2TB 10K 6Gbps SAS 2.5" G3HS HDD (00AJ146)

- ▶ Models with E7 v4 processors: Two Intel S3710 400GB Enterprise Performance SATA G3HS 2.5" SSDs (00YC325)
- ▶ Models with E7 v3 or v2 processors: Two S3700 400GB SATA 2.5" MLC G3HS Enterprise SSDs (00AJ161)
- ▶ One Intel I350-T4 ML2 Quad Port GbE Adapter (00D1998)
- ▶ Models with E7 v4 processors: Two Mellanox ConnectX-3 10 GbE Adapters (00D9690) plus four SFP+ SR Transceivers (46C3447)
- ▶ Models with E7 v3 or v2 processors: Two Mellanox ConnectX-3 40GbE / FDR IB VPI Adapters (00D9550) plus four Mellanox QSA Adapters (QSFP to SFP+) (90Y3842) and four SFP+ SR Transceivers (46C3447)
- ▶ Models with 12 standard slots: Two X6 Half-length I/O Books (44X4049)
- ▶ Lenovo solution for SAP HANA media

**Note:** The operating system software is not included with the SAP HANA models. Operating system selection must be a separate line item included in order: SLES for SAP with standard or priority support. The SAP HANA Software is included, but the license is sold separately by SAP or an SAP business partner. VMware Enterprise Plus license sold separately. IBM Spectrum Scale (GPFS) is available from Lenovo separately.

### 3.4.2 x3950 X6 HANA models

Table 3-5 lists the SAP HANA models of x3950 X6, machine type 6241.

Table 3-5 SAP HANA Workload Optimized Solution models for x3950 X6

Model	Intel Xeon Processor	Memory	RAID	Drive bays Drives	Ethernet	I/O slots	Power supplies
<b>x3950 X6 models with E7 v4 processors (TruDDR4 memory) - HANA models</b>							
6241-8Gx	4x E7-8880 v4 22C 2.2GHz 55M 150W	512 GB 32x 16GB DDR4 RDIMM	1x M5210 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3710 SSD	2x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	12 std; 24 max	8x 1400W HS / 8
6241-8Hx	4x E7-8880 v4 22C 2.2GHz 55M 150W	1024 GB 32x 32GB DDR4 RDIMM	2x M5210 + upgrades <sup>a</sup>	16x 2.5-inch HS 12x 1.2 TB SAS HDD 4x 400 GB S3710 SSD	2x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	12 std; 24 max	8x 1400W HS / 8
6241-8Jx	8x E7-8880 v4 22C 2.2GHz 55M 150W	2048 GB 64x 32GB DDR4 RDIMM	2x M5210 + upgrades <sup>a</sup>	16x 2.5-inch HS 12x 1.2 TB SAS HDD 4x 400 GB S3710 SSD	2x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	18 std; 24 max	8x 1400W HS / 8
<b>x3950 X6 HANA models with E7 v4 processors (TruDDR4 memory) - TopSeller - North America only</b>							
6241-ELU	8x E7-8880 v4 22C 2.2GHz 55M 150W	512 GB 32x 16GB DDR4 RDIMM	1x M5210 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3710 SSD	2x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	12 std; 24 max	8x 1400W HS / 8
<b>x3950 X6 HANA models with E7 v3 processors and TruDDR4 memory (withdrawn from marketing)</b>							
6241-HHx	4x E7-8880 v3 18C 2.3GHz 45MB 150W	512 GB 32x 16 GB DDR4 RDIMM	1x M5210 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3700 SSD	2x 4x1GbE ML2 2x Mellanox 10G <sup>a</sup>	12 std 24 max	8x 1400 W HS / 8
6241-Hlx	4x E7-8880 v3 18C 2.3GHz 45MB 150W	1024 GB 32x 32 GB DDR4 RDIMM	2x M5210 + upgrades <sup>a</sup>	16x 2.5-inch HS 12x 1.2 TB SAS HDD 4x 400 GB S3700 SSD	2x 4x1GbE ML2 2x Mellanox 10G <sup>a</sup>	12 std 24 max	8x 1400 W HS / 8

Model	Intel Xeon Processor	Memory	RAID	Drive bays Drives	Ethernet	I/O slots	Power supplies
6241-HJx	8x E7-8880 v3 18C 2.3GHz 45MB 150W	2048 GB 64x 32 GB DDR4 RDIMM	2x M5210 + upgrades <sup>a</sup>	16x 2.5-inch HS 12x 1.2 TB SAS HDD 4x 400 GB S3700 SSD	2x 4x1GbE ML2 2x Mellanox 10G <sup>a</sup>	18 std 24 max	8x 1400 W HS / 8
<b>x3950 X6 HANA models with E7 v3 processors and DDR3 memory (withdrawn from marketing)</b>							
6241-HEx	4x E7-8880 v2 15C 2.5GHz 37.5MB 130W	512 GB 32x 16 GB DDR3 RDIMM	1x M5210 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3700 SSD	2x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	12 std 24 max	8x 1400 W HS / 8
6241-HFx	4x E7-8880 v2 15C 2.5GHz 37.5MB 130W	1024 GB 32x 32 GB DDR3 LRDIMM	2x M5210 + upgrades <sup>a</sup>	16x 2.5-inch HS 12x 1.2 TB SAS HDD 4x 400 GB S3700 SSD	2x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	12 std 24 max	8x 1400 W HS / 8
6241-HGx	8x E7-8880 v2 15C 2.5GHz 37.5MB 130W	2048 GB 64x 32 GB DDR3 LRDIMM	2x M5210 + upgrades <sup>a</sup>	16x 2.5-inch HS 12x 1.2 TB SAS HDD 4x 400 GB S3700 SSD	2x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	18 std 24 max	8x 1400 W HS / 8
<b>x3950 X6 HANA models with E7 v2 processors and DDR3 memory (withdrawn from marketing)</b>							
6241-HBx <sup>b</sup>	4x E7-8880 v2 15C 2.5GHz 37.5MB 130W	512 GB 32x 16 GB DDR3 RDIMM	1x M5210 + upgrades <sup>a</sup>	8x 2.5-inch HS 6x 1.2 TB SAS HDD 2x 400 GB S3700 SSD	2x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	12 std 24 max	8x 1400 W HS / 8
6241-HCx <sup>b</sup>	4x E7-8880 v2 15C 2.5GHz 37.5MB 130W	1024 GB 32x 32 GB DDR3 LRDIMM	2x M5210 + upgrades <sup>a</sup>	16x 2.5-inch HS 12x 1.2 TB SAS HDD 4x 400 GB S3700 SSD	2x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	12 std 24 max	8x 1400 W HS / 8
6241-HDx <sup>b</sup>	8x E7-8880 v2 15C 2.5GHz 37.5MB 130W	2048 GB 64x 32 GB DDR3 LRDIMM	2x M5210 + upgrades <sup>a</sup>	16x 2.5-inch HS 12x 1.2 TB SAS HDD 4x 400 GB S3700 SSD	2x 4x1GbE ML2 2x Mellanox 10Gb <sup>a</sup>	18 std 24 max	8x 1400 W HS / 8

a. See the list of specific components in the next section.

b. Withdrawn from marketing

Each of these HANA models include the following components

- ▶ One or two ServeRAID M5210 controller (46C9110) with 2GB cache upgrade with flash backup (47C8664), plus Performance Accelerator (47C8710) and SSD Caching Enabler (47C8712)
- ▶ Twelve 1.2TB 10K 6Gbps SAS 2.5" G3HS HDD (00AJ146)
- ▶ Models with E7 v4 processors: Four Intel S3710 400GB Enterprise Performance SATA G3HS 2.5" SSDs (00YC325)
- ▶ Models with E7 v3 or v2 processors: Four S3700 400GB SATA 2.5" MLC G3HS Enterprise SSDs (00AJ161)
- ▶ Two Intel I350-T4 ML2 Quad Port GbE Adapter (00D1998)
- ▶ Models with E7 v4 processors: Two Mellanox ConnectX-3 10 GbE Adapters (00D9690) plus four SFP+ SR Transceivers (46C3447)
- ▶ Models with E7 v3 or v2 processors: Two Mellanox ConnectX-3 40GbE / FDR IB VPI Adapters (00D9550) plus four Mellanox QSA Adapters (QSFP to SFP+) (90Y3842) and four SFP+ SR Transceivers (46C3447)
- ▶ Models with 18 standard slots: Two X6 Half-length I/O Books (44X4049)
- ▶ Lenovo solution for SAP HANA media

**Note:** The operating system software is not included with the SAP HANA models. Operating system selection must be a separate line item included in order: SLES for SAP with standard or priority support. The SAP HANA Software is included, but the license is sold separately by SAP or an SAP business partner. VMware Enterprise Plus license sold separately. IBM Spectrum Scale (GPFS) is available from Lenovo separately.

## 3.5 Physical design

The x3850 X6 and x3950 X6 servers are based on a new design in which all major server components are in easy-to-service modules called “books”. A book can be inserted into the server chassis from the front or rear of the system (depending on the type of book). In the middle of the server is a passive midplane, which connects all of the server’s components to each other.

The following types of books are available:

- ▶ **Compute Books:** This book includes a module that is accessible from the front of the server and has one processor and up to 24 DIMM modules.
- ▶ **Storage Books:** This book features the server’s modules, which contain the set of HDDs or SSDs, several PCIe RAID adapters, USB and video ports, and an operator panel. Storage Books are in front of the server.
- ▶ **I/O Books:** Accessible from the rear of the server, these books provide I/O capabilities. Each I/O module provides up to three PCIe 3.0 slots for half-length or full-length PCIe adapters.

Figure 3-2 shows the front view of the x3850 X6 server.




Figure 3-2 The x3850 X6 front view

As shown in Figure 3-2, each x3850 X6 server can have up to four Compute Books (two Compute Books minimum), and one Storage Book, which can have up to 8x 2.5-inch HDDs or SSDs or up to 16x 1.8-inch SSDs. The x3850 X6 server also has three USB 3.0 ports accessible from the front, one video port, and the control panel with LCD screen.

Figure 3-3 shows the front view of the x3950 X6 server.




Figure 3-3 Front view of the x3950 X6

As shown in Figure 3-3, the x3950 X6 is the equivalent of two x3850 X6 servers. It can have up to eight Compute Books (four Compute Books minimum), and two Storage Books, each of which can have up to 8x 2.5-inch drives or up to 16x 1.8-inch SSDs.

For more information about the Compute Book, see 3.7, “Compute Books” on page 75. For more information about the Storage Book, see 3.11, “Storage subsystem” on page 94.

Figure 3-4 shows the rear of the x3850 X6 server.




Figure 3-4 Rear view of the x3850 X6

The rear of the x3850 X6 has up to three I/O Books. The books feature the following configuration:

- ▶ One primary I/O Book with 3 PCIe 3.0 slots, one ML2-slot for Ethernet adapters, two hot-swappable fans, one IMM2 Ethernet port, 4 USB 2.0 ports, video port, and serial port
- ▶ Two other I/O Books, each with 3 PCIe 3.0 ports, which can be used for half-length or full-length PCIe cards

For more information about the I/O Books, see 3.12, “I/O subsystem” on page 100.

At the bottom of the server at the rear are bays for up to four power supply modules. For more information, see 3.24, “Power subsystem” on page 122.

Figure 3-5 shows the rear of the x3950 X6.



Figure 3-5 Rear view of the x3950 X6

## 3.6 Ports and controls

The x3850 X6 and x3950 X6 servers feature a broad set of I/O ports and controls, including an LCD system information display panel located on the front operator panel. In this section, the ports, controls, and LEDs on the X6 server are described.

### 3.6.1 Front operator panel

The front operator panel is a panel on the front side of the Storage Book, which contains controls and LEDs that provide information about the status of the server.



Figure 3-6 shows ports, controls, and LEDs on the front operator panel.




Figure 3-6 Ports, controls, and LEDs on the front operator panel

The front operator panel contains the following controls:

► **Power button/LED**

Press the power button to turn the server on and off manually or to wake the server from a reduced-power state. The LED features the following states:

- Off: Input power is not present, or the power supply or LED failed.
- Flashing rapidly (3 times per second): The server is turned off and is not ready to be turned on. The power-on button is disabled. This state lasts approximately 10 seconds after input power was applied or restored.
- Flashing slowly (once per second): The server is turned off and is ready to be turned on. You can press the power-on button to turn on the server.
- Lit: The server is turned on.

► **Locate button/LED**

Press the locate button to visually locate the server among other servers. When you press the locate button, the LED is lit and it continues to be lit until you press it again to turn it off.

► **Check log LED**

When this LED is lit (yellow), it indicates that a noncritical event occurred, such as the wrong I/O Book was inserted in a bay, or power requirements exceeded the capacity of the installed power supplies.

► **System-error LED**

When this yellow LED is lit, it indicates that a system error occurred. A system-error LED is also on the rear of the server. This LED is controlled by the IMM2.

► **Reset button:** Press this button to reset the server and run the power-on self-test (POST). You might have to use a pen point or the end of a straightened paper clip to press the button. The Reset button is in the upper-right corner of the front operator panel.

- ▶ PCIe slots 11 and 12 error LEDs  
When these LEDs are lit, they indicate that an error occurred in PCIe slots 11 and 12 on the storage book board.
- ▶ USB 2.0 port  
Connect a USB device, such as a USB mouse, keyboard, or other device, to this connector.
- ▶ USB 3.0 ports  
Connect a USB device, such as a USB mouse, keyboard, or other device, to any of these connectors.
- ▶ Video connector  
Connect a monitor to this connector. The video connectors on the front and rear of the server can be used simultaneously.

### 3.6.2 LCD system information panel

The LCD system information panel displays various types of information about the server. It is on the Storage Book on the front of the server (see Figure 3-6 on page 71). The information panel provides quick access to system status, firmware, network, and health information.

Figure 3-7 shows the LCD system information panel.



Figure 3-7 The LCD system information panel

The LCD system information display panel contains the following buttons:

- ▶ Scroll up button: Press this button to scroll up or scroll to the left in the main menu to locate and select the system information that you want displayed.
- ▶ Select button: Press this button to make your selection from the menu options.
- ▶ Scroll down button: Press this button to scroll down or scroll to the right in the main menu to location and select the system information that you want displayed.

Figure 3-8 shows the available options that can be shown on the LCD system information panel.

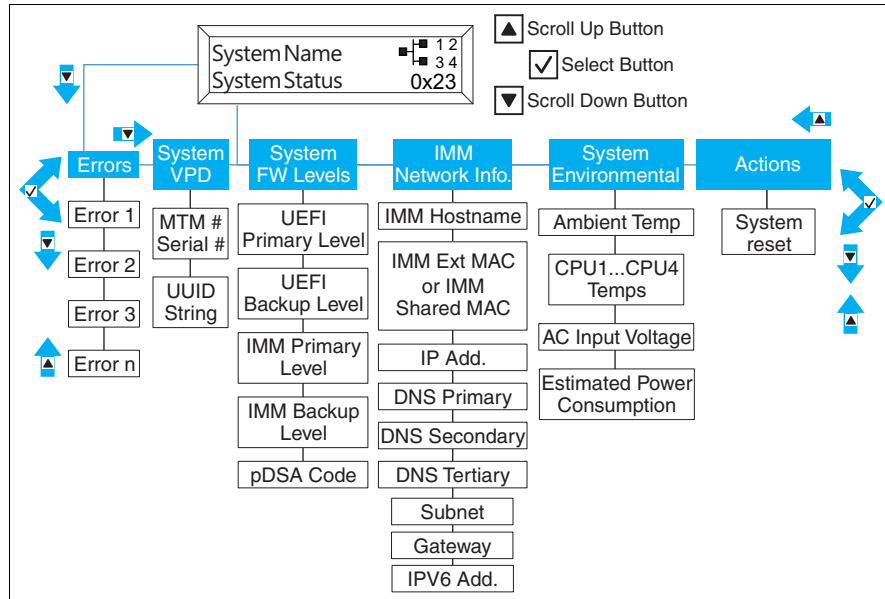


Figure 3-8 System properties that are shown on the LCD system information panel

Figure 3-9 shows an example of the information that you might see on the display panel.

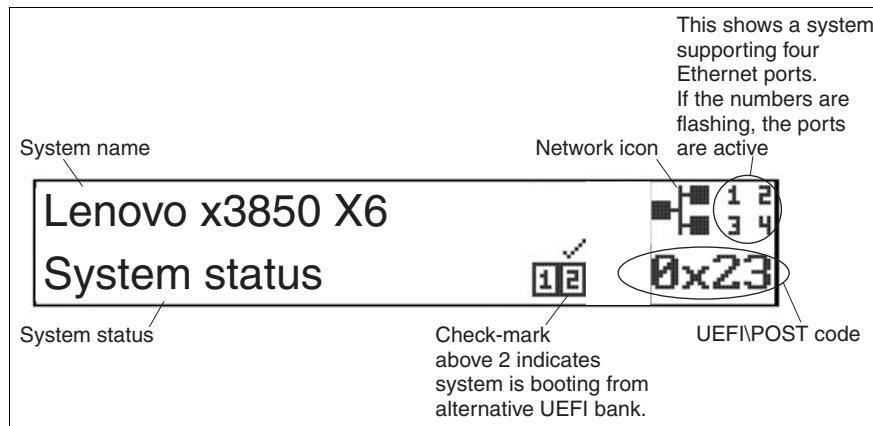


Figure 3-9 LCD system information panel example

You can use the Scroll up and Scroll down buttons to navigate inside the menu. You can use the Select button to choose an appropriate submenu.

For the Errors submenu set, if only one error occurs, the LCD display panel displays that error. If more than one error occurs, the LCD display panel displays the number of errors that occurred.

The LCD system information display panel displays the following types of information about the server:

- ▶ IMM system error log
- ▶ System VPD information:
  - Machine type and serial number

- Universal Unique Identifier (UUID) string
- ▶ System firmware levels:
  - UEFI code level
  - IMM code level
  - pDSA code level
- ▶ IMM network information:
  - IMM host name
  - IMM dedicated MAC address. Only the MAC address that is in use is displayed (dedicated or shared)
  - IMM shared MAC address
  - IP v4 information
  - IP v6 address
- ▶ System environmental information:
  - Ambient temperature
  - CPU temperature
  - AC input voltage
  - Estimated power consumption

### 3.6.3 Rear ports and LEDs

Figure 3-10 shows the ports and LEDs that are on the rear of the server.




Figure 3-10 Rear ports and LEDs on the Primary I/O book

The following components are featured in Figure 3-10:

- ▶ 3V battery error LED: When this LED is lit, it indicates that a Standard I/O Book battery error occurred.
- ▶ I/O board error LED: When this LED is lit, it indicates that an error occurred on the Primary I/O Book board.

- ▶ **System-error LED:** When this LED is lit, it indicates that a system error occurred. An LED on the front operator panel is also lit to help isolate the error. This LED is functionally equivalent to the system-error LED on the front of the server.
- ▶ **Locate LED:** Use this LED to visually locate the server among other servers. You can use IBM Systems Director to light this LED remotely. IMM can also be used to turn this LED on and off. This LED is functionally equivalent to the locate LED on the front of the server.
- ▶ **NMI button:** Press this button to force a nonmaskable interrupt to the microprocessor. You might have to use a pen or the end of a straightened paper clip to press the button. You can also use it to force a blue-screen memory dump. Use this button only when you are directed to do so by Lenovo Support.
- ▶ **Serial connector:** Connect a 9-pin serial device to this connector. The serial port is shared with the IMM. The IMM can take control of the shared serial port to redirect serial traffic by using Serial over LAN (SOL).
- ▶ **Systems-management Ethernet connector:** Use this connector to manage the server by using a dedicated management network. If you use this connector, the IMM cannot be accessed directly from the production network. A dedicated management network provides more security by physically separating the management network traffic from the production network. You can use the Setup utility to configure the server to use a dedicated systems-management network or a shared network.
- ▶ **USB 2.0 connectors:** Connect a USB device, such as a USB mouse, keyboard, or other device, to any of these connectors.
- ▶ **Video connector:** Connect a monitor to this connector. The video connectors on the front and rear of the server can be used simultaneously.

The x3850 X6 system has the following number of ports:

- ▶ **Front:** Two USB 3.0, one USB 2.0, and one DB-15 video ports.
- ▶ **Rear:** Four USB 2.0, one DB-15 video, one DB-9 serial, and one 1 GbE RJ-45 IMM2 systems management.
- ▶ **Internal:** One USB 2.0 port for embedded hypervisor.

The x3950 X6 system has the following number of ports when not partitioned (partitioning allows you to split the x3950 X6 into two separate virtual 4-socket systems):

- ▶ **Front:** Four USB 3.0, two USB 2.0, and one DB-15 video ports.
- ▶ **Rear:** Eight USB 2.0, one DB-15 video, two DB-9 serial, and one 1 GbE RJ-45 IMM2 systems management.
- ▶ **Internal:** One USB 2.0 port for embedded hypervisor.

The x3950 X6 system has the following number of ports when partitioned:

- ▶ **Front:** Four USB 3.0, two USB 2.0, and two DB-15 video ports.
- ▶ **Rear:** Eight USB 2.0, two DB-15 video, two DB-9 serial, and two 1 GbE RJ-45 IMM2 systems management.
- ▶ **Internal:** Two USB 2.0 port for embedded hypervisor.

## 3.7 Compute Books

In this section, we describe the central component of the X6 family server: the Compute Book. We describe the Compute Book's design and its main components and the Compute Book's placement order for the 4U and 8U chassis. This section includes the following topics:

- ▶ 3.7.1, "Compute Book design" on page 76
- ▶ 3.7.2, "Compute Book population order" on page 78

- ▶ 3.7.3, “Upgrading Compute Books” on page 79

The processor options are preinstalled in a Compute Book. For more information, see 3.8, “Processor options” on page 80.

### 3.7.1 Compute Book design

The main component of x3850 X6 and x3950 X6 servers is the Compute Book. It contains the following components:

- ▶ One Intel Xeon E7-4800 v3 or v4 processor, or E7-8800 v3 or v4 processor
- ▶ 24 DDR3 or DDR4 memory slots (model dependent), 12 on each side
- ▶ Two hot-swap dual-motor fans

#### **E7 v2 and v3 processors:**

- ▶ Models with E7 v3 processors are now withdrawn from marketing, however compute books are still available for field upgrades.
- ▶ Compute Books and server models with E7 v2 processors are now withdrawn from marketing.

Figure 3-11 shows the Compute Book.




Figure 3-11 Compute Book

As shown in Figure 3-11, the Compute Book has a cover with a transparent window, with which you can check light path diagnostic LEDs and view installed DIMMs without removing the cover.

The internal components of the Compute Book are shown in Figure 3-12. The system board of the Compute Book has two sides on which all components are installed.




Figure 3-12 Compute Book with the side covers removed

Figure 3-13 shows the Compute Book from the front where two hot-swap fans are installed.




Figure 3-13 Front view of the Compute Book, front and rear views of fans

As shown in Figure 3-13, the Compute Book has two buttons that are hidden behind the upper fan module: a light path diagnostic button and a slider to unlock the release handle.

Use the light path diagnostic button to determine a failed DIMM module on the Compute Book. The appropriate DIMM error LED on the Compute Book board should be lit. Use the slider to unlock the release handle so that you can remove the Compute Book from the chassis.

The system board of the Compute Book also contains a set of error LEDs for each memory module, which show failed modules when the light path diagnostic button is pressed. Figure 3-14 shows the DIMM error LEDs of the Compute Book's left side.

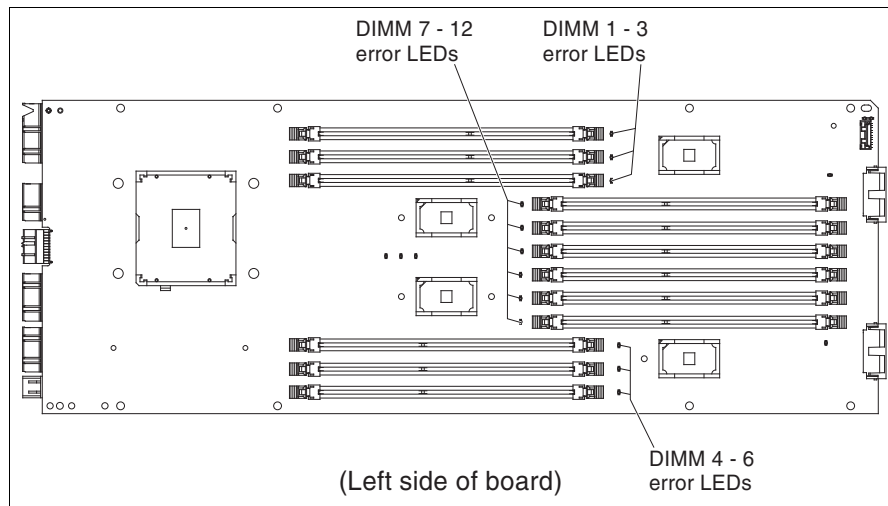


Figure 3-14 DIMM error LEDs placed on the left side of the system board

### 3.7.2 Compute Book population order

The x3850 X6 server supports the following combinations:

- ▶ 2 Compute Books
- ▶ 4 Compute Books

A configuration of 3 Compute Books is not supported. A configuration of one processor is supported in model A4x only. No other server model or configure to order (CTO) configuration supports one processor.

The x3950 X6 supports the following combinations:

- ▶ 4 Compute Books
- ▶ 6 Compute Books
- ▶ 8 Compute Books

Other combinations are not supported.



Figure 3-15 shows how the Compute Books are numbered.




Figure 3-15 Compute Book numbering

The Compute Books feature the following installation sequences:

- ▶ x3850 X6:
  - 1 Compute Book: Bay 1
  - 2 Compute Books: Bays 1 and 2
  - 4 Compute Books: All bays
- ▶ x3950 X6:
  - 4 Compute Books: Bays 1, 2, 5, and 6
  - 6 Compute Books: Bays 1, 2, 3, 4, 5 and 6
  - 8 Compute Books: All bays

### 3.7.3 Upgrading Compute Books

With the x3850 X6, machine type 6241, we support upgrading the server from Intel Xeon E7 v3 processors to Intel Xeon E7 v4 processors by replacing the Compute Books. The upgrade process is made simple by the innovative book design of the x3850 X6.

Consider the following points when upgrading from v3 to v4 processors:

- ▶ eXFlash DIMMs are not compatible with the v4 processors.
- ▶ Intel E7 v4 processors require minimum firmware levels. For more information, see the documentation that is included with the Compute Books.
- ▶ The TruDDR4 and DDR3 memory that is compatible with the v3 processors is not supported with the v4 processors. Only the new TruDDR4 memory options are compatible with the v4 processors. This means you will need to plan the following steps to upgrade the server:
  1. Purchase up to four new Compute Books.
  2. Order the appropriate memory DIMMs to match your workload requirements.

3. Check and upgrade all firmware to at least the minimum levels that are needed to support E7 v4 processors, if necessary.
4. Install the new TruDDR4 memory DIMMs to the new E7 v4 Compute Books.
5. Power off the server and remove the E7 v3 Compute Books.
6. Install the new E7 v4 Compute Books.

Consider the following points when upgrading from v2 to v3 processors:

- ▶ eXFlash DIMMs are not supported with TruDDR4 memory; eXFlash DIMMs are only supported with DDR3 memory.
- ▶ Intel Xeon E7 v3 processors require minimum firmware levels. For more information, see the documentation that is included with the Compute Books.

If you plan to use your existing DDR3 memory, complete the following steps to upgrade the server:

1. Purchase up to four new Compute Books.
2. Check and upgrade all firmware to at least the minimum levels that are needed to support E7 v3 processors, if necessary.
3. Power off the server and remove the old Compute Books.
4. Transfer all DDR3 memory DIMMs to the new E7 v3 Compute Books.
5. Reinstall the new E7 v3 Compute Books.

If you plan to upgrade to TruDDR4 memory, order the appropriate memory DIMMs to match your workload requirements.

## 3.8 Processor options

The x3850 X6 and x3950 X6 servers use Intel Xeon E7-4800 v3 or v4, or E7-8800 v3 or v4 processors, formerly known by the code names “Haswell-EX” for the v3 processors, and “Broadwell-EX” for the v4 processors.

### **E7 v2 and v3 processors:**

- ▶ Models with E7 v3 processors are now withdrawn from marketing, however compute books are still available for field upgrades.
- ▶ Compute Books and server models with E7 v2 processors are now withdrawn from marketing.

The Intel Xeon E7-4800 v3 series processors are available with up to 18 cores and 45 MB of last-level cache and can form a 4-socket configuration. The Intel Xeon E7-8800 v3 series processors are also available with up to 18 cores and 45 MB of last-level cache, but can be used in 4-socket and 8-socket configurations. Using E7-8800 processors enables Compute Books to be swapped between x3850 X6 and x3950 X6 servers.

The Intel Xeon E7-4800 v4 series processors are available with up to 16 cores and 40 MB of last-level cache and can form a 4-socket configuration. The Intel Xeon E7-8800 v4 series processors are also available with up to 24 cores and 60 MB of last-level cache, but can be used in 4-socket and 8-socket configurations.

Using E7-8800 processors enables Compute Books to be swapped between x3850 X6 and x3950 X6 servers.

Table 3-6 lists the processor options for the x3850 X6 and x3950 X6 servers, and which server models have each processor standard. If there is no corresponding where-used model for a particular processor, this processor is available through CTO only.

Table 3-6 also lists the processor options that are grouped in the following manner:

- ▶ E7-4800 v3 processors in Compute Books with TruDDR4 support
- ▶ E7-8800 v3 processors in Compute Books with TruDDR4 support
- ▶ E7-4800 v3 processors in Compute Books with DDR3 support
- ▶ E7-8800 v3 processors in Compute Books with DDR3 support

The processor options are shipped preinstalled in a Compute Book. All Compute Books in a server must have identical processors.

Table 3-6 Processor

Part number Feature codes <sup>a</sup>	Description (processor installed in a Compute Book)	L3 cache	Memory support	Memory bus speed (RAS / Perf) <sup>b</sup>	Supported quantities x3850 X6 <sup>c</sup>	Supported quantities x3950 X6
<b>Intel Xeon E7-4800 v4 processors (not supported in the x3950 X6) with support for TruDDR4 memory</b>						
00WH302 / ATX1 / ATXR	X6 Compute Book Intel Xeon Processor E7-4809 v4 8C 2.1GHz 20M 115W	20 MB	TruDDR4	1866 MHz / 1333 MHz	2, 4	No support
00WH306 / ATX2 / ATXS	X6 Compute Book Intel Xeon Processor E7-4820 v4 10C 2.0GHz 25M 115W	25 MB	TruDDR4	1866 MHz / 1333 MHz	2, 4	No support
00WH310 / ATX3 / ATXT	X6 Compute Book Intel Xeon Processor E7-4830 v4 14C 2.0GHz 35M 115W	35 MB	TruDDR4	1866 MHz / 1333 MHz	2, 4	No support
00WH314 / ATX4 / ATXU	X6 Compute Book Intel Xeon Processor E7-4850 v4 16C 2.1GHz 40M 115W	40 MB	TruDDR4	1866 MHz / 1333 MHz	2, 4	No support
<b>Intel Xeon E7-8800 v4 processors with support for TruDDR4 memory (supported in x3850/x3950 X6)</b>						
00KH384 / ATWZ / ATYF	X6 Compute Book Intel Xeon Processor E7-8855 v4 14C 2.1GHz 35M 140W	35 MB	TruDDR4	1866 MHz / 1333 MHz	2, 4	4, 6, 8
00WH322 / ATX6 / ATXW	X6 Compute Book Intel Xeon Processor E7-8860 v4 18C 2.2GHz 45M 140W	45 MB	TruDDR4	1866 MHz / 1600 MHz	2, 4	4, 6, 8
00WH318 / ATX5 / ATXV	X6 Compute Book Intel Xeon Processor E7-8867 v4 18C 2.4GHz 45M 165W	45 MB	TruDDR4	1866 MHz / 1600 MHz	2, 4	4, 6, 8
00WH326 / ATX7 / ATXX	X6 Compute Book Intel Xeon Processor E7-8870 v4 20C 2.1GHz 50M 140W	50 MB	TruDDR4	1866 MHz / 1600 MHz	2, 4	4, 6, 8
00WH330 / ATX8 / ATXY	X6 Compute Book Intel Xeon Processor E7-8880 v4 22C 2.2GHz 55M 150W	55 MB	TruDDR4	1866 MHz / 1600 MHz	2, 4	4, 6, 8
00WH334 / ATX9 / ATXZ	X6 Compute Book Intel Xeon Processor E7-8890 v4 24C 2.2GHz 60M 165W	60 MB	TruDDR4	1866 MHz / 1600 MHz	2, 4	4, 6, 8

Part number Feature codes <sup>a</sup>	Description (processor installed in a Compute Book)	L3 cache	Memory support	Memory bus speed (RAS / Perf) <sup>b</sup>	Supported quantities x3850 X6 <sup>c</sup>	Supported quantities x3950 X6
00WH342 / ATXB / ATY1	X6 Compute Book Intel Xeon Processor E7-8891 v4 10C 2.8GHz 60M 165W	60 MB	TruDDR4	1866 MHz / 1600 MHz	2, 4	4, 6, 8
00WH346 / ATXC / ATY2	X6 Compute Book Intel Xeon Processor E7-8893 v4 4C 3.2GHz 60M 140W	60 MB	TruDDR4	1866 MHz / 1600 MHz	2, 4	4, 6, 8
00YG935 AVR6 / AVR7	X6 Compute Book Intel Xeon Processor E7-8894 v4 24C 2.4GHz 60M 165W	60 MB	TruDDR4	1866 MHz / 1600 MHz	2, 4	4, 6, 8
<b>Intel Xeon E7-4800 v3 processors (not supported in the x3950 X6) with support for DDR3 memory</b>						
00ML902 AS7T / AS8H	X6 DDR3 Compute Book Intel Xeon Processor E7-4809 v3 8C 2.0GHz 115W	20 MB	DDR3	1333 MHz / 1333 MHz	2, 4	No support
00ML906 AS7U / AS8J	X6 DDR3 Compute Book Intel Xeon Processor E7-4820 v3 10C 1.9GHz 115W	25 MB	DDR3	1333 MHz / 1333 MHz	2, 4	No support
00ML910 AS7V / AS8K	X6 DDR3 Compute Book Intel Xeon Processor E7-4830 v3 12C 2.1GHz 115W	30 MB	DDR3	1600 MHz / 1333 MHz	2, 4	No support
00ML914 AS7W / AS8L	X6 DDR3 Compute Book Intel Xeon Processor E7-4850 v3 14C 2.2GHz 115W	35 MB	DDR3	1600 MHz / 1333 MHz	2, 4	No support
<b>Intel Xeon E7-8800 v3 processors (supported in x3850/x3950 X6) with support for DDR3 memory</b>						
00ML922 AS7X / AS8M	X6 DDR3 Compute Book Intel Xeon Processor E7-8860 v3 16C 2.2GHz 140W	40 MB	DDR3	1600 MHz / 1600 MHz	2, 4	4, 6, 8
00FP692 AS7R / AS8F	X6 DDR3 Compute Book Intel Xeon Processor E7-8867 v3 16C 2.5GHz 165W	45 MB	DDR3	1600 MHz / 1600 MHz	2, 4	4, 6, 8
00ML926 AS7Y / AS8N	X6 DDR3 Compute Book Intel Xeon Processor E7-8870 v3 18C 2.1GHz 140W	45 MB	DDR3	1600 MHz / 1600 MHz	2, 4	4, 6, 8
00ML930 AS7Z / AS8P	X6 DDR3 Compute Book Intel Xeon Processor E7-8880 v3 18C 2.3GHz 150W	45 MB	DDR3	1600 MHz / 1600 MHz	2, 4	4, 6, 8
00ML938 AS81 / AS8R	X6 DDR3 Compute Book Intel Xeon Processor E7-8880L v3 18C 2.0GHz 115W	45 MB	DDR3	1600 MHz / 1600 MHz	2, 4	4, 6, 8
00ML934 AS80 / AS8Q	X6 DDR3 Compute Book Intel Xeon Processor E7-8890 v3 18C 2.5GHz 165W	45 MB	DDR3	1600 MHz / 1600 MHz	2, 4	4, 6, 8
00ML942 AS82 / AS8S	X6 DDR3 Compute Book Intel Xeon Processor E7-8891 v3 10C 2.8GHz 165W	45 MB	DDR3	1600 MHz / 1600 MHz	2, 4	4, 6, 8

Part number Feature codes <sup>a</sup>	Description (processor installed in a Compute Book)	L3 cache	Memory support	Memory bus speed (RAS / Perf) <sup>b</sup>	Supported quantities x3850 X6 <sup>c</sup>	Supported quantities x3950 X6
00ML946 AS83 / AS8T	X6 DDR3 Compute Book Intel Xeon Processor E7-8893 v3 4C 3.2GHz 140W	45 MB	DDR3	1600 MHz / 1600 MHz	2, 4	4, 6, 8
<b>Intel Xeon E7-4800 v3 processors (not supported in the x3950 X6) with support for TruDDR4 memory</b>						
00ML950 AS84 / AS8U	X6 DDR4 Compute Book Intel Xeon Processor E7-4809 v3 8C 2.0GHz 115W	20 MB	TruDDR4	1866 MHz / 1333 MHz	2, 4	No support
00ML954 AS85 / AS8V	X6 DDR4 Compute Book Intel Xeon Processor E7-4820 v3 10C 1.9GHz 115W	25 MB	TruDDR4	1866 MHz / 1333 MHz	2, 4	No support
00ML958 AS86 / AS8W	X6 DDR4 Compute Book Intel Xeon Processor E7-4830 v3 12C 2.1GHz 115W	30 MB	TruDDR4	1866 MHz / 1333 MHz	2, 4	No support
00ML962 AS87 / AS8X	X6 DDR4 Compute Book Intel Xeon Processor E7-4850 v3 14C 2.2GHz 115W	35 MB	TruDDR4	1866 MHz / 1333 MHz	2, 4	No support
<b>Intel Xeon E7-8800 v3 processors (supported in x3850/x3950 X6) with support for TruDDR4 memory</b>						
00ML970 AS88 / AS8Y	X6 DDR4 Compute Book Intel Xeon Processor E7-8860 v3 16C 2.2GHz 140W	40 MB	TruDDR4	1866 MHz / 1600 MHz	2, 4	4, 6, 8
00FP696 AS7S / AS8G	X6 DDR4 Compute Book Intel Xeon Processor E7-8867 v3 16C 2.5GHz 165W	45 MB	TruDDR4	1866 MHz / 1600 MHz	2, 4	4, 6, 8
00ML974 AS89 / AS8Z	X6 DDR4 Compute Book Intel Xeon Processor E7-8870 v3 18C 2.1GHz 140W	45 MB	TruDDR4	1866 MHz / 1600 MHz	2, 4	4, 6, 8
00ML978 AS8A / AS90	X6 DDR4 Compute Book Intel Xeon Processor E7-8880 v3 18C 2.3GHz 150W	45 MB	TruDDR4	1866 MHz / 1600 MHz	2, 4	4, 6, 8
00ML986 AS8C / AS92	X6 DDR4 Compute Book Intel Xeon Processor E7-8880L v3 18C 2.0GHz 115	45 MB	TruDDR4	1866 MHz / 1600 MHz	2, 4	4, 6, 8
00ML982 AS8B / AS91	X6 DDR4 Compute Book Intel Xeon Processor E7-8890 v3 18C 2.5GHz 165W	45 MB	TruDDR4	1866 MHz / 1600 MHz	2, 4	4, 6, 8
00ML990 AS8D / AS93	X6 DDR4 Compute Book Intel Xeon Processor E7-8891 v3 10C 2.8GHz 165W	45 MB	TruDDR4	1866 MHz / 1600 MHz	2, 4	4, 6, 8
00ML994 AS8E / AS94	X6 DDR4 Compute Book Intel Xeon Processor E7-8893 v3 4C 3.2GHz 140W	45 MB	TruDDR4	1866 MHz / 1600 MHz	2, 4	4, 6, 8

a. The first feature code is for the first Compute Book; the second feature code is for all additional Compute Books.

b. The processors support two memory modes: RAS mode (also known as Lockstep mode), and Performance mode (also known as Independent mode). In Performance mode, the SMI2 link operates at twice the memory bus speed that is shown.

c. A configuration of one processor is supported in models 6241-A4x and 6241-D5x only. No other server model or CTO configuration supports one processor.

## 3.9 Memory

The x3850 X6 and x3950 X6 (6241) support DDR3 or TruDDR4 memory, depending on the Compute Books that are used in the server.

Lenovo TruDDR4 Memory uses the highest-quality components that are sourced from Tier 1 DRAM suppliers and only memory that meets the strict requirements of Lenovo is selected. It is compatibility tested and tuned for optimal System x performance and throughput.

TruDDR4 Memory has a unique signature that is programmed into the DIMM that enables System x servers to verify whether the memory that is installed is qualified or supported by Lenovo. Because TruDDR4 Memory is authenticated, certain extended memory performance features can be enabled to extend performance over industry standards.

Lenovo DDR3 memory is compatibility tested and tuned for optimal System x performance and throughput. Lenovo memory specifications are integrated into the light path diagnostics for immediate system performance feedback and optimum system uptime

Lenovo memory specifications are integrated into the light path diagnostics for immediate system performance feedback and optimum system uptime. From a service and support standpoint, Lenovo memory automatically assumes the system warranty, and Lenovo provides service and support worldwide.

As described in 2.4, “Memory” on page 33, the x3850 X6 and x3950 X6 support TruDDR4 memory operating at speeds up to 1866 MHz and DDR3 memory at speeds up to 1600 MHz (model dependent).

**Tip:** TruDDR4 2133 MHz and 2400 MHz DIMMs operate with Intel Xeon E7 processors, but only at up to 1866 MHz speeds.

The x3850 X6 supports up to 96 DIMMs when all processors are installed, 24 DIMMs per processor. The x3950 X6 supports up to 192 DIMMs when all processors are installed, 24 DIMMs per processor. Each processor has four memory channels that are implemented by using Scalable Memory Interface generation 2 (SMI2) chips, and the server implements three DIMMs per channel. The processor and the corresponding memory DIMM slots are on the Compute Book.

The Intel Xeon E7 processors support two memory modes: Performance mode and RAS (or Lockstep) mode. For more information, see 2.4.1, “Operational modes” on page 34.

The following rules apply when the memory configuration is selected:

- ▶ The server supports RDIMMs and LRDIMMs.
- ▶ Mixing different types of memory (RDIMMs and LRDIMMs) is not supported.
- ▶ Mixing different generations of memory (DDR3 and TruDDR4) is not supported.
- ▶ The X6 servers support 3 different memory types and frequencies; DDR3 1600 MHz, TruDDR4 2133 MHz and TruDDR4 2400 MHz.

The following rules apply when installing DIMMs in Compute Books:

- DDR3 1600 MHz DIMMs can be installed in Compute Books with v2 and v3 processors. They can not be intermixed with TruDDR4 memory.
- TruDDR4 2133 MHz can be installed in Compute Books with v3 or v4 processors. They *can* be intermixed with TruDDR4 2400 MHz memory with Compute Books running v4 processors.

- TruDDR4 2400 MHz can be installed in Compute Books with v4 processors.
- ▶ The maximum number of ranks per one memory channel is six with RDIMMs or 24 with LRDIMMs.
- ▶ In RAS (Lockstep) mode, DIMMs must be installed in a pair.
- ▶ The maximum quantity of DIMMs that can be installed in the server depends on the number of processors, DIMM type, rank, and operating voltage, as shown in the “Maximum qty supported” row in Table 3-10 on page 88.
- ▶ All DIMMs in the server operate at the same speed, which is determined as the lowest value of one of the following options:
  - Memory speed that is supported by the specific processor.
  - Lowest of maximum operating speeds for selected memory configuration that depends on rated speed, operating voltage, and quantity of DIMMs per channel, as shown under the “Maximum operating speed” section in Table 3-10 on page 88.
- ▶ Configurations with DDR3 memory also support eXFlash DIMMs (now withdrawn), which are also installed in the DIMM slots. For more information about these eXFlash DIMMs and the configuration rules, see 3.10, “Memory-channel storage” on page 92.

### 3.9.1 Memory options

In this section, we describe the available memory options, memory characteristics, and speeds depending on the selected memory mode.

Table 3-7 lists the memory options that are available for x3850 X6 and x3950 X6 servers.

Table 3-7 Memory options

Part number	Feature code	Description	Maximum supported (x3850 X6 / x3950 X6)	Standard models where used
<b>TruDDR4 RDIMMs and LRDIMMs - 2400 MHz (for use with E7 v4 processors)</b>				
46W0821	ATC8	8GB TruDDR4 Memory (1Rx4, 1.2V) PC4-19200 CL17 2400MHz LP RDIMM	96 / 192 (24 per CPU)	14x, 1RC
46W0829	ATCA	16GB TruDDR4 Memory (2Rx4, 1.2V) PC4-19200 CL17 2400MHz LP RDIMM	96 / 192 (24 per CPU)	All other v4 models
46W0833	ATCB	32GB TruDDR4 Memory (2Rx4, 1.2V) PC4-19200 CL17 2400MHz LP RDIMM	96 / 192 (24 per CPU)	8Cx, 8Fx, 8Hx, 8Jx
46W0841	ATGG	64GB TruDDR4 Memory (4Rx4, 1.2V) PC4-19200 PC4 2400MHz LP LRDIMM	96 / 192 (24 per CPU)	-
<b>TruDDR4 RDIMMs and LRDIMMs - 2133 MHz (for use with E7 v3 processors, also supported with E7 v4; 00KH391 only supported on certain E7 v3 processors)</b>				
46W0788	A5B5	8GB TruDDR4 Memory (1Rx4, 1.2V) PC4-17000 CL15 2133MHz LP RDIMM	96 / 192 (24 per CPU)	D5x
46W0796	A5B7	16GB TruDDR4 Memory (2Rx4, 1.2V) PC4-17000 CL15 2133MHz LP RDIMM	96 / 192 (24 per CPU)	F2x, F4x, G2x, G4x
95Y4808	A5UJ	32GB TruDDR4 Memory (2Rx4, 1.2V) PC4-17000 CL15 2133MHz LP RDIMM	96 / 192 (24 per CPU)	-

Part number	Feature code	Description	Maximum supported (x3850 X6 / x3950 X6)	Standard models where used
00KH391 <sup>a</sup>	AUF3	32GB TruDDR4 Memory (4Rx4, 1.2V) PC4-17000 CL15 2133MHz LP LRDIMM	96 / 192 (24 per CPU)	-
95Y4812	A5UK	64GB TruDDR4 Memory (4Rx4,1.2V) PC4-17000 CL15 2133MHz LP LRDIMM	96 / 192 (24 per CPU)	-
<b>DDR3 RDIMMs</b>				
00D5024 <sup>b</sup>	A3QE	4GB (1x4GB, 1Rx4, 1.35V) PC3L-12800 CL11 ECC DDR3 1600MHz LP RDIMM	96 / 192 (24 per CPU)	-
00D5036	A3QH	8GB (1x8GB, 1Rx4, 1.35V) PC3L-12800 CL11 ECC DDR3 1600MHz LP RDIMM	96 / 192 (24 per CPU)	A4x, B1x, B3x, C1x, C4x, D4x
46W0672	A3QM	16GB (1x16GB, 2Rx4, 1.35V) PC3L-12800 CL11 ECC DDR3 1600MHz LP RDIMM	96 / 192 (24 per CPU)	F1x, F3x, G1x, G3x
<b>DDR3 LRDIMMs</b>				
46W0676	A3SR	32GB (1x32GB, 4Rx4, 1.35V) PC3L-12800 CL11 ECC DDR3 1600MHz LP LRDIMM	96 / 192 (24 per CPU)	-

a. 00KH391 is supported only in Compute Books with E7-8880 v3 or E7-8890 v3 processors; no other memory DIMM is supported if 00KH391 is installed

b. 00D5024 is supported only in Compute Books with E7 v2 processors

Table 3-10 (for DDR3 memory) and Table 3-9 (for TruDDR4 memory) show the characteristics of the supported DIMMs and the memory speeds. The cells that are highlighted in gray indicate that the X6 servers support higher memory frequencies or larger memory capacity (or both) than the Intel processor specification defines.

**Memory speed:** In Performance mode, memory channels operate independently, and the SMI2 link operates at twice the DIMM speed. In RAS mode, two channels operate synchronously, and the SMI2 link operates at the DIMM speed.

Table 3-8 Maximum memory speeds: 2133 MHz TruDDR4 memory - x3850 / x3950 X6

Specification	TruDDR4 RDIMMs		TruDDR4 LRDIMMs
	Single rank	Dual rank	Quad rank
Rank	Single rank	Dual rank	Quad rank
Part number	46W0788 (8 GB)	46W0796 (16 GB) 95Y4808 (32 GB)	95Y4812 (64 GB)
Rated speed	2133 MHz	2133 MHz	2133 MHz
Rated voltage	1.2 V	1.2 V	1.2 V
Operating voltage	1.2 V	1.2 V	1.2 V
Maximum quantity <sup>a</sup>	96 / 192	96 / 192	96 / 192
Largest DIMM	8 GB	16 GB	64 GB
Max memory capacity	0.75 TB / 1.5 TB	1.5 TB / 3 TB	6 TB / 12 TB
<b>Maximum operating speed: Performance mode (2:1 mode; SMI2 link operates at twice the speed shown)</b>			
1 DIMM per channel	1600 MHz	1600 MHz	1600 MHz



Specification	TruDDR4 RDIMMs		TruDDR4 LRDIMMs
2 DIMM per channel	1600 MHz	1600 MHz	1600 MHz
3 DIMM per channel	1600 MHz	1600 MHz <sup>b</sup>	1600 MHz
<b>Maximum operating speed: RAS mode (1:1 mode; SMI2 link operates at the speed shown)</b>			
1 DIMM per channel	1866 MHz	1866 MHz	1866 MHz
2 DIMM per channel	1866 MHz	1866 MHz	1866 MHz
3 DIMM per channel	1600 MHz	1600 MHz <sup>b</sup>	1600 MHz

a. Maximum quantity that is supported is shown for all processors that are installed for x3850 / x3950 X6.

b. This speed is achieved only when Lenovo TruDDR4 memory is used.

Table 3-9 Maximum memory speeds: 2400 MHz TruDDR4 memory - x3850 / x3950 X6

Specification	TruDDR4 RDIMMs		TruDDR4 LRDIMMs
Rank	Single rank	Dual rank	Quad rank
Part number	46W0821 (8 GB)	46W0829 (16 GB) 46W0833 (32 GB)	46W0841 (64 GB)
Rated speed	2400 MHz	2400 MHz	2400 MHz
Rated voltage	1.2 V	1.2 V	1.2 V
Operating voltage	1.2 V	1.2 V	1.2 V
Maximum quantity <sup>a</sup>	96 / 192	96 / 192	96 / 192
Largest DIMM	8 GB	16 GB	64 GB
Max memory capacity	0.75 TB / 1.5 TB	1.5 TB / 3 TB	6 TB / 12 TB
<b>Maximum operating speed: Performance mode (2:1 mode; SMI2 link operates at twice the speed shown)</b>			
1 DIMM per channel	1600 MHz	1600 MHz	1600 MHz
2 DIMM per channel	1600 MHz	1600 MHz	1600 MHz
3 DIMM per channel	1600 MHz	1600 MHz <sup>b</sup>	1600 MHz
<b>Maximum operating speed: RAS mode (1:1 mode; SMI2 link operates at the speed shown)</b>			
1 DIMM per channel	1866 MHz	1866 MHz	1866 MHz
2 DIMM per channel	1866 MHz	1866 MHz	1866 MHz
3 DIMM per channel	1600 MHz	1600 MHz <sup>b</sup>	1600 MHz

a. Maximum quantity that is supported is shown for all processors that are installed for x3850 / x3950 X6.

b. This speed is achieved only when Lenovo TruDDR4 memory is used.

Table 3-10 Maximum memory speeds: DDR3 memory

DIMM specification	RDIMM				LR-DIMM			
	Single-rank DIMM		Dual-rank DIMM		Quad-rank LR-DIMM		8-rank LR-DIMM	
Ranks	Single-rank DIMM		Dual-rank DIMM		Quad-rank LR-DIMM		8-rank LR-DIMM	
Rated speed	1600 MHz		1600 MHz		1600 MHz		1333 MHz	
Rated voltage	1.35 V		1.35 V		1.35 V		1.35 V	
Operating voltage	1.35 V	1.5 V	1.35 V	1.5 V	1.35 V	1.5 V	1.35 V	1.5 V
Max. DIMM capacity	8 GB	8 GB	16 GB	16 GB	32 GB	32 GB	64 GB	64 GB
Max. memory capacity <sup>a</sup>	0.75 TB / 1.5 TB	0.75 TB / 1.5 TB	1.5 TB / 3 TB	1.5 TB / 3 TB	3 TB / 6 TB	3 TB / 6 TB	6 TB / 12 TB	6 TB / 12 TB
<b>Maximum operating speed - Performance mode (2:1 mode - SMI2 link operates at twice the DDR3 speed shown)</b>								
1 DIMM / channel	1333 MHz	1333 MHz	1333 MHz	1333 MHz	1333 MHz	1333 MHz	1333 MHz	1333 MHz
2 DIMMs / channel	1333 MHz	1333 MHz	1333 MHz	1333 MHz	1333 MHz	1333 MHz	1333 MHz	1333 MHz
3 DIMMs / channel	1066 MHz	1333 MHz	1066 MHz	1333 MHz	1333 MHz	1333 MHz	1333 MHz	1333 MHz
<b>Maximum operating speed - RAS mode (1:1 mode - SMI2 link operates at the DDR3 speed shown)</b>								
1 DIMM / channel	1333 MHz	1600 MHz	1333 MHz	1600 MHz	1333 MHz	1600 MHz	1333 MHz	1333 MHz
2 DIMMs / channel	1333 MHz	1600 MHz	1333 MHz	1600 MHz	1333 MHz	1600 MHz	1333 MHz	1333 MHz
3 DIMMs / channel	1066 MHz	1333 MHz	1066 MHz	1333 MHz	1333 MHz	1333 MHz	1333 MHz	1333 MHz

a. Maximum memory capacity for the x3850 X6 server / Maximum memory capacity for the x3950 X6 server.

The following memory protection technologies are supported:

- ▶ ECC
- ▶ Chipkill (for x4-based memory DIMMs)
- ▶ Redundant bit steering (Double Device Data Correction)
- ▶ Memory mirroring
- ▶ Memory rank sparing

Chipkill and Redundant Bit Steering are supported in RAS mode. Chipkill is supported in Performance mode.

If memory mirroring is used, DIMMs must be installed in pairs for Performance mode (minimum of one pair per each processor) and quads for RAS mode. DIMMs in the pair or quad must be identical in type and size.

If memory rank sparing is used, a minimum of 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.

### 3.9.2 Memory population order

The x3850 X6 and x3950 X6 servers have two main memory operation modes: Performance mode and RAS (Lockstep) mode, and also two other modes: Mirroring and Rank Sparring.

The main operation modes can be combined with extra modes; therefore, the following operation modes are available:

- ▶ Performance (also referred to as Independent mode)
- ▶ RAS (also referred to as Lockstep mode)
- ▶ Performance + Mirroring
- ▶ Performance + Rank Sparring
- ▶ RAS + Mirroring
- ▶ RAS + Rank Sparring

For more information about memory operation modes, see 2.4.1, “Operational modes” on page 34 and 2.4.2, “Memory mirroring and rank sparing” on page 38.

Depending on the selected operating mode, you should use the appropriate DIMMs installation order. Table 3-11 lists the DIMMs placement order that depends on the operation mode.

The DIMM installation order for Rank Sparring mode follows the Performance mode or RAS (Lockstep) mode installation order based on the operation mode selected.

Table 3-11 DIMMs placement order

DIMM installation order	DIMM slots			
	Performance Mode	Performance + Mirroring	RAS Mode	RAS + Mirroring
1	DIMM 9	DIMM 9, 19	DIMM 9, 15	DIMM 1, 9, 15, 19
2	DIMM 6	DIMM 6, 16	DIMM 6, 24	DIMM 6, 10, 16, 24
3	DIMM 1	DIMM 1, 15	DIMM 1, 19	DIMM 2, 8, 14, 20
4	DIMM 10	DIMM 10, 24	DIMM 10, 16	DIMM 5, 11, 17, 23
5	DIMM 15	DIMM 8, 20	DIMM 8, 14	DIMM 3, 7, 13, 21
6	DIMM 24	DIMM 5, 17	DIMM 5, 23	DIMM 4, 12, 18, 22
7	DIMM 19	DIMM 2, 14	DIMM 2, 20	Not applicable
8	DIMM 16	DIMM 11, 23	DIMM 11, 17	Not applicable
9	DIMM 8	DIMM 7, 21	DIMM 7, 13	Not applicable
10	DIMM 5	DIMM 4, 18	DIMM 4, 22	Not applicable
11	DIMM 2	DIMM 3, 13	DIMM 3, 21	Not applicable
12	DIMM 11	DIMM 12, 22	DIMM 12, 18	Not applicable
13	DIMM 14	Not applicable	Not applicable	Not applicable
14	DIMM 23	Not applicable	Not applicable	Not applicable
15	DIMM 20	Not applicable	Not applicable	Not applicable
16	DIMM 17	Not applicable	Not applicable	Not applicable

DIMM installation order	DIMM slots			
	Performance Mode	Performance + Mirroring	RAS Mode	RAS + Mirroring
17	DIMM 7	Not applicable	Not applicable	Not applicable
18	DIMM 4	Not applicable	Not applicable	Not applicable
19	DIMM 3	Not applicable	Not applicable	Not applicable
20	DIMM 12	Not applicable	Not applicable	Not applicable
21	DIMM 13	Not applicable	Not applicable	Not applicable
22	DIMM 22	Not applicable	Not applicable	Not applicable
23	DIMM 21	Not applicable	Not applicable	Not applicable
24	DIMM 18	Not applicable	Not applicable	Not applicable

As shown in Table 3-11, DIMMs must be installed in pairs (minimum of one pair per CPU) in Mirroring mode. Both DIMMs in a pair must be identical in type and size.

Chipkill is supported in RAS mode only (Performance mode is not supported), and DIMMs must be installed in a pair.

Figure 3-16 shows the DIMM slot numbering on each side of the Compute Book.

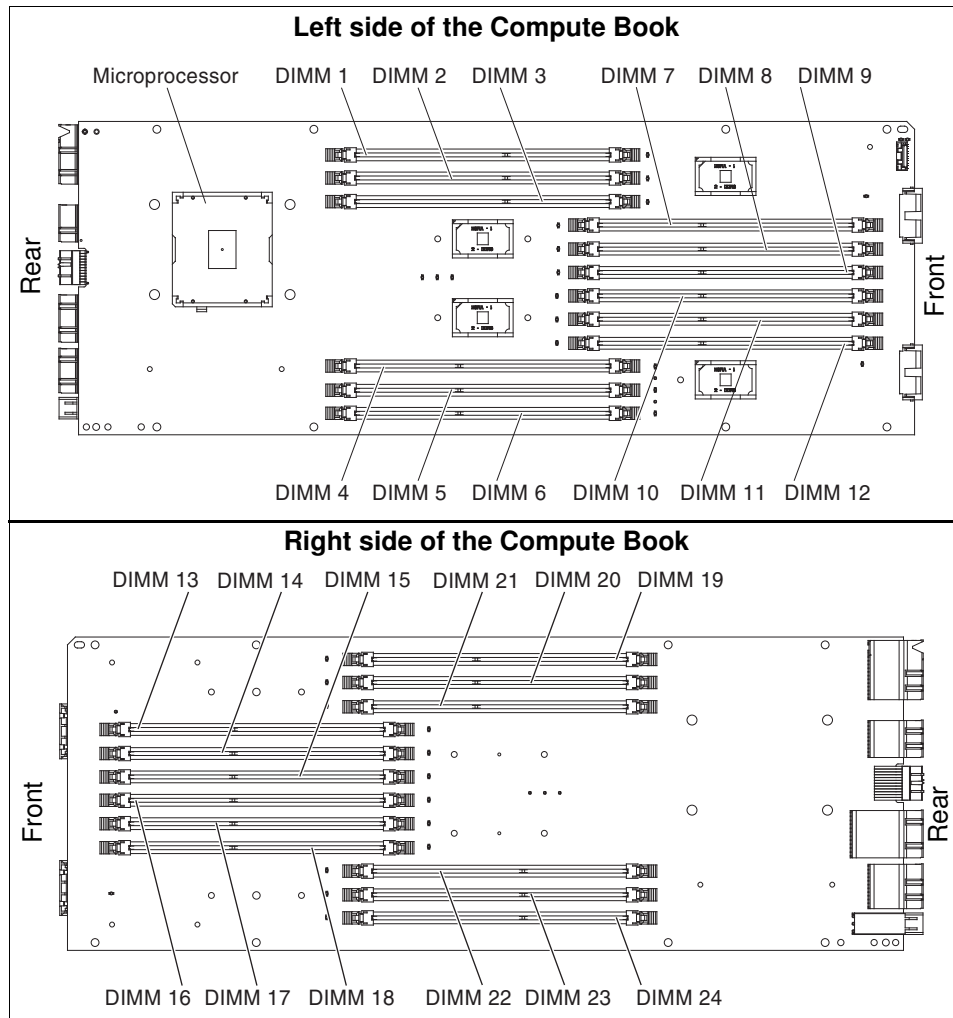


Figure 3-16 DIMM slot numbering

For Rank Sparring, the spare rank must have identical or larger memory capacity than all the other ranks (sparing source ranks) on the same channel. A minimum of two DIMMs per channel is required if only single rank DIMMs are used. The DIMM installation order must be modified to accommodate single rank only configurations.

Rank Sparring can be supported by 1 DIMM per channel if multi-rank DIMMs are used. Configurations that consist of multiple CPU sockets must have matching memory configurations on all CPUs because memory sparing is globally applied across all CPU sockets when enabled.

If you use eXFlash DIMMs, see 3.10, "Memory-channel storage" on page 92 for more information about the population order of eXFlash DIMMs.

## 3.10 Memory-channel storage

**Note:** The eXFlash DIMM options were withdrawn from marketing, effective March 31, 2016. The following information is for reference only.

Lenovo eXFlash memory-channel storage is an innovative ultrafast flash storage offering from Lenovo in the physical form of memory DIMMs. eXFlash DIMMs are installed in memory DIMM sockets but appear to the operating system and applications as block storage devices.

Data transfers between processors and eXFlash DIMMs run directly without any extra controllers, such as PCIe controllers and SAS/SATA controllers. This brings storage electrically closer to the processor subsystem, which significantly reduces latency and improves performance. Figure 3-17 shows the eXFlash DIMM.




Figure 3-17 eXFlash DIMM

Table 3-12 shows the available Lenovo eXFlash DIMM options.

Table 3-12 Lenovo eXFlash Memory Channel Storage options

Part number	Feature code	Name and description	Maximum supported x3850 X6	Maximum supported x3950 X6
00FE000 <sup>a</sup>	A4GX	eXFlash 200GB DDR3 Storage DIMM	32	32
00FE005 <sup>a</sup>	A4GY	eXFlash 400GB DDR3 Storage DIMM	32	32

a. Withdrawn from marketing

eXFlash DIMM modules are available in DDR3 form-factor only and can be installed in the same DIMM slots on the Compute Book as regular DDR3 DIMMs. Compute Books with TruDDR4 memory (v3 or v4 processors) do not support eXFlash DIMMs.

Figure 3-18 on page 93 shows one eXFlash DIMM installed with RDIMMs in the Compute Book.



Figure 3-18 eXFlash DIMM installed in the Compute Book

The eXFlash memory-channel storage includes the following key features:

- Less than 5 microseconds response time
- Less wait time between transactions
- Deterministic response time across varying workloads
- Tight standard deviation on performance
- Consistent performance for highest throughput and speed
- Offers up to 12.8 TB total flash storage capacity per server with 200 GB and 400 GB eXFlash DIMMs.
- Provides ultrahigh overall performance by using array of eXFlash DIMMs in a parallel manner.

The following rules apply when a server configuration is built with eXFlash DIMMs:

- ▶ Only DDR3 Compute Books are supported. Compute Books with TruDDR4 memory (v3 or v4 processors) do not support eXFlash DIMMs.
- ▶ 200 GB and 400 GB eXFlash DIMMs cannot be mixed.
- ▶ Performance memory mode must be selected; RAS (Lockstep) memory mode is not supported.
- ▶ Only RDIMMs are supported by eXFlash DIMMs; LRDIMMs are not supported.
- ▶ The following maximum quantities of eXFlash DIMMs in an x3850 X6 are supported:
  - 1 processor: 8 eXFlash DIMMs
  - 2 processors: 16 eXFlash DIMMs
  - 4 processors: 32 eXFlash DIMMs
- ▶ The following maximum quantities of eXFlash DIMMs in an x3950 X6 are supported:
  - 4 processors: 16 eXFlash DIMMs
  - 6 processors: 24 eXFlash DIMMs
  - 8 processors: 32 eXFlash DIMMs

## 3.11 Storage subsystem

In this section, we describe the storage subsystem of X6 servers, including Storage Book, supported RAID adapters, HDDs and SSDs. We also describe high-performance High IOPS adapters and PCIe Flash Storage Adapters that are designed for heavy I/O workloads.

This section includes the following topics:

- ▶ 3.11.1, “Storage Book”
- ▶ 3.11.2, “Backplanes” on page 96
- ▶ 3.11.3, “RAID controllers” on page 97
- ▶ 3.11.4, “Disk drive options” on page 99
- ▶ 3.11.5, “PCIe Flash Storage Adapters” on page 99

### 3.11.1 Storage Book

The Storage Book is another key component of module architecture of the X6 servers and provides the internal drive storage components. The x3850 X6 server contains one Storage Book and the x3950 X6 has two Storage Books.

The Storage Book is accessible from the front of the server and contains the following components:

- ▶ Up to eight 2.5-inch drives or 16 1.8-inch SSDs or up to four NVMe drives
- ▶ Up to two PCIe 3.0 x8 RAID-adapters
- ▶ Operator panel with LCD system information panel
- ▶ Two USB 3.0 ports
- ▶ One USB 2.0 port
- ▶ One SVGA port

Figure 3-19 shows the front view of the Storage Book.




Figure 3-19 Front view of a Storage Book



The Storage Book supports up to eight HDDs (2.5-inch format), or up to 16 1.8-inch SSDs, or up to 4 NVMe drives (2.5-inch format) in the following combinations:

- ▶ 4x 2.5-inch HDD hot-swap drives
- ▶ 8x 2.5-inch HDD hot-swap drives
- ▶ 8x 1.8-inch hot-swap SSDs
- ▶ 16x 1.8-inch hot-swap SSDs
- ▶ 4x 2.5-inch HDD hot-swap drives + 8x 1.8-inch hot-swap SSD bays
- ▶ 4x 2.5-inch PCIe NVMe SSDs
- ▶ 2x 2.5-inch PCIe NVMe SSDs + 8x 1.8-inch HDD hot-swap drives
- ▶ 2x 2.5-inch PCIe NVMe SSDs + 4x 2.5-inch HDD hot-swap drives

Figure 3-20 shows all available combinations of drives in the Storage Book.

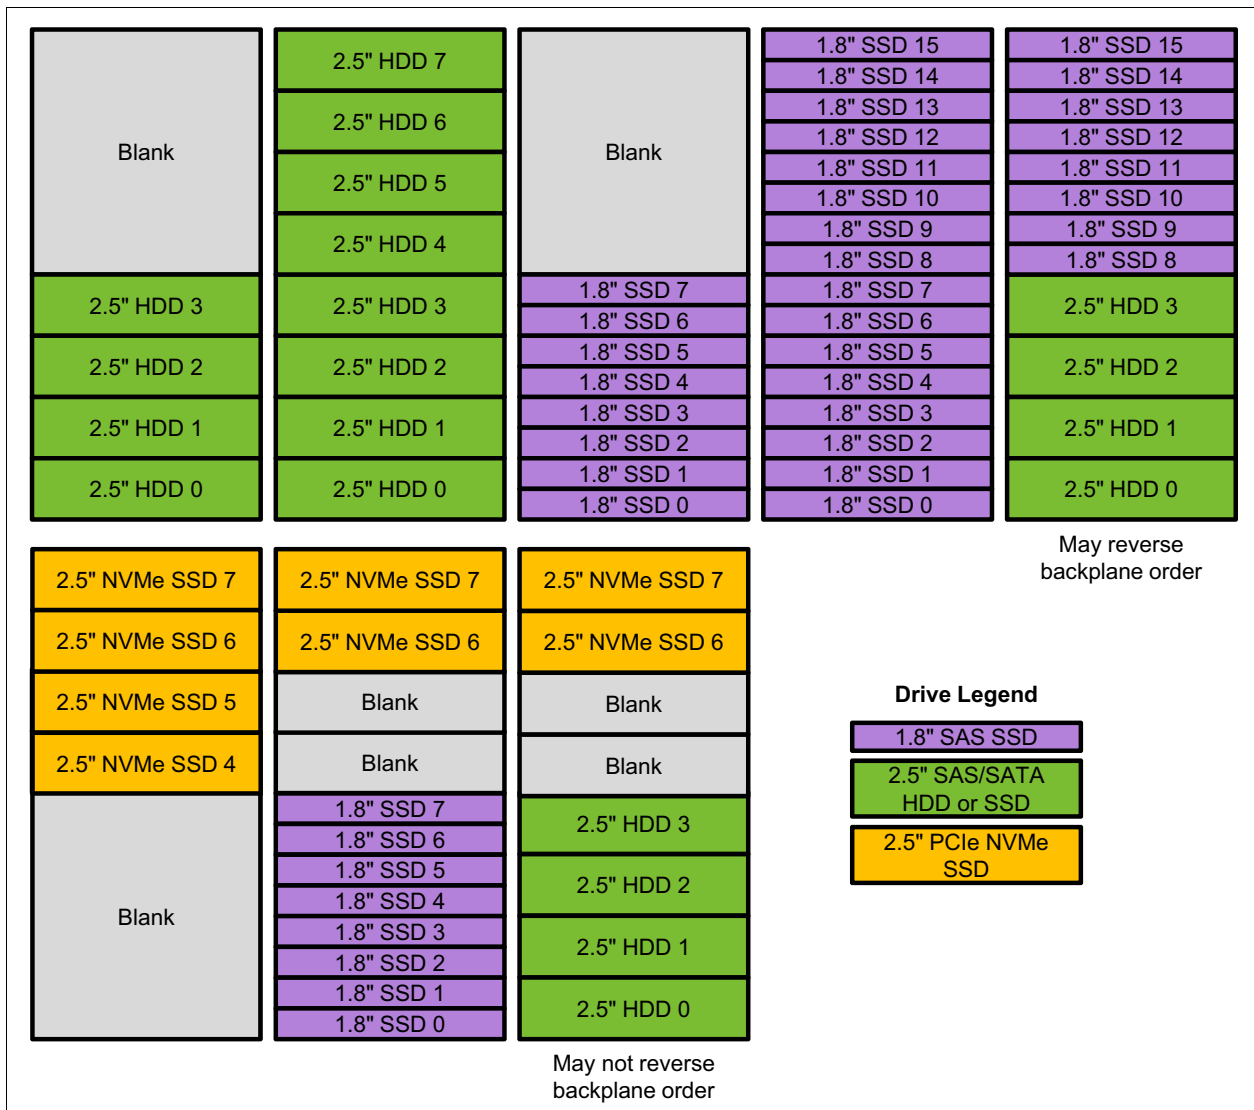


Figure 3-20 Storage Book population capabilities

### 3.11.2 Backplanes

The drives in each Storage Book are connected by backplanes. The following types of backplanes are available:

- ▶ Connect four 2.5-inch hot-swap HDDs
- ▶ Connect eight 1.8-inch hot-swap SSDs
- ▶ Connect four 2.5-inch PCIe NVMe SSDs

Each Storage Book can house two backplanes. All standard models ship with one backplane with four 2.5-inch hot-swap drive bays.

Figure 3-21 shows the location of SAS backplanes in the Storage Book and the RAID controllers that connect to the backplanes.




Figure 3-21 SAS backplanes location in the Storage Book

Table 3-13 lists the three backplanes.

Table 3-13 Backplane options for internal storage expansion

Part number	Feature code	Name and description	Maximum supported x3850 X6	Maximum supported x3950 X6
44X4104	A4A6	4x 2.5-inch HS 12Gb SAS HDD Backplane	2	4
44X4106	A4A7	8x 1.8-inch HS 12Gb SAS HDD Backplane	2	4
44X4108	A4A8	1x4 2.5-inch NVMe PCIe Gen3 SSD Backplane	1	2

Table 3-14 lists the required SAS backplanes for all combinations of drives that are supported by the X6 Storage Book.

Table 3-14 Drive and backplane combinations per Storage Book

Drive bay combination	Backplanes required	Controllers needed
4x 2.5-inch HDD/SSD	One 4x 2.5-inch HS 12Gb SAS HDD Backplane, 44X4104	1
8x 2.5-inch HDD/SSD	Two 4x 2.5-inch HS 12Gb SAS HDD Backplane, 44X4104	1
4x 2.5-inch HDD/SSD + 8x 1.8-inch SSD	<ul style="list-style-type: none"> <li>▶ One 4x 2.5-inch HS 12Gb SAS HDD Backplane, 44X4104</li> <li>▶ One 8x 1.8-inch HS 12Gb SAS HDD Backplane, 44X4106</li> </ul>	2
8x 1.8-inch SSD	One 8x 1.8-inch HS 12Gb SAS HDD Backplane, 44X4106	1
16x 1.8-inch SSD	Two 8x 1.8-inch HS 12Gb SAS HDD Backplane, 44X4106	2
4x 2.5-inch NVMe	One 4x 2.5-inch NVMe PCIe Gen3 SSD Backplane, 44X4108	1 or 2 <sup>a</sup>
2x 2.5-inch NVMe + 4x 2.5-inch HDD/SSD	<ul style="list-style-type: none"> <li>▶ One 4x 2.5-inch NVMe PCIe Gen3 SSD Backplane, 44X4108</li> <li>▶ One 4x 2.5-inch HS 12Gb SAS HDD Backplane, 44X4104</li> </ul>	2
2x 2.5-inch NVMe + 8x 1.8-inch SSD	<ul style="list-style-type: none"> <li>▶ One 4x 2.5-inch NVMe PCIe Gen3 SSD Backplane, 44X4108</li> <li>▶ One 8x 1.8-inch HS 12Gb SAS HDD Backplane, 44X4106</li> </ul>	2

a. The number of controllers that is required depends on the number of drives installed: one or two drives requires one controller; three or four drives requires two controllers.

### 3.11.3 RAID controllers

The Storage Book contains two PCIe 3.0 x8 slots for RAID controllers or SAS host bus adapters (HBAs) for use with the internal drives. If two SSD units (16x 1.8-inch SSDs) are used, two controllers are required.

Figure 3-22 shows the Storage Book's internal components and one of the RAID adapters.




Figure 3-22 Storage Book side view

Table 3-15 lists the RAID controllers, HBAs, and other hardware and feature upgrades that are used for internal disk storage. The adapters are installed in slots in the Storage Book.

Table 3-15 RAID controllers and HBAs for internal storage

Part number	Feature code	Description	Maximum supported x3850 X6	Maximum supported x3950 X6	Models where used
<b>Controllers</b>					
46C9110	A3YZ	ServeRAID M5210 SAS/SATA Controller	2	4	All
47C8675	A3YY	N2215 SAS/SATA HBA	2	4	-
00ML997	AS95	NVMe PCIe SSD Extender Adapter	2	4	-
<b>Hardware upgrades for the M5210</b>					
47C8656	A3Z0	ServeRAID M5200 Series 1 GB Cache/RAID 5 Upgrade	2	4	-
47C8660	A3Z1	ServeRAID M5200 Series 1 GB Flash/RAID 5 Upgrade	2	4	-
47C8664	A3Z2	ServeRAID M5200 Series 2 GB Flash/RAID 5 Upgrade	2	4	-
47C8668	A3Z3	ServeRAID M5200 Series 4 GB Flash/RAID 5 Upgrade	2	4	-
<b>Feature on-Demand upgrades for the M5210</b>					
47C8708	A3Z6	ServeRAID M5200 Series Zero Cache/RAID 5 Upgrade	1	2	-
47C8706	A3Z5	ServeRAID M5200 Series RAID 6 Upgrade	1	2 <sup>a</sup>	-
47C8710	A3Z7	ServeRAID M5200 Series Performance Accelerator	1	2 <sup>a</sup>	-
47C8712	A3Z8	ServeRAID M5200 Series SSD Caching Enabler	1	2 <sup>a</sup>	-

a. Requires cache memory upgrade (47C8656, 47C8660, or 47C8664)

The ServeRAID M5210 SAS/SATA Controller features the following specifications:

- ▶ Eight internal 12 Gbps SAS/SATA ports
- ▶ Two x4 HD mini-SAS internal connectors (SFF-8643)
- ▶ Supports connections to SAS/SATA drives and SAS Expanders
- ▶ Supports RAID levels 0, 1, and 10
- ▶ Supports RAID levels 5 and 50 with optional M5200 Series RAID 5 upgrades
- ▶ Supports RAID 6 and 60 with the optional M5200 Series RAID 6 Upgrade
- ▶ Supports 1 GB non-backed cache or 1 GB or 2 GB flash-backed cache
- ▶ Up to 12 Gbps throughput per port
- ▶ PCIe 3.0 x8 host interface
- ▶ Based on the LSI SAS3108 12 Gbps ROC controller

The N2215 SAS/SATA HBA features the following specifications:

- ▶ Eight internal 12 Gbps SAS/SATA ports
- ▶ Two x4 HD mini-SAS internal connectors (SFF-8643)
- ▶ Supports connections to SAS/SATA HDDs and SATA SSDs
- ▶ Optimized for SSD performance
- ▶ No RAID support
- ▶ Up to 12 Gbps throughput per port
- ▶ PCIe 3.0 x8 host interface
- ▶ Based on the LSI SAS3008 12 Gbps controller

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

<http://lenovopress.com/systemx/raid>

The NVMe PCIe SSD Extender Adapter is a PCIe adapter that routes the PCIe 3.0 x8 signal from the slot in the Storage Book, splits the signal into two x4 links, and routes those via cables to the drive backplane. The extender adapter is shown in Figure 3-23 on page 99.

One adapter is required for every two drives that are installed in the Storage Book. Because there are only two PCIe slots in the Compute Book, only two adapters can be installed; therefore, only four NVMe drives can be installed in a Storage Book.

For the x3950 X6 with two Compute Books, a total of four NVMe PCIe SSD Extender Adapter can be installed plus a total of four NVMe drives.



Figure 3-23 NVMe PCIe SSD Extender Adapter

### 3.11.4 Disk drive options

For a list of currently available internal drive options, see the Lenovo Press product guides:

- ▶ x3850 X6: <https://lenovopress.com/tips1250-x3850-x6-6241#internal-drive-options>
- ▶ x3950 X6: <https://lenovopress.com/tips1251-x3950-x6-6241#internal-drive-options>

**Note:** The drive IDs that are assigned by IMM2 match the IDs that are indicated on the server front bezel.

The operating system and UEFI report the HDDs that are attached to the 4x2.5-inch NVMe PCIe Gen3 SSD backplane as PCI devices.

### 3.11.5 PCIe Flash Storage Adapters

The PCIe Flash Storage Adapters (also known as High IOPs adapters and flash adapters) are high-performance flash and NVMe SSD storage adapters with low latency. Depending on the workload, the Flash Storage Adapter can deliver throughput of up to 900,000 IOPS and latency less than 100 microseconds.

Table 3-16 lists available High IOPS adapter options.

**Withdrawn:** All supported Flash Storage Adapters are now withdrawn from marketing.

Table 3-16 High IOPS adapter options

Part number	Feature code	Description	Maximum supported x3850 X6	Maximum supported x3950 X6
00YA800 <sup>a</sup>	AT7N	io3 1.25TB Enterprise Mainstream Flash Adapter	6	12
00YA803 <sup>a</sup>	AT7P	io3 1.6TB Enterprise Mainstream Flash Adapter	6	12
00YA806 <sup>a</sup>	AT7Q	io3 3.2TB Enterprise Mainstream Flash Adapter	6	12
00YA809 <sup>a</sup>	AT7R <sup>b</sup>	io3 6.4TB Enterprise Mainstream Flash Adapter	Nine	18
00YA812 <sup>a</sup>	AT7L	Intel P3700 1.6TB NVMe Enterprise Performance Flash Adapter	Nine	18
00YA815 <sup>a</sup>	AT7M	Intel P3700 2.0TB NVMe Enterprise Performance Flash Adapter	Nine	18
00AE995 <sup>a</sup>	ARYP	1000GB Enterprise io3 Flash Adapter	6	12
00AE998 <sup>a</sup>	ARYQ	1300GB Enterprise io3 Flash Adapter	6	12
00JY004 <sup>a</sup>	ARYS <sup>b</sup>	5200GB Enterprise io3 Flash Adapter	Nine	18
00JY001 <sup>a</sup>	ARYR	2600GB Enterprise io3 Flash Adapter	6	12

a. Withdrawn from marketing

b. These adapters cannot be ordered via CTO. The adapters can be ordered as an option only. The server cannot be shipped with this adapter installed. For more information, see this page:

<https://support.lenovo.com/documents/serv-io3>

For more information about the Flash Storage Adapters, see 2.6.4, “Flash Storage Adapters” on page 47.

## 3.12 I/O subsystem

In this section, we describe I/O Books, which are the third type of modular component that is used in the x3850 X6 and x3950 X6.

Subsequent sections also include relevant topics:

- ▶ 3.13, “Primary I/O Book” on page 102
- ▶ 3.14, “Half-length I/O Books and Full-length I/O Books” on page 104
- ▶ 3.15, “Hot-swap adapter support” on page 108
- ▶ 3.16, “Network adapters” on page 109
- ▶ 3.17, “SAS adapters for external storage” on page 112
- ▶ 3.18, “Fibre Channel host bus adapters” on page 114
- ▶ 3.19, “GPU adapters and co-processors” on page 114

The I/O Books provide many of the server ports and most of the PCIe adapter slots. The following types of I/O Books are available:

- ▶ Primary I/O Book, which is a core component of the server and consists of four PCIe slots and system logic, such as IMM2 and UEFI:
  - The x3850 X6 has one Primary I/O Book standard
  - The x3950 X6 has two Primary I/O Books standard

- ▶ Optional I/O Books: Half-length I/O Books or Full-length I/O Books, which provide three slots each:
  - The x3850 X6 supports two optional I/O Books
  - The x3950 X6 supports four optional I/O Books

The I/O Books are accessible from the rear of the server. To access the slots, release the locking handle and pull each book out from the rear of the server.

Figure 3-24 shows the rear view of the x3850 X6 server where you can see the Primary I/O Book and Optional I/O Books.




Figure 3-24 I/O Books accessible from the rear of the server

The Storage Book also contains two PCIe slots, as described in 3.11.1, “Storage Book” on page 94.

The PCIe lanes that are used in the I/O Books and Storage Book are connected to installed processors in the following configurations:

- ▶ The slots in the Primary I/O Book and Storage I/O Book connect to processor 1 or processor 2.
- ▶ The slots in the optional I/O Book in bay 1 connect to processor 4.
- ▶ The slots in the optional I/O Book in bay 2 connect to processor 3.

This PCIe connectivity is shown in Figure 3-25.

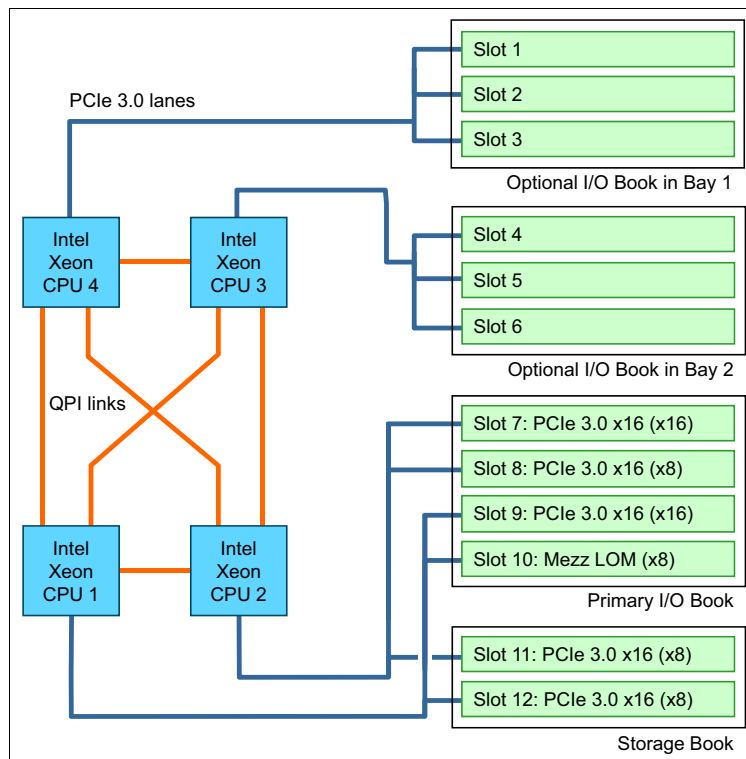


Figure 3-25 PCIe connectivity in the x3850 X6

### 3.13 Primary I/O Book

The Primary I/O Book is the core module, which provides base I/O connectivity with the following PCIe slots:

- ▶ Two PCIe 3.0 x16 slots (x16-wired), half-length, full-height, up to 75 W of power
- ▶ One PCIe 3.0 x16 (x8-wired), half-length, full-height, up to 75 W of power
- ▶ One Mezzanine LOM Generation 2 (ML2) slot for network adapters with the new ML2 form factor (PCIe 3.0 x8)

As shown in Figure 3-25, the use of all slots in the Primary I/O Book requires two processors (two Compute Books) installed by using the following configuration:

- ▶ Processor 1 drives a PCIe x16 slot and the ML2 slot
- ▶ Processor 2 drives the other PCIe 16 slot and the PCI x8 slot

Remember, the x3850 X6 has *one* Primary I/O Book, while the x3950 X6 has *two* Primary I/O Books.

The Primary I/O Book also contains the following core logic:

- ▶ IMM2
- ▶ TPMs
- ▶ Two hot-swap fan modules
- ▶ Internal USB port for an embedded hypervisor
- ▶ Dedicated 1 Gigabit Ethernet port dedicated for IMM2 connectivity
- ▶ Four USB 2.0 ports



- ▶ VGA port
- ▶ Serial port

Figure 3-26 shows the Primary I/O Book location and its components.




Figure 3-26 Primary I/O Book location and ports

Rather than a fixed network controller integrated on the system board, the x3850 X6 and x3950 X6 offer a dedicated ML2 (mezzanine LOM 2) adapter slot where you can select (for example) a 4-port Gigabit controller or a 2-port 10Gb controller from various networking vendors, as listed in Table 3-19 on page 110. This slot also supports out-of-band connectivity to the IMM2 management controller.

The Primary I/O Book has three PCIe 3.0 slots for optional PCIe adapters. All three ports have PCIe x16 physical form-factor, but only PCIe slots 7 and 9 have 16 PCIe 3.0 lines; PCIe slot 8 works as PCIe 3.0 x8. All PCIe ports in the Primary I/O Book support only half-length full-height PCIe adapters. Maximum power consumption for each PCIe slot is 75 W. To use full-length adapters, add one or two Full-Length I/O Books.

Figure 3-27 shows the Primary I/O Book removed from the server and the internal components.




Figure 3-27 Primary I/O Book removed showing the internal components

As shown in Figure 3-27, the Primary I/O Book has an internal USB port for the embedded hypervisor.

The Primary I/O Book also has a large plastic air baffle inside (the baffle is raised on a hinge in Figure 3-27), which routes hot air from the Storage Book through the two fans in the Primary I/O Book.

The dedicated IMM2 systems management port is a dedicated Gigabit Ethernet port that connects to the IMM2. This port is useful if you have an isolated management network. If you conduct management activities over your production network, you might want instead to use Port 1 of the ML2 adapter, which can be shared between the operating system and IMM2.

Although the x3950 X6 has two Primary I/O Books, the second video, IMM2, and internal USB hypervisor ports are used only when the x3950 X6 is partitioned into two virtual 4-socket servers.

The Primary I/O Book is *not* a hot-swappable component of the server.

## 3.14 Half-length I/O Books and Full-length I/O Books

The Half-length I/O Books and Full-length I/O Books, collectively referred to as optional I/O Books, each add three PCIe slots to the server.

As with the Primary I/O Book, the optional I/O Books are also installed from the rear side of the server. Figure 3-28 on page 105 shows the locations of the optional I/O Books in the x3850 X6 server.

As shown in Figure 3-25 on page 102, the optional I/O Books require the following processors (Compute Books) to be installed:

- ▶ The I/O Book in bay 1 requires processor 4 be installed

- ▶ The I/O Book in bay 2 requires processor 3 be installed




Figure 3-28 Optional I/O Book locations

The part numbers for the optional I/O Books are listed in Table 3-17.

The table also lists the maximum number of books that is supported. The x3850 X6 supports up to *two* optional I/O Books, and they both can be Half-length I/O Books or both Full-length I/O Books or one of each. Similarly, the x3950 X6 supports up to *four* Optional I/O Books in any combination.

Table 3-17 Optional I/O Book part numbers

Part number	Feature code	Description	Maximum supported x3850 X6	Maximum supported x3950 X6
44X4049	A4A2	X6 Half-Length I/O Book	2	4
44X4051	A4A3 <sup>a</sup>	X6 Full-Length I/O Book		

a. The Full-length I/O Book can be ordered as an option only. It is not available via CTO because the Full-length I/O Book cannot be shipped as installed in the server.

Figure 3-29 shows the two Optional I/O Books.




Figure 3-29 Optional I/O Books

As described in 2.2, “System architecture” on page 20, the first optional I/O Book connects to the processor in Compute Book 3 and the second optional I/O Book connects to the processor in Compute Book 4. This bus connectivity results in the following requirements:

- ▶ In the x3850 X6, four Compute Books are required to use one or two optional I/O Books.
- ▶ In the x3950 X6, the use of two optional I/O Books requires six Compute Books, and the use of four optional I/O Books requires all eight Compute Books.

### 3.14.1 Half-length I/O Book

The Half-length I/O Book provides the following PCIe slots to the server:

- ▶ Two PCIe 3.0 x8 slots (x8 wired)
- ▶ One PCIe 3.0 x16 slot (x16 wired)

All slots support half-length full-height adapters (full-length adapters are not supported) and the maximum power consumption for each PCIe slot is 75 W. Figure 3-30 shows a top-down view of the Half-length I/O Book.




Figure 3-30 Half-length I/O Book PCIe ports location

The Half-length I/O Book supports hot-swap PCIe adapters. For more information, see 3.15, “Hot-swap adapter support” on page 108.

### 3.14.2 Full-length I/O Book

The Full-length I/O Book provides the following PCIe slots to the server:

- ▶ Two PCIe 3.0 x16 (x16 wired)
- ▶ One PCIe 2.0 x8 slot (x4 wired)

The Full-length I/O Book also includes two auxiliary power connectors. With the use of these connectors and the supplied power cords, the I/O Book supports one double-wide adapter up to 300 W. The following auxiliary power connectors are available:

- ▶ One 2x4 power connector; supplies up to 150 W more power to the adapter
- ▶ One 2x3 power connector; supplies up to 75 W more power to the adapter

The cables that connect to these auxiliary connectors are shipped with the Full-length I/O Book.

The combined power consumption of all adapters that are installed in the Full-length I/O Book cannot exceed 300 W.

**Note:** The 2x3 connector is intended to be used only when one adapter is installed in the first x16 slot (the up-most slot that is shown in Figure 3-31), which requires 225 W or 300 W of power. The location of the 2x3 connector prevents an adapter from being installed in the other x16 slot.

Figure 3-31 shows the Full-length I/O Book.




Figure 3-31 Full-length I/O Book

Figure 3-32 shows the Full-length I/O Book with Intel Xeon Phi coprocessor installed.



Figure 3-32 Full-length I/O Book with Intel Xeon Phi card installed

The Half-length I/O Book installs flush with the Primary I/O Book at the rear of the server. When installed, the Full-length I/O Book adds a 99 mm (4 in) mechanical extension to the base length dimension of the chassis.

**Rack support:** Some rack chassis are not deep enough to support the servers with the Full-length I/O Book installed. For more information, see 4.2, “Rack selection and rack options” on page 131.

Figure 3-33 shows a Full-length I/O Book and a Half-length I/O Book installed in the server.




Figure 3-33 I/O Books installed in the x3850 X6

The Full-length I/O Book supports hot-swap PCIe adapters. For more information, see 3.15, “Hot-swap adapter support” on page 108.

## 3.15 Hot-swap adapter support

The Half-length I/O Book and Full-length I/O Book are hot-swap capable. Because of this capability, you can remove the I/O Book while the operating system is still running, add or replace an adapter, then reinsert the I/O Book.

The following requirements must be met for hot-swap:

- ▶ The operating system must support hot-swap. At the time of this writing, the following operating systems support hot-swap adapters:
  - Windows Server 2008 R2
  - Windows Server 2012
  - Windows Server 2012 R2
  - Red Hat Enterprise Linux 6 Server x64 Edition
  - Red Hat Enterprise Linux 7
  - SUSE Linux Enterprise Server 11 for AMD64/EM64T
  - SUSE Linux Enterprise Server 11 with Xen for AMD64/EM64T
  - SUSE Linux Enterprise Server 12
  - SUSE Linux Enterprise Server 12 with XEN
- ▶ For an I/O Book to be hot-swappable, all adapters that are installed in the book must support hot-swap.
- ▶ Removing a filler blank and replacing it with a Half-length I/O Book is supported; however, replacing a filler blank with a Full-length I/O Book while the server is running is not supported. A restart is required.

- ▶ A Half-length I/O Book cannot be hot-swapped with a Full-length I/O Book and a Full-length I/O Book cannot be hot-swapped with a Half-length I/O Book. A restart is required.
- ▶ Only certain adapters support hot-swap. Table 3-18 on page 109 lists the adapters that supported hot-swap at the time of this writing.

Table 3-18 Adapters that support hot-swap

Part number	Feature code	Description
00D8540	A4M9	Emulex Dual Port 10GbE SFP+ VFA IIIr for System x <sup>a</sup>
49Y7960	A2EC	Intel X520 Dual Port 10GbE SFP+ Adapter for System x
49Y7970	A2ED	Intel X540-T2 Dual Port 10GBaseT Adapter for System x
49Y4230	5767	Intel Ethernet Dual Port Server Adapter I340-T2 for System x
49Y4240	5768	Intel Ethernet Quad Port Server Adapter I340-T4 for System x
00AG500	A56K	Intel I350-F1 1xGbE Fibre Adapter for System x
00AG510	A56L	Intel I350-T2 2xGbE BaseT Adapter for System x
00AG520	A56M	Intel I350-T4 4xGbE BaseT Adapter for System x

a. Hot-swap is supported only when the adapter is in pNIC mode. Hot-swap is not supported in either vNIC™ mode (Virtual Fabric mode or Switch Independent mode).

## 3.16 Network adapters

The x3850 X6 and x3950 X6 servers support ML2 adapters that are installed in the custom ML2 slot. This slot supports adapters with two 10 Gb ports or four ports that are 1 Gb. It also supports direct connectivity to the IMM2 service processor for out-of-band systems management.

As described in 3.3, “Standard models of X6 servers” on page 60, Models B3x, F3x, and F4x include the Broadcom NetXtreme II ML2 Dual Port 10GbE SFP+ adapter as standard. All other standard models include Intel I350-T4 ML2 Quad Port GbE Adapter adapter (I350-AM4 based).

The Broadcom NetXtreme II ML2 Dual Port 10GbE SFP+ Adapter has the following specifications:

- ▶ Dual-port 10 Gb Ethernet connectivity
- ▶ Broadcom BCM57810S ASIC
- ▶ SFP+ ports that support fiber optic and direct-attach copper (DAC) cables

For more information about this adapter, see *Broadcom NetXtreme 10 GbE SFP+ Network Adapter Family for System x*, TIPS1027, which is available at this website:

<http://lenovopress.com/tips1027>

The Intel I350-T4 ML2 Quad Port GbE Adapter has the following specifications:

- ▶ Quad-port 1 Gb Ethernet connectivity
- ▶ Intel I350-AM4 ASIC
- ▶ RJ45 ports for copper cables

For more information about this adapter, see *Intel I350 Gigabit Ethernet Adapter Family for System x*, TIPS1155, which is available at this website:

<http://lenovopress.com/tips1155>

The supported ML2 adapters are listed in Table 3-19 on page 110.

Table 3-19 ML2 adapter options

Part number	Feature code	Description	Maximum supported x3850 X5	Maximum supported x3950 X5
<b>25 Gb Ethernet</b>				
00MN990 <sup>a</sup>	ATZR	Mellanox ConnectX-4 Lx ML2 1x25GbE SFP28 Adapter	1	2
<b>10 Gb Ethernet</b>				
00D2026	A40S	Broadcom NetXtreme II ML2 Dual Port 10GbaseT	1	2
00D2028	A40T	Broadcom NetXtreme II ML2 Dual Port 10GbE SFP+	1	2
00AG560	AT7U	Emulex VFA5.2 ML2 Dual Port 10GbE SFP+ Adapter	1	2
00D8544	A4NZ	Emulex VFA5 ML2 FCoE/iSCSI license (FoD) (Upgrade for 00D1996 - one for each adapter)	1	2
00D1994	A40P	Intel X540 ML2 Dual Port 10GbaseT Adapter	1	2
00JY940 <sup>a</sup>	ATRH	Intel X710-DA2 ML2 2x10GbE SFP+ Adapter	1	2
94Y5200	AS74	Intel X710 ML2 4x10GbE SFP+ Adapter for System x	1	2
<b>Gigabit Ethernet</b>				
00D1998	A40R	Intel I350-T4 ML2 Quad Port GbE Adapter	1	2
<b>InfiniBand</b>				
00FP650	A5RK	Mellanox ConnectX-3 Pro ML2 2x40GbE/FDR VPI Adapter for System x	1	2

a. Not supported in servers with E7 v2 compute books

The server also supports various other Ethernet and InfiniBand network adapters, as listed in Table 3-20. The maximum quantity that is listed is for configurations with all processors and I/O books installed.

Table 3-20 Network adapters

Part number	Feature code	Description	Hot-swap capable	Maximum supported <sup>a</sup>
<b>100 Gb Ethernet</b>				
00MM960 <sup>b</sup>	ATRP	Mellanox ConnectX-4 2x100GbE/EDR IB QSFP28 VPI Adapter	No	4 / 8
<b>40 Gb Ethernet</b>				
00MM950 <sup>b</sup>	ATRN	Mellanox ConnectX-4 Lx 1x40GbE QSFP+ Adapter	No	9 / 18
00D9550	A3PN	Mellanox ConnectX-3 40GbE/ FDR IB VPI Adapter for System x	No	9 / 18
<b>25 Gb Ethernet</b>				
01GR250	AUAJ	Mellanox ConnectX-4 Lx 2x25GbE SFP28 Adapter	No	9 / 18



Part number	Feature code	Description	Hot-swap capable	Maximum supported <sup>a</sup>
<b>10 Gb Ethernet</b>				
94Y5180	A4Z6	Broadcom NetXtreme Dual Port 10GbE SFP+ Adapter	No	9 / 18
44T1370	A5GZ	Broadcom NetXtreme 2x10GbE BaseT Adapter for System x	No	9 / 18
00AG570	AT7S	Emulex VFA5.2 2x10 GbE SFP+ PCIe Adapter	No	9 / 18
00AG580	AT7T	Emulex VFA5.2 2x10 GbE SFP+ Adapter and FCoE/iSCSI SW	No	9 / 18
00JY820	A5UT	Emulex VFA5 2x10 GbE SFP+ PCIe Adapter for System x	No	9 / 18
00JY830 <sup>c</sup>	A5UU	Emulex VFA5 2x10 GbE SFP+ Adapter and FCoE/iSCSI SW for System x	No	9 / 18
00JY824	A5UV	Emulex VFA5 FCoE/iSCSI SW for PCIe Adapter for System x (FoD) (upgrade for 00AG570 or 00JY820)	License	9 / 18
00D8540 <sup>c</sup>	A4M9	Emulex Dual Port 10GbE SFP+ VFA III-R for System x	Yes	9 / 18
49Y7960	A2EC	Intel X520 Dual Port 10GbE SFP+ Adapter for System x	Yes	9 / 18
00MM850 <sup>d</sup>	ATRY	Intel X550-T1 Single Port 10GBase-T Adapter	No	9 / 18
49Y7970	A2ED	Intel X540-T2 Dual Port 10GBaseT Adapter for System x	Yes	9 / 18
00MM860 <sup>b</sup>	ATPX	Intel X550-T2 Dual Port 10GBase-T Adapter	No	9 / 18
81Y3520 <sup>c</sup>	AS73	Intel X710 2x10GbE SFP+ Adapter for System x	No	9 / 18
00D9690	A3PM	Mellanox ConnectX-3 10 GbE Adapter for System x	No	9 / 18
90Y4600 <sup>c</sup>	A3MR	QLogic 8200 Dual Port 10GbE SFP+ VFA for System x	No	9 / 18
00Y5624 <sup>c</sup>	A3MT	QLogic 8200 VFA FCoE/iSCSI license for System x (FoD) for System x (upgrade for 90Y4600)	License	9 / 18
<b>Gigabit Ethernet</b>				
90Y9370	A2V4	Broadcom NetXtreme I Dual Port GbE Adapter for System x	No	9 / 18
90Y9352	A2V3	Broadcom NetXtreme I Quad Port GbE Adapter for System x	No	9 / 18
49Y4230 <sup>c</sup>	5767	Intel Ethernet Dual Port Server Adapter I340-T2 for System x	Yes	9 / 18
49Y4240 <sup>c</sup>	5768	Intel Ethernet Quad Port Server Adapter I340-T4 for System x	Yes	9 / 18
00AG500	A56K	Intel I350-F1 1xGbE Fibre Adapter for System x	Yes	9 / 18
00AG510	A56L	Intel I350-T2 2xGbE BaseT Adapter for System x	Yes	9 / 18
00AG520	A56M	Intel I350-T4 4xGbE BaseT Adapter for System x	Yes	9 / 18
<b>InfiniBand</b>				
00D9550	A3PN	Mellanox ConnectX-3 40GbE/ FDR IB VPI Adapter for System x	No	9 / 18
00MM960 <sup>b</sup>	ATRP	Mellanox ConnectX-4 2x100GbE/EDR IB QSFP28 VPI Adapter	No	4 / 8
00KH924	ASWQ	Mellanox ConnectX-4 EDR IB VPI Single-port x16 PCIe 3.0 HCA	No	4 / 8
<b>Omni-Path Architecture (OPA)</b>				
00WE027	AU0B	Intel OPA 100 Series Single-port PCIe 3.0 x16 HFA	No	3 / 6

Part number	Feature code	Description	Hot-swap capable	Maximum supported <sup>a</sup>
00WE023	AU0A	Intel OPA 100 Series Single-port PCIe 3.0 x8 HFA	No	2 / 4

- a. Quantities for x3850 X6 / x3950 X6
- b. Not supported in servers with E7 v2 compute books
- c. Withdrawn from marketing
- d. Only supported in servers with E7 v4 compute books

For more information, see the list of Lenovo Press Product Guides in the Networking adapters category that is available at this website:

<http://lenovopress.com/systemx/networkadapters>

## 3.17 SAS adapters for external storage

Table 3-21 on page 112 lists the SAS HBAs and external RAID adapters that are supported by the x3850 X6 and x3950 X6. The maximum quantity listed is for configurations with all processors and I/O books installed.

Table 3-21 lists the supported RAID controller and HBA for external storage connectivity.

Table 3-21 SAS HBAs, RAID controllers and options for external disk storage expansion

Part number	Feature code	Description	Maximum supported x3850 X6	Maximum supported x3950 X6
<b>SAS HBAs</b>				
46C9010 <sup>a</sup>	A3MV	N2125 SAS/SATA HBA	9	18
00AE912	A5M0	N2225 SAS/SATA HBA	9	18
00AE916	A5M1	N2226 SAS/SATA HBA	5	10
<b>RAID adapters</b>				
00AE938	A5ND	ServeRAID M5225-2GB SAS/SATA Controller	3	6
81Y4478 <sup>a</sup>	A1WX	ServeRAID M5120 SAS/SATA Controller	2	5
<b>Features on-Demand upgrades for the M5225</b>				
47C8706	A3Z5	ServeRAID M5200 Series RAID 6 Upgrade FoD	1	1
47C8710	A3Z7	ServeRAID M5200 Series Performance Accelerator FoD	1	1
47C8712	A3Z8	ServeRAID M5200 Series SSD Caching Enabler FoD	1	1
<b>Supported hardware upgrades for the M5120</b>				
81Y4487 <sup>a</sup>	A1J4	ServeRAID M5100 Series 512MB Flash/RAID 5 Upgrade	1	2
81Y4559 <sup>a</sup>	A1WY	ServeRAID M5100 Series 1GB Flash/RAID 5 Upgrade	1	2
47C8670 <sup>a</sup>	A4G6	ServeRAID M5100 Series 2GB Flash/RAID 5 Upgrade	1	2
<b>Supported Features on-Demand upgrades for the M5120</b>				
90Y4318 <sup>a</sup>	A2MD	ServeRAID M5100 Series SSD Caching Enabler	1 <sup>b</sup>	2 <sup>b</sup>

Part number	Feature code	Description	Maximum supported x3850 X6	Maximum supported x3950 X6
90Y4273 <sup>a</sup>	A2MC	ServeRAID M5100 Series SSD Performance Key	1 <sup>b</sup>	2 <sup>b</sup>
81Y4546 <sup>a</sup>	A1X3	ServeRAID M5100 Series RAID 6 Upgrade	1 <sup>b</sup>	2 <sup>b</sup>

a. Withdrawn from marketing

b. These M5120 features upgrades require a cache memory upgrade (part number 81Y4487 or 81Y4559).

Table 3-22 compares the specifications of the external SAS HBAs and RAID adapters.

Table 3-22 Specifications comparison

Feature	N2125	N2225	N2226	ServeRAID M5120	ServeRAID M5225
Adapter type	SAS HBA	SAS HBA	SAS HBA	RAID adapter	RAID adapter
Part number	46C9010	00AE912	00AE916	81Y4478	00AE938
Form factor	Low profile	Low profile	Low profile	Low profile	Low profile
Controller chip	LSI SAS2308	LSI SAS3008	LSI SAS3008	LSI SAS2208	LSI SAS3108
Host interface	6 Gbps SAS	12 Gbps SAS	12 Gbps SAS	6 Gbps SAS	12 Gbps SAS
Number of external ports	8	8	8	8	8
External port connectors	2x Mini-SAS (SFF-8088)	2x Mini-SAS HD (SFF-8644)	4x Mini-SAS HD (SFF-8644)	2x Mini-SAS (SFF-8088)	2x Mini-SAS HD (SFF-8644)
Drive interface	SAS, SATA	SAS, SATA	SAS, SATA	SAS, SATA	SAS, SATA
Drive type	HDD, SSD	HDD, SSD	HDD, SSD	HDD, SED, SSD	HDD, SED, SSD
Maximum number of devices	1024	1024	1024	240	240
RAID levels	None	None	None	0/1/10; Optional 5/50 and 6/60	0/1/10/5/50; Optional 6/60 (47C8706)
JBOD mode	Yes	Yes	Yes	No	No
Cache	None	None	None	Optional; 512MB, 1GB or 2GB	2 GB (included)
CacheVault cache protection	None	None	None	Flash (with cache upgrades)	Flash (included)
FastPath	None	None	None	Optional (90Y4273)	Optional (47C8710)
CacheCade Pro 2.0	None	None	None	Optional (90Y4318)	Optional (47C8712)

**Cache upgrade required:** The ServeRAID M5120 SAS/SATA Controller ships standard without a cache. One of the available cache upgrades (part number 81Y4487 or 81Y4559) is required for the M5120 adapter operations, and it must be purchased with the controller.

For more information about the adapters, see these Lenovo Press Product Guides:

- ▶ N2125: <http://lenovopress.com/tips1062>
- ▶ N2225 & N2226: <http://lenovopress.com/tips1175>
- ▶ ServeRAID M5120: <http://lenovopress.com/tips0858>
- ▶ ServeRAID M5225: <http://lenovopress.com/tips1258>

### 3.18 Fibre Channel host bus adapters

Table 3-23 lists Fibre Channel HBAs that are supported by the x3850 X6 and x3950 X6 servers. The maximum quantity listed is for configurations with all processors and I/O Books installed.

Table 3-23 Fibre Channel HBA options

Part number	Feature code	Description	Hot-swap capable	Maximum supported x3850 X6	Maximum supported x3950 X6
<b>Fibre Channel: 16 Gb (New generation)</b>					
01CV830 <sup>a</sup>	ATZU	Emulex 16Gb FC Single-port HBA	No	9	18
01CV840 <sup>a</sup>	ATZV	Emulex 16Gb FC Dual-port HBA	No	9	18
01CV750 <sup>a</sup>	ATZB	QLogic 16Gb FC Single-port HBA	No	9	18
01CV760 <sup>a</sup>	ATZC	QLogic 16Gb FC Dual-port HBA	No	9	18
<b>Fibre Channel: 16 Gb</b>					
81Y1655	A2W5	Emulex 16Gb FC Single-port HBA	No	9	18
81Y1662	A2W6	Emulex 16Gb FC Dual-port HBA	No	9	18
00Y3337	A3KW	QLogic 16Gb FC Single-port HBA	No	9	18
00Y3341	A3KX	QLogic 16Gb FC Dual-port HBA	No	9	18
<b>Fibre Channel: 8 Gb</b>					
42D0485	3580	Emulex 8 Gb FC Single-port HBA	No	9	18
42D0494	3581	Emulex 8 Gb FC Dual-port HBA	No	9	18
42D0501	3578	QLogic 8 Gb FC Single-port HBA	No	9	18
42D0510	3579	QLogic 8 Gb FC Dual-port HBA	No	9	18

a. Not supported in servers with E7 v2 compute books

### 3.19 GPU adapters and co-processors

The server supports co-processors and graphics processing units (GPUs) that are listed in Table 3-24; each is installed in a Full-length I/O Book. No other adapter can be installed in the Full-length I/O Book. The Full-length I/O Book includes the necessary auxiliary power cables.

The GPUs and coprocessors have server memory minimums and maximums, as indicated in the Table 3-24.

**Not available via CTO:** These adapters are not available via CTO and cannot be shipped installed in the server. The adapters cannot be shipped installed because they are installed in the Full-length I/O Book, which extends beyond the rear of the chassis (see Figure 3-33 on page 108). These adapters must be shipped separately from the server.

Table 3-24 GPU adapters and coprocessors

Part number	Feature code	Description	Min / Max system RAM	Hot-swap capable	Maximum supported x3850 X6	Maximum supported x3950 X6
00FP676 <sup>a</sup>	A5FG <sup>b</sup>	NVIDIA Tesla K40c PCIe x16 for System x3850/x3950 X6	24 GB / 1 TB	No	2	4
00FP675 <sup>a</sup>	A3YU <sup>b</sup>	NVIDIA Quadro K4000 PCIe x16 for System x3850/x3950 X6	6 GB / 1 TB	No	2	4
00FP672 <sup>a</sup>	A3YV <sup>b</sup>	NVIDIA Quadro K6000 PCIe x16 for System x3850/x3950 X6	24 GB / 1 TB	No	2	4
90Y2495	AU3W <sup>b</sup>	NVIDIA Quadro M6000 24GB GPU, PCIe (active)	24 GB	No	2	4
00YL378	ATZF	NVIDIA Quadro M5000 GPU, PCIe (active)	16 GB / 1 TB	No	2	4

- a. Withdrawn from marketing
- b. Not available via CTO.

**Note:** If an NVIDIA GPU is installed, a maximum of 1 TB of system memory can be installed.

## 3.20 Partitioning

The 8-socket x3950 X6 server can form a 2-node configuration by using a feature called *node partitioning*. The x3950 X6 server electronically splits into two logical independent servers or nodes, which use their own set of resources and have their own operating system installed.

When partitioning is enabled, the 8-socket server is seen by the operating systems as two independent 4-socket servers, as shown in Figure 3-34.




Figure 3-34 x3950 X6 server with partitioning

For details about how to implement partitioning, see 5.7, “Partitioning the x3950 X6” on page 205.

## 3.21 Standard onboard features

In this section, we describe standard onboard features in the X6 servers, such as Integrated Management Module II, UEFI, Integrated TPM, and light path diagnostics.

This section includes the following topics:

- ▶ 3.21.1, “Integrated Management Module II” on page 117
- ▶ 3.21.2, “Unified Extensible Firmware Interface” on page 117
- ▶ 3.21.3, “Integrated Trusted Platform Module” on page 117
- ▶ 3.21.4, “Light path diagnostics” on page 118

### 3.21.1 Integrated Management Module II

Each X6 server has an IMM2 (version 2.1) service processor onboard. The IMM2 provides the following standard major features:

- ▶ IPMI v2.0 compliance
- ▶ Remote configuration of IMM2 and UEFI settings without the need to power on the server
- ▶ Remote access to 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, by using Simple Network Management Protocol, SNMP)
- ▶ Remote presence, including remote control of server by using a Java or Active x client
- ▶ Operating system failure window (blue screen) capture and display through the web interface
- ▶ Virtual media that allow 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 2.8, “Integrated Management Module” on page 49.

For more information about configuring the IMM2, see 5.1, “Configuring the IMM2 settings” on page 156.

### 3.21.2 Unified Extensible Firmware Interface

The x3850 X6 and x3950 X6 servers use an integrated UEFI next-generation BIOS, which includes the following capabilities:

- ▶ Human-readable event logs; no more beep codes
- ▶ Complete setup solution by allowing adapter configuration function to be moved to UEFI
- ▶ Complete out-of-band coverage by Advanced Settings Utility™ to simplify remote setup

The use of all of the features of UEFI requires an UEFI-aware operating system and adapters. UEFI is fully compatible with an earlier version of BIOS.

For more information about UEFI, see 2.7, “Unified Extensible Firmware Interface” on page 49 and the Lenovo white paper, *Introducing UEFI-Compliant Firmware on System x and BladeCenter Servers*, which is available at this address:

<https://support.lenovo.com/docs/UM103225>

For more information about the UEFI menu setup, see 5.2, “Unified Extensible Firmware Interface settings” on page 165 and 5.3, “UEFI common settings” on page 169.

### 3.21.3 Integrated Trusted Platform Module

Trusted computing is an industry initiative that provides a combination of secure software and secure hardware to create a trusted platform. It is a specification that increases network security by building unique hardware IDs into computing devices. The X6 servers implement TPM version 1.2 support.

The x3850 X6 has two TPM chips in the Primary I/O Book; the x3950 X6 has four TPM chips.

Full disk encryption applications, such as the BitLocker Drive Encryption feature of Microsoft Windows Server 2008, can use this technology. The operating system uses it to protect the keys that encrypt the computer's operating system volume and provide integrity authentication for a trusted boot pathway (BIOS, boot sector, and others). Several vendor full-disk encryption products also support the TPM chip.

For more information about this technology, see the Trusted Computing Group (TCG) TPM Main Specification at this website:

[http://www.trustedcomputinggroup.org/resources/tpm\\_main\\_specification](http://www.trustedcomputinggroup.org/resources/tpm_main_specification)

### 3.21.4 Light path diagnostics

Light path diagnostics is a set of LEDs on server's components that help to determine failed components. If a server's component failed or has an error, appropriate LEDs light up on the front operator panel, the light path diagnostics panel, and on the failed component. By checking the LEDs, you can determine the failed component.

If an error occurs, check the LEDs on the front operator panel first. If the check log LED is lit, check IMM event log or system-event log. If the system-error LED is lit, find the appropriate component with lit LED. Figure 3-35 shows the location of LEDs on the front operator panel.




Figure 3-35 Front operator panel LEDs location



For example, if a Compute Book error occurs, the Error LED on the appropriate Compute Book should light up. Figure 3-36 shows the LED locations on the Compute Book.




Figure 3-36 Compute Book LEDs location

If the Compute Book has a memory error, appropriate LEDs on the Compute Book’s system board can help to determine the failed DIMM. To use this feature, you must disconnect the server from power, remove the Compute Book from the server, press the Light path button on the front panel of the Compute Book, and find the lit LED on the system board. The DIMM LED that is lit indicates the failed memory module. Figure 3-37 shows the location of the DIMM LEDs on one side of the Compute Book.




Figure 3-37 DIMMs connectors location

In addition to the light path diagnostic LEDs, you can use the LCD system information display on the front operator panel, which displays errors from the IMM event log.

For more information about the LCD system information display and the front operator panel, see 3.6.2, “LCD system information panel” on page 72.

For more information about Light path diagnostics, see the *System x3850 X6 Installation and Service Guide*, which is available at this website:

[http://publib.boulder.ibm.com/infocenter/systemx/documentation/index.jsp?topic=/com.lenovo.sysx.6241.doc/product\\_page.html](http://publib.boulder.ibm.com/infocenter/systemx/documentation/index.jsp?topic=/com.lenovo.sysx.6241.doc/product_page.html)

## 3.22 Integrated virtualization

The X6 servers support VMware ESXi that is installed on a USB memory key. The key is installed in a USB port that is on the Primary I/O Book. Figure 3-38 shows the location of the USB port for the embedded hypervisor on the Primary I/O Book system board.




Figure 3-38 Internal USB port location for embedded hypervisor (view from above)

The virtualization options are listed in Table 3-25.

Table 3-25 Virtualization options

Part number	Feature code	Description	Maximum supported <sup>a</sup>
Blank USB keys			
00WH140	ATRM	Blank USB Memory Key 4G SLC for VMware ESXi Downloads (4GB capacity)	1 / 2
41Y8298	A2G0	Blank USB Memory Key for VMware ESXi Downloads (2GB capacity)	1 / 2
USB keys preloaded with Lenovo custom image			
00ML235	ASN7	USB Memory Key for VMware ESXi 5.5 Update 2	1 / 2
00WH150 <sup>b</sup>	ATZG	USB Memory Key for VMware ESXi 5.5 Update 3B	1 / 2
00WH138	ATRL	USB Memory Key 4G for VMware ESXi 6.0 Update 1A	1 / 2
00WH151 <sup>b</sup>	ATZH	USB Memory Key for VMware ESXi 6.0 Update 2	1 / 2
CTO only	AVNW	USB Memory Key for VMware ESXi 6.5	1 / 2

- a. Two hypervisor keys are supported by only the x3950 X6 and only if the x3950 X6 is configured to be partitioned, where the two halves of the server operate as two independent four-socket servers. CTO orders can include only one hypervisor key.
- b. Not supported in servers with E7 v2 compute books

For more information about the use of the embedded hypervisor, see 6.4, “Use of embedded VMware ESXi” on page 239.

### 3.23 Hot-swap capabilities

The X6 servers follow the standard color coding scheme that is used by Lenovo for touch points and hot-swap components.

Hot-swap components have orange handles or touch points. Orange tabs are found on fan modules, power supplies, extra I/O Books, and disk drives. The orange designates that the items are hot-swap, and can be removed and replaced while the chassis is powered.

Touch points that are blue cannot be hot-swapped; the server must be powered off before removing these devices. Blue touch points can be found on components, such as Compute Books, Primary I/O Book, and Storage Book.

Table 3-26 lists which components are hot-swap.

*Table 3-26 Hot-swap capabilities*

Component	Hot-swap ability
Compute Book	No
Storage Book	No
Primary I/O Book	No
Optional I/O Books	Yes
Disk drives	Yes
Compute Book fans	Yes
Primary I/O Book fans	Yes
Power supplies	Yes

For more information about the hot-swap capabilities of the Optional I/O Books, see 3.15, “Hot-swap adapter support” on page 108.

## 3.24 Power subsystem

The x3850 X6 server supports up to four power supplies and the x3950 X6 supports up to eight power supplies. Power supplies are hot-swap. The available options are listed in Table 3-27.

Table 3-27 Power supply options

Part number	Feature code	Description	Supported quantities x3850 X6	Supported quantities x3950 X6	Models where used
44X4132	A4R0	900W Power Supply	2, 4 <sup>a</sup>	4, 8	A4x, B1x, B3x, C1x, C4x, BAx, CAx, CCx
44X4152	A54E	1400W HE Redundant Power Supply (not certified for altitudes above 5000 m)	2, 4	4, 8	-
44X4150	A54D	1400W HE Redundant Power Supply for altitudes above 5000 meters	2, 4	4, 8	All HANA models
88Y7433	A2EA	750W High Efficiency -48 V DC Power Supply	4	8	-

a. The only exception to the rule of a minimum of two power supplies is model A4x, which has a minimum of one power supply.

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

Consider the following configuration rules:

- ▶ Power supplies must be in pairs (x3850 X6) or quads (x3950 X6), except for model A4x, where one power supply is supported.
- ▶ The AC power supplies can be mixed. For example, the following combinations are valid for the x3850 X6:
  - Two 44X4132 and two 44X4152
  - Two 44X4132 and two 44X4150
- ▶ It is not advised that you mix the two 1400 W options (44X4152 for low altitude and 44X4150 for high altitude) because you will nullify the high-altitude capabilities of 44X4150.
- ▶ You cannot mix the AC power supplies with the DC power supply.

Use the Power Configurator that is available at the following website to determine the power that your server needs:

<http://support.lenovo.com/documents/LNVO-PWRCONF>

For more information about power supplies and power planning, see 4.5, “Power guidelines” on page 136.

For additional information on powering and cooling the X6 system refer to the power and cooling reference guides located here:

<http://support.lenovo.com/documents/LNVO-POWINF>

For assistance with power and cooling for the X6 server, email [power@lenovo.com](mailto:power@lenovo.com).

Figure 3-39 shows the 1400 W AC power supply rear view and highlights the LEDs. There is a handle for removal and insertion of the power supply.




Figure 3-39 1400W AC power supply

The following rear LEDs are available:

- ▶ AC Power: When lit green, this LED indicates that AC power is being supplied to the PSU inlet.
- ▶ DC Power: When lit green, this LED indicates that DC power is being supplied to the chassis midplane.
- ▶ Fault: When lit amber, this LED indicates a fault with the PSU.

Total output power is only 95% additive because of loss from current sharing (for example, two 1400 W supplies provide 2660 W instead of 2800 W).

The power supply supports EPA Platinum, and 80Plus certification. The standard has several ratings, such as Bronze, Silver, Gold, and Platinum. To meet the 80 PLUS Platinum standard, the power supply must have a power factor (PF) of 0.95 or greater at 50% rated load and efficiency equal to or greater than the following values:

- ▶ 90% at 20% of rated load
- ▶ 94% at 50% of rated load
- ▶ 91% at 100% of rated load

For more information about 80 PLUS, see this website:

<http://www.plugloadsolutions.com>

The four power supply bays are divided into two power domains to support N+N power supply and power feed redundancy, where N = 1 or 2 (depending upon system configuration and load). Power supplies that are installed in bays 1 and 3 belong to Group A; power supplies that are installed in bays 2 and 4 belong to Group B.

The x3850 X6 server supports the following modes of redundancy based on the power supply configuration, system load, and the Power Policy configuration controlled by the IMM:

- ▶ Non-redundant
- ▶ Fully system redundant
- ▶ Redundant with reduced performance (throttling)

The IMM2 must be used to set and change the Power Policy and System Power configurations. The power configurations and policies can be changed via the web interface, CIM, and ASU interfaces. These settings cannot be changed by UEFI. The default configuration setting for AC and DC models is Non-Redundant with Throttling enabled.

For more information about power supply redundancy and power supply installation order, see 4.5, “Power guidelines” on page 136.

## 3.25 Fans and cooling

The x3850 X6 servers feature the following types of fans:

- ▶ Up to 10 hot-swap, counter-rotating system fans and two fans (motors) per fan pack that are used for Compute Books and Storage Book cooling. Two fan packs are used for each Compute Book cooling and are installed only when a corresponding Compute Book is installed. Two HDD fan packs in Primary I/O Book must always be installed for Storage Book cooling.
- ▶ Power supply 1 - 4 internal fans at the rear of the power supply.

Figure 3-40 shows fan locations and numbering in the server.




Figure 3-40 Fan locations and numbering

Fan packs 1 - 8 are in front of the Compute Books that are numbered left to right in Figure 3-40, looking at the front of the server. Fan packs 9 and 10 are in the rear of the server, numbered left to right in Figure 3-40.

The fan zones are available:

- ▶ Fan Zone 1: Fan packs 9 and 10, which are used to cool the Storage Book.
- ▶ Fan Zone 2: Fan packs 1 and 2, which are controlled by the temperature readings near Compute Book 1, ML2 slot, and PCIe slot 9.
- ▶ Fan Zone 3: Fan packs 3 and 4, which are controlled by the temperature readings near Compute Book 2 and PCIe slots 7 and 8.
- ▶ Fan Zone 4: Fan packs 5 and 6, which are controlled by the temperature readings near Compute Book 3 and PCIe slots 4, 5, and 6.
- ▶ Fan Zone 5: Fan packs 7 and 8, which are controlled by the temperature readings near Compute Book 4 and PCI slots 1, 2, and 3.
- ▶ Fan Zone 6: Power Supply 1 - 4 fans.

Fan speed for all fan zones is controlled by several parameters, such as inlet ambient temperature, CPU temperature, DIMMs, and PCIe card temperatures.

Figure 3-41 shows one of the Compute Book fans and Primary I/O Book fan (for Storage Book cooling).



Figure 3-41 Compute Book and Primary I/O Book fans

## 3.26 Upgrading to an 8-socket X6 server

The x3850 X6 server has a flexible modular design with which you can increase the server's compute power and I/O capabilities by adding Compute Books and I/O Books. The modular design also means that if your business needs more processing or I/O capability within the same system image, you can migrate to an eight-socket x3950 X6.

Lenovo supports upgrading a 4-socket X6 server to an 8-socket X6 server. The following methods are recommended:

- ▶ Start with a 4-socket x3950 X6 (8U chassis) and add processors when needed.
- ▶ Start with a 4-socket x3850 X6 (4U chassis) and upgrade it when needed by using a request for price quotation (RPQ) process.

### Option 1: Start with a 4-socket x3950 X6

With this option, you plan for the need for future 8-socket performance and purchase an x3950 X6 with only four Compute Books installed. The initial order includes the following components:

- ▶ Four Compute Books with E7-8800 with the same processor models (v2, v3, or v4). E7-4800 series processors are not supported.
- ▶ Four Power Supplies (preferably 1400W or combination of 900W and 1400W)
- ▶ Two Primary IO Book standard
- ▶ Two Storage Book Standard

When the need arises to upgrade the server to six or eight processors, purchase more Compute Books with the same processor model and more power supplies (as determined by using the Power Configurator), I/O Books, adapters, and drives as needed.

**Note:** Compute Books can be installed in a x3950 X6 from a System x3850 X6 (machine type 3837) or a System x3850 X6 (machine type 6241), as long as the same Intel Xeon E7-8800 series processors are used in those Compute Books as in the x3950 X6's Compute Books.

## Option 2: Upgrade your 4-socket x3850 X6 by using an RPQ

With this option, you do not have to plan for 8-socket capacity up front. You start with an x3850 X6 and when you are ready to upgrade to a 6-socket or 8-socket server, you purchase an upgrade offering via an RPQ.

As part of the RPQ, a service engineer comes onsite with the new mechanical chassis and performs the field upgrade by transferring all components to the new chassis. This method also requires the x3850 X6 compute books be the same E7-8800 processors as ordered for the RPQ; however, in this scenario, the server maintain the original serial number.

**Use of E7-4800 processors:** Intel Xeon E7-4800 v2, v3, and v4 processors cannot be used in an x3950 X6. If your x3850 X6 has Compute Books with E7-4800 processors, these components must be replaced with Compute Books with E7-8800 v2, v3, or v4 processors if you plan to upgrade to an x3950 X6. In this instance, memory may also need to be replaced. Refer to 3.9, "Memory" on page 84 for additional information.

For this method, submit an RPQ for assessment and pricing. The x3850 X6 configuration is evaluated and recommendations are made based on the workload requirements.

The major parts of the 4U to 8U upgrade are the 8U chassis, Storage Book, and Primary I/O Book. All of the components in the package are installed in the top portion of the chassis. The 4U system's components are transferred to the bottom section of the chassis.

Although this upgrade requires a new 8U chassis replacing the 4U chassis, most of the internal components can be moved from the x3850 X6 to the x3950 X6.

The following x3850 X6 components can be migrated to the x3950 X6 as part of the RPQ upgrade:

- ▶ Compute Books if they use Intel Xeon E7-8800 processors
- ▶ All memory DIMMs
- ▶ Storage Book
- ▶ All internal drives
- ▶ Primary I/O Book (and associated fans)
- ▶ Half-length I/O Books
- ▶ Full-length I/O Books
- ▶ All adapters
- ▶ All power supplies

The RPQ upgrade includes the following parts:

- ▶ New 8U chassis and 8-socket midplane
- ▶ A second Storage Book
- ▶ A second Primary I/O Book

The RPQ upgrade might also require the following new parts:

- ▶ More Compute Books (a minimum of four Compute Books required in the x3950 X6)
- ▶ More power supplies (a minimum of four are required in the x3950 X6)
- ▶ More I/O Books, network adapters, drives as needed



Consider the following key points regarding this upgrade:

- ▶ Intel Xeon E7-4800 processors cannot be used in an x3950 X6. If your x3850 X6 has Compute Books with E7-4800 v2, v3, or v4 processors, these components must be replaced with Compute Books with E7-8800 v2, v3, or v4 processors if you plan to upgrade to an x3950 X6. The memory in the Compute Books can be reused in the x3950 X6 if you are using the same series processor.
- ▶ All processors that are used in the x3950 X6 must be identical; for example, all E7-8850 v4 processors. A minimum of four processors are required.
- ▶ The upgrade results in the following parts no longer being used (“parts on the floor”):
  - Existing 4U chassis and 4-socket midplane
  - Compute Books that are based on E7-4800 processors
- ▶ To minimize the number of parts that cannot be used in the upgraded system, the original x3850 X6 should be configured with Compute Books that include E7-8800 processors. Because many standard models of the x3850 X6 (see Table 3-2 on page 60) contain E7-4800 processors, you might need to use CTO (by using a configurator, such as x-config) or order the server by using Special Bid to create a server configuration with E7-8800 processors.
- ▶ The RPQ upgrade process also involves transferring the x3850 X6 serial number to the x3950 X6 chassis. This transfer makes the upgrade simpler from an asset or depreciation management perspective. This transfer also means that the old 4U chassis is retired because it does not have a valid serial number.
- ▶ Ideally, all power supplies are the 1400 W variant or a combination of 1400W and 900W. Regardless of the selection, the power supplies coexistence rules that are described in 3.24, “Power subsystem” on page 122 must be followed.
- ▶ Depending on your workload and configuration, you might need to provision for more PDU outlets, cables, and power capacity for your x3950 X6 server. Use the Lenovo Power Configurator to determine your total power draw to assist you in provisioning adequate power. The Power Configurator is available at this website:  
<http://support.lenovo.com/documents/LNVO-PWRCONF>
- ▶ Another 4U of rack space is required when the x3850 X6 is upgraded to the x3950 X6 for a total of 8U of rack space.
- ▶ To upgrade the x3850 X6 server to a x3950 X6 server, you must allow for downtime. The server must be powered off and have some of its components removed for reinstallation into the new x3950 X6 server.



# Infrastructure planning

The Lenovo System x3850 X6 and x3950 X6 are enterprise-class Intel processor-based servers for mission-critical applications. These servers can replace an entire rack of conventional servers.

In this chapter, we describe infrastructure planning and considerations. This chapter includes the following topics:

- ▶ 4.1, “Physical and electrical specifications” on page 130
- ▶ 4.2, “Rack selection and rack options” on page 131
- ▶ 4.3, “Floor clearance” on page 134
- ▶ 4.4, “Using Rear Door Heat eXchanger” on page 135
- ▶ 4.5, “Power guidelines” on page 136
- ▶ 4.6, “Cooling considerations” on page 145
- ▶ 4.7, “Uninterruptible power supply units” on page 146
- ▶ 4.8, “PDU and line cord selection” on page 147

## 4.1 Physical and electrical specifications

The x3850 X6 and x3950 X6 servers have the following physical, electrical, and environmental specifications:

- ▶ x3850 X6 dimensions:
  - Height: 173 mm (6.8 in)
  - Width: 482 mm (19.0 in)
  - Depth: 804 mm (31.6 in)
  - Depth with cable management brackets installed: 836 mm (32.9 in)
  - Depth with Full-length I/O Book installed: 921 mm (36.2 in)
- ▶ x3950 X6 dimensions:
  - Height: 351 mm (13.7 in)
  - Width: 482 mm (19.0 in)
  - Depth: 804 mm (31.6 in)
  - Depth with cable management brackets installed: 836 mm (32.9 in)
  - Depth with Full-length I/O Book installed: 921 mm (36.2 in)
- ▶ x3850 X6 weight:
  - Minimum configuration: 35.9 kg (79.2 lb)
  - Typical: 46.4 kg (102.3 lb)
  - Maximum: 54.7 kg (120 lb)
- ▶ x3950 X6 weight:
  - Minimum configuration: 84.5 kg (186.3 lb)
  - Typical: 88.2 kg (194.5 lb)
  - Maximum: 110.0 kg (242.6 lb)
- ▶ Air temperature support:
  - Server on: 5 - 40 °C (41 - 104 °F); altitude: 0 - 3,050 m (10,000 ft).
  - Server off: 5 - 45 °C (41 - 113 °F); altitude: 0 - 3,050 m (10,000 ft).
  - Shipment: -40 - 60 °C (-40 - 140 °F)
- ▶ Humidity support:
  - Server on: 8% to 85%, maximum dew point 24 °C
  - Server off: 8% to 85%, maximum dew point 27 °C
- ▶ Electrical: Models with 1400 W AC power supplies:
  - 100 - 127 VAC; 50 Hz or 60 Hz; 10 A (900 W DC output)
  - 200 - 240 VAC; 50 Hz or 60 Hz; 8 A (1400 W DC output)
  - Input kilovolt-amperes (kVA) (approximately):
    - Minimum configuration: 0.16 kVA
    - Maximum configuration: 3.2 kVA
- ▶ Electrical: Models with 900 W AC power supplies:
  - 100 - 127 VAC; 50 Hz or 60 Hz; 10.7 A
  - 200 - 240 VAC; 50 Hz or 60 Hz; 5.3 A
  - Input kilovolt-amperes (kVA) (approximately):
    - Minimum configuration: 0.16 kVA
    - Maximum configuration: 2.0 kVA

- ▶ Electrical: Models with 750 W DC power supplies:
  - -40 to -75 VDC
  - Input kilovolt-amperes (kVA) (approximately):
    - Minimum configuration: 0.16 kVA
    - Maximum configuration: 1.7 kVA
- ▶ BTU output:
  - Minimum configuration: 546 Btu/hr (160 watts)
  - Maximum configuration: 10,912 Btu/hr (3,200 watts)
- ▶ Noise level:
  - 6.6 bels (operating)
  - 6.4 bels (idle)

## 4.2 Rack selection and rack options

All models of the x3850 X6 and x3950 X6 include a rail kit and cable management bracket kit for installation in a 19-inch rack cabinet. The contents of the kits are shown in Figure 4-1.




Figure 4-1 x3850 X6 and x3950 X6 rail kit and cable management bracket kit

The server supports the rack cabinets that are listed in Table 4-1.

**Full-length I/O Book:** As indicated in Table 4-1, some racks are not deep enough to support the servers with Full-length I/O Books installed.

Table 4-1 Rack cabinets

Part number	Description	Supported	Maximum supported	
			x3850 X6	x3950 X6
2018-86X	11U Office Enablement Kit	No	Not applicable	

Part number	Description	Supported	Maximum supported	
			x3850 X6	x3950 X6
9307-2PX	25U Static S2 Standard Rack	No	Not applicable	
9307-2RX	25U Standard Rack	No	Not applicable	
9307-4RX	42U Standard Rack	No	Not applicable	
9307-4XX	42U Standard Rack Extension	No	Not applicable	
9308-4EX	42U Enterprise Expansion Rack	Yes <sup>a</sup>	10	5
9308-4PX	42U Enterprise Rack	Yes <sup>a</sup>	10	5
9360-4PX	42U 1200 mm Deep Dynamic Rack	Yes	10	5
9360-4EX	42U 1200 mm Deep Dynamic Expansion Rack	Yes	10	5
9361-4PX	42U 1200 mm Deep Static Rack	Yes	10	5
9361-4EX	42U 1200 mm Deep Static Expansion Rack	Yes	10	5
9362-4PX	47U 1200 mm Deep Static Rack	Yes	11	5
9362-4EX	47U 1200 mm Deep Static Expansion Rack	Yes	11	5
9363-4CX	PureFlex System 42U Rack	Yes	10	5
9363-4DX	PureFlex System 42U Expansion Rack	Yes	10	5
9363-4PX	42U 1100 mm Dynamic Rack	Yes	10	5
9363-4EX	42U 1100 mm Dynamic Expansion Rack	Yes	10	5
1410-2RX	Intelligent Cluster™ 25U Rack Family	No	Not applicable	
1410-4RX	Intelligent Cluster 42U Rack Family	Yes <sup>a</sup>	10	5
1410-PRA	Intelligent Cluster 42U 1200 mm Deep Rack Family	Yes	10	5
1410-PRB	Intelligent Cluster 42U 1100 mm Enterprise V2 Rack Family	Yes	10	5
7200-4PX	Smart Analytics 42U Rack Family	Yes <sup>a</sup>	10	5

a. This rack has insufficient clearance at the back of the rack to accommodate the installation of the Full-length I/O Book in the x3850 X6 or x3950 X6 system.

For more information, see the list of Lenovo Press Product Guides in the Rack cabinets and options category at this website:

<http://lenovopress.com/systemx/rack>

The server supports the rack console switches and monitor kits that are listed in Table 4-2.

Table 4-2 Rack options

Part number	Feature code	Description
<b>Monitor kits and keyboard trays</b>		
17238BX	A3EK	1U 18.5-inch Standard Console
<b>Console keyboards</b>		

Part number	Feature code	Description
00MW310	AU8B	Lenovo UltraNav® Keyboard USB - US Eng
46W6712	A50G	Keyboard w/ Int. Pointing Device USB - US Eng 103P RoHS v2
46W6713	A50H	Keyboard w/ Int. Pointing Device USB - Arabic 253 RoHS v2
46W6714	A50J	Keyboard w/ Int. Pointing Device USB - Belg/UK 120 RoHS v2
46W6715	A50K	Keyboard w/ Int. Pointing Device USB - Chinese/US 467 RoHS v2
46W6716	A50L	Keyboard w/ Int. Pointing Device USB - Czech 489 RoHS v2
46W6717	A50M	Keyboard w/ Int. Pointing Device USB - Danish 159 RoHS v2
46W6718	A50N	Keyboard w/ Int. Pointing Device USB - Dutch 143 RoHS v2
46W6719	A50P	Keyboard w/ Int. Pointing Device USB - French 189 RoHS v2
46W6720	A50Q	Keyboard w/ Int. Pointing Device USB - Fr/Canada 445 RoHS v2
46W6721	A50R	Keyboard w/ Int. Pointing Device USB - German 129 RoHS v2
46W6722	A50S	Keyboard w/ Int. Pointing Device USB - Greek 219 RoHS v2
46W6723	A50T	Keyboard w/ Int. Pointing Device USB - Hebrew 212 RoHS v2
46W6724	A50U	Keyboard w/ Int. Pointing Device USB - Hungarian 208 RoHS v2
46W6725	A50V	Keyboard w/ Int. Pointing Device USB - Italian 141 RoHS v2
46W6726	A50W	Keyboard w/ Int. Pointing Device USB - Japanese 194 RoHS v2
46W6727	A50X	Keyboard w/ Int. Pointing Device USB - Korean 413 RoHS v2
46W6728	A50Y	Keyboard w/ Int. Pointing Device USB - LA Span 171 RoHS v2
46W6729	A50Z	Keyboard w/ Int. Pointing Device USB - Norwegian 155 RoHS v2
46W6730	A510	Keyboard w/ Int. Pointing Device USB - Polish 214 RoHS v2
46W6731	A511	Keyboard w/ Int. Pointing Device USB - Portugese 163 RoHS v2
46W6732	A512	Keyboard w/ Int. Pointing Device USB - Russian 441 RoHS v2
46W6733	A513	Keyboard w/ Int. Pointing Device USB - Slovak 245 RoHS v2
46W6734	A514	Keyboard w/ Int. Pointing Device USB - Spanish 172 RoHS v2
46W6735	A515	Keyboard w/ Int. Pointing Device USB - Swed/Finn 153 RoHS v2
46W6736	A516	Keyboard w/ Int. Pointing Device USB - Swiss F/G 150 RoHS v2
46W6737	A517	Keyboard w/ Int. Pointing Device USB - Thai 191 RoHS v2
46W6738	A518	Keyboard w/ Int. Pointing Device USB - Turkish 179 RoHS v2
46W6739	A519	Keyboard w/ Int. Pointing Device USB - UK Eng 166 RoHS v2
46W6740	A51A	Keyboard w/ Int. Pointing Device USB - US Euro 103P RoHS v2
46W6741	A51B	Keyboard w/ Int. Pointing Device USB - Slovenian 234 RoHS v2
<b>Console switches</b>		
1754D2X	6695	Global 4x2x32 Console Manager (GCM32)

Part number	Feature code	Description
1754D1X	6694	Global 2x2x16 Console Manager (GCM16)
1754A2X	0726	Local 2x16 Console Manager (LCM16)
1754A1X	0725	Local 1x8 Console Manager (LCM8)
<b>Console cables</b>		
43V6147	3757	Single Cable USB Conversion Option (UCO)
39M2895	3756	USB Conversion Option (4 Pack UCO)
46M5383	5341	Virtual Media Conversion Option Gen2 (VCO2)
46M5382	5340	Serial Conversion Option (SCO)

For more information, see the list of Lenovo Press Product Guides in the Rack cabinets and options category at this website:

<http://lenovopress.com/systemx/rack>

### Installing in a non Lenovo rack

The x3850 X6 and x3950 X6 servers can be installed in non Lenovo racks. The X6 systems can be installed in most industry standard 483 cm (19 in) rack cabinets that conform to the Electronic Industries Alliance (EIA) standard EIA-310-D Cabinets, Racks, Panels, and Associated Equipment (1992).

### Shipping requirements for X6

The Full-length I/O Book cannot be shipped when it is installed in the x3850 X6 or x3950 X6 server. This I/O Book and its adapters must be shipped separately from the server.

## 4.3 Floor clearance

The server is a traditional rack design that is installed in the rack rails from the front of the rack. Sufficient space to lift and mount the server is required. For more information about racking, see the *x3850 X6 Installation and Service Guide*.

When the server is mounted in a rack, it is on non-sliding rails and fixed to the rack. There is no need for more floor clearance to pull the server out of the rack for maintenance or upgrades. This extra clearance is not needed because all of the serviceable servers and components that can be upgraded can be accessed from the front or rear of the server.

Having components that are accessed from the front or rear without sliding the server is beneficial for the following reasons:

- ▶ The modular design makes the system easier to service because you need to pull only the affected subsystem without having to pull the entire server out from the rack.
- ▶ Because there is no requirement to slide the server in and out of the rack, there are no cable management arms with which to be concerned.
- ▶ The system is easier to install because you can unload all pluggable parts for lower weight when you are lifting in to the rack.
- ▶ The system is easier to upgrade from 2S to 4S by adding Compute Books.



- ▶ The system is easier to add I/O by hot-swapping and hot-adding an I/O Book to add adapters or by adding an I/O Book to the server.
- ▶ The system is easier to add memory by removing the appropriate Compute Book, adding the memory, then reinserting the Compute Book.

After the server is installed in the rack, the only floor clearance you need is for pulling out or installing pluggable components, such as the Compute Book, Storage Book, I/O Books, or power supplies.

## 4.4 Using Rear Door Heat eXchanger

A Rear Door Heat eXchanger (RDHX) is available for the 42U 1100 mm Enterprise V2 Dynamic Racks.

The RDHX has the following features:

- ▶ It attaches in place of the perforated rear door and adds 10 mm, which makes the overall package 1200 mm. This depth is the depth of two standard data center floor tiles.
- ▶ The doors use 3/4 inch quick connect couplers, which include automatic valves that restrict water leakage (often a few drops at most) when the doors are connected or disconnected.
- ▶ Each door has a capacity of 9 liters (2.4 gallons), and supports flow rates of 22.7 liters (6 gallons) to 56.8 liters (15 gallons) per minute.
- ▶ The doors have no moving parts; the fans in the equipment move air through the heat exchanger as easily as a standard rack door.

If there is a disruption in water flow, the rack reverts to standard air cooling.

- ▶ Each door can remove 100% of the heat that is generated by servers that use 30 kW of power and 90% of the heat that is generated by servers that use 40 kW, by using 18° C (64° F) water at a 27° C (81° F) server inlet air temperature.

Figure 4-2 shows the capability of the RDHX.




Figure 4-2 Heat removal performance with a 30 kW load

Although more heat can be extracted if the water is cooler, the water temperature cannot be below the dew point in the server room, or condensation forms on the rear door.

- ▶ Some standard computer room air conditioning is provisioned to control humidity and enable doors to be disconnected for maintenance or other requirements.
- ▶ The reduced air conditioning requirement typically saves approximately 1 KW per rack that is used to compress refrigerant and move air.

The reduction in air conditioner noise coupled with the acoustic dampening effect of the heat exchangers and the decrease in high velocity cold air has the added benefit of making the data center environment less hostile.

The ordering information for the rear door heat exchanger is listed in Table 4-3.

Table 4-3 Part number for the Rear Door Heat eXchanger for the 42U 1100 m rack

Part number	Description
175642X	Rear Door Heat eXchanger for 42U 1100 mm Enterprise V2 Dynamic Racks

For more information, see the *Rear Door Heat eXchanger V2 Type 1756 Installation and Maintenance Guide*, which is available at this website:

<https://support.lenovo.com/docs/UM103398>

## 4.5 Power guidelines

In this section, we provide power planning considerations, describe the power subsystem in the x3850 X6 and x3950 X6 servers, and show examples of best practices for configuring power connections and power distribution. We also provide advice for adequate cooling for racks containing these systems.

This section includes the following topics:

- ▶ 4.5.1, “Considerations” on page 136
- ▶ 4.5.2, “Power supply redundancy” on page 137
- ▶ 4.5.3, “Rules for achieving redundancy” on page 139
- ▶ 4.5.4, “Power supply installation order” on page 140
- ▶ 4.5.5, “Power policy” on page 141
- ▶ 4.5.6, “More power settings in the IMM2” on page 143
- ▶ 4.5.7, “Examples of power connections” on page 145

### 4.5.1 Considerations

When you are planning your power source for an x3850 X6 or an x3950 X6 system rack installation, consider the following variables:

- ▶ Input voltage range: 100 - 120 VAC or 200 - 240 VAC
- ▶ Power Distribution Unit (PDU) input: Single-phase or three-phase
- ▶ Power redundancies: AC source feed (power feed) N+N, power supply N+1, or power supply N (no redundancy)
- ▶ PDU control: Switched and monitored, monitored, or non-monitored
- ▶ Hardware: Quantity of components and component power draw

The following examples provide guidance about selecting PDUs, power input line cords, and PDU to server power jumper cord connections.

The following approaches can be used to provision power:

- ▶ Provision to the label rating of the power supplies so that any configuration can be supported; this approach covers all hot-swap components that can be added later.
- ▶ Provision to the maximum, calculated, or observed power that the systems can draw.

**Note:** The official power planning tool is the Power Configurator. You can determine the total power draw of your server configuration with this tool. This tool will validate N+N redundancy based on your particular configuration. For more information, see this website:

<http://support.lenovo.com/documents/LNVO-PWRCONF>

For assistance with selecting appropriate PDUs to configure, see the PDU planning guides that available at this website:

<http://support.lenovo.com/documents/LNVO-POWINF>

For assistance with rack, power, thermal and mechanical, and quoting appropriate PDU and UPSs for this system, email the Lenovo power team at [power@lenovo.com](mailto:power@lenovo.com).

## 4.5.2 Power supply redundancy

The four x3850 X6 power supply bays are divided into two power groups to support N+N power supply and power feed redundancy (where N = 1 or 2, depending on your system configuration and load).

The power supply bays are numbered from left to right when viewed from the rear of the chassis, as shown in Figure 4-3. Power supplies that are installed in bays 1 and 3 belong to Group A (red); power supplies that are installed in bays 2 and 4 belong to Group B (blue).




Figure 4-3 Grouping of the power supplies in the x3850 X6

Figure 4-4 shows the power supply bay numbering for the x3950 X6.

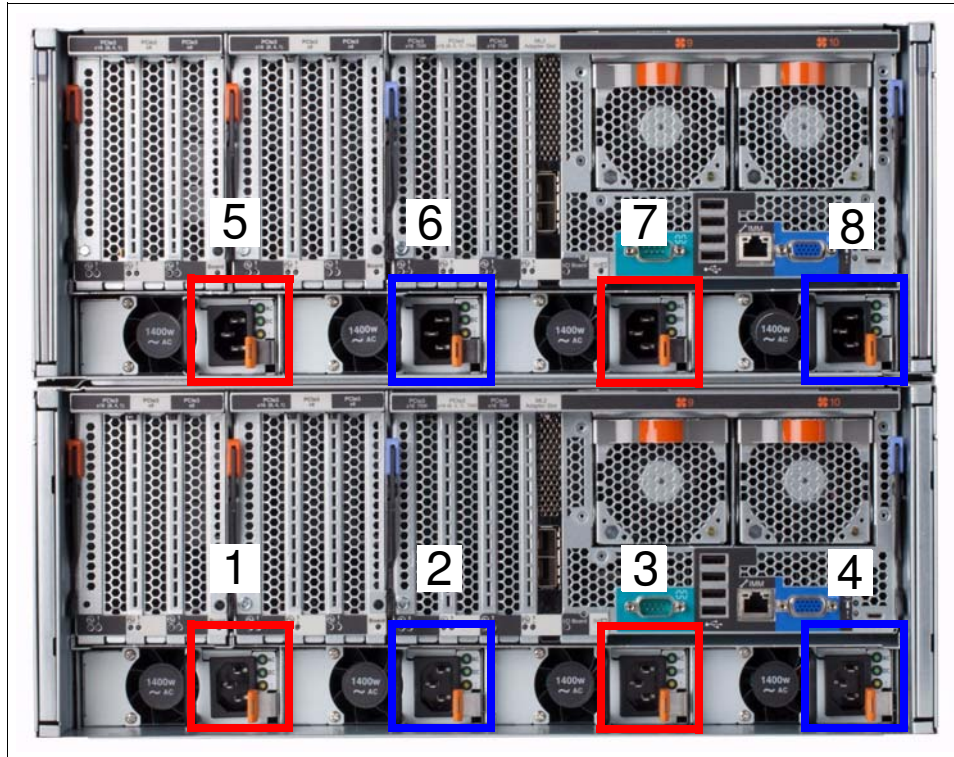


Figure 4-4 Grouping of the power supplies in the x3950 X6

The following types of power supplies are supported:

- ▶ 900 W AC
- ▶ 1400 W AC
- ▶ 1400 W AC for high altitude above 5,000 meters
- ▶ 750 W DC

The x3850 X6 and x3950 X6 support the following modes of redundancy based on the power supply configuration, system load, and the Power Policy configuration that is controlled by the Integrated Management Module 2 (IMM2):

- ▶ Non-redundant
- ▶ Fully system redundant
- ▶ Redundant with reduced performance (throttling)

The default configuration setting for AC and DC models is Non-Redundant with Throttling enabled.

You can set and change the Power Policy and System Power Configurations by using the IMM2 web interface. The power configurations and policies can be changed via the web, Advanced Settings Utility (ASU), Common Information Model (CIM), and Unified Extensible Firmware Interface (UEFI) interfaces. These settings cannot be changed by UEFI.

For more information about how to connect to the IMM2, see 7.2, “Integrated Management Module II” on page 244. For information about the power policy in the IMM2 and changing the settings, see 4.5.5, “Power policy” on page 141.

Depending on your load and power requirements, the x3850 X6 server supports the following power supply installations:

- ▶ One 900 W
- ▶ One 1400 W
- ▶ Two 900 W
- ▶ Two 1400 W
- ▶ Two 900 W and two 1400 W
- ▶ Four 900 W
- ▶ Four 1400 W
- ▶ Four 750 W DC

There are two different 1400 W power supply units available: a standard 1400 W power supply unit (PSU) and a high-altitude 1400 W PSU. The high-altitude PSU is used when the server is at an altitude higher than 5000 meters from sea level. It is not advised to mix the high-altitude PSUs with the standard 1400 W or 900 W PSUs because this configuration nullifies the high-altitude capabilities.

Table 4-4 lists the different power supply installation configurations, policies, and level of redundancy that is achieved, depending on the type and number of power supplies that are installed.

Table 4-4 Power supply installation and redundancy

Power supply 1	Power supply 2	Power supply 3	Power supply 4	Fully redundant AC source	Power module redundant	Non-redundant
1400 W	1400 W	1400 W	1400 W	Yes	Yes	Yes
1400 W	900 W	900 W	1400 W	Yes <sup>a</sup>	No	Yes
900 W	1400 W	1400 W	900 W	Yes <sup>a</sup>	No	Yes
None	1400 W	1400 W	None	Yes <sup>a</sup>	No	Yes
None	None	1400 W	None	No	No	Yes
900 W	900 W	900 W	900 W	Yes <sup>a</sup>	Yes	Yes
None	900 W	900 W	None	Yes <sup>a</sup>	No	Yes
None	None	900 W	None	No	No	Yes
750 W DC	750 W DC	750 W DC	750 W DC	Yes <sup>a</sup>	Yes	Yes

a. Workload and configuration dependent. Use the Power Configurator to plan for your level of redundancy, which is available at this website: <http://support.lenovo.com/us/en/documents/LNVO-PWRCONF>

### 4.5.3 Rules for achieving redundancy

To achieve power supply and power feed redundancy, observe the following rules:

- ▶ Power supplies in bays 1 and 3 (Group A) and bays 2 and 4 (Group B) must be balanced in quantity and power output (A = B).
- ▶ Power feeds to Group A and Group B must be different for redundancy (see Figure 4-5).
- ▶ Mixed wattages must also be balanced on Group A (900 W/1400 W) and Group B (900 W/1400 W).
- ▶ IMM2 Code requires power supplies in bays 1 and 3 to be wired to feed 1, and power supplies in bays 2 and 4 to be wired to feed 2. This configuration is necessary because of

how the power supplies are internally connected and how the IMM2 determines whether the system is redundant based on the selected policy. If the power supplies are not correctly connected, the selected IMM2 policy might not function as intended.

Figure 4-5 shows a correctly balanced and redundant system with four 900 W power supplies installed that are connected to two separate feeds.

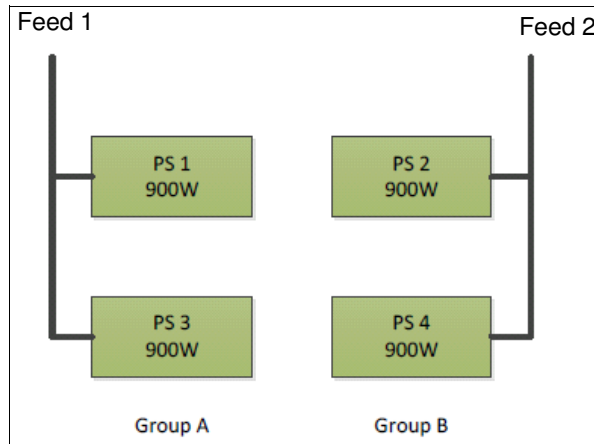


Figure 4-5 Correctly balanced and redundant power

If a mix of 900 W and 1400 W power supplies is used, ensure that power input feeds 1 and 2 have a mix of wattage. This configuration ensures source redundancy and is the only way that mixed power supplies are supported.

The following power feed configuration is supported for mixed power supplies:

- ▶ Power supply bay 1: 1400 W, Power feed 1
- ▶ Power supply bay 2: 900 W, Power feed 2
- ▶ Power supply bay 3: 900 W, Power feed 1
- ▶ Power supply bay 4: 1400 W, Power feed 2

#### 4.5.4 Power supply installation order

The power supplies are installed in the following order for the x3850 X6:

1. Power supply bay 3
2. Power supply bay 2
3. Power supply bays 1 and 4

The power supplies are installed in the following order for the x3950 X6:

1. Power supply bays 3, 2, 7 and 6
2. Power supply bays 1, 4, 5 and 8

The following rules apply when power supplies are installed in the x3850 X6:

- ▶ For one power supply configuration, install the power supply in bay 3. This configuration does not support any form of power supply redundancy. A power supply filler must be installed in bays 1, 2, and 4.
- ▶ For two power supply configurations, install the power supplies in bays 2 and 3, with each power supply on separate input power feeds. Ensure that both power supplies are of the same type in terms of wattage and both must be AC or DC.
- ▶ A configuration of three power supplies is not supported.

The following rules apply when power supplies are installed in the x3950 X6:

- ▶ For four power supply configurations, install the power supplies in bays 2, 6, 3, and 7 with each power supply on separate input power feeds. Power supplies must be all AC or all DC. Pairs of power supplies must be the same wattage.
- ▶ Only configurations of four or eight power supplies are supported. Any other combination (1, 2, 3, 5, 6, and 7) is not supported.

## 4.5.5 Power policy

The x3850 X6 and x3950 X6 servers support three modes of redundancy based on the power supply configuration, system load, and Power Policy configuration that is controlled by the IMM2, as described in 4.5.2, “Power supply redundancy” on page 137.

The default configuration setting for AC and DC models is Non-Redundant with Throttling enabled. The power configurations and policies can be changed via the IMM2, CIM, and ASU interfaces.

This section describes changing the Power Policy by using the IMM2 interface.

### Configuring the power policy by using the IMM2

Complete the following steps to configure the power policy by using the IMM2 web interface:

1. Connect to the IMM2 interface as described in 7.2, “Integrated Management Module II” on page 244. Access the Server Power Management page by selecting **Server Management** → **Power Management**. The panel in Figure 4-6 is opened.

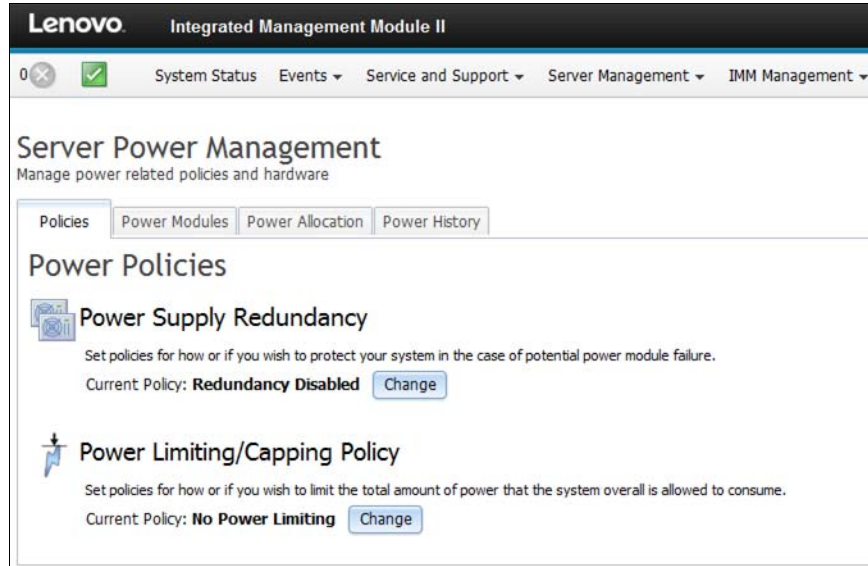


Figure 4-6 Server Power Management page in the IMM2

- To set the Power Policy of your system, ensure that you are on the Policies tab and click **Change** next to Current Policy Power Supply Redundancy.

Figure 4-7 shows the power policies window in the IMM2 for an x3850 X6 system with 4 x 1400 W power supplies installed.

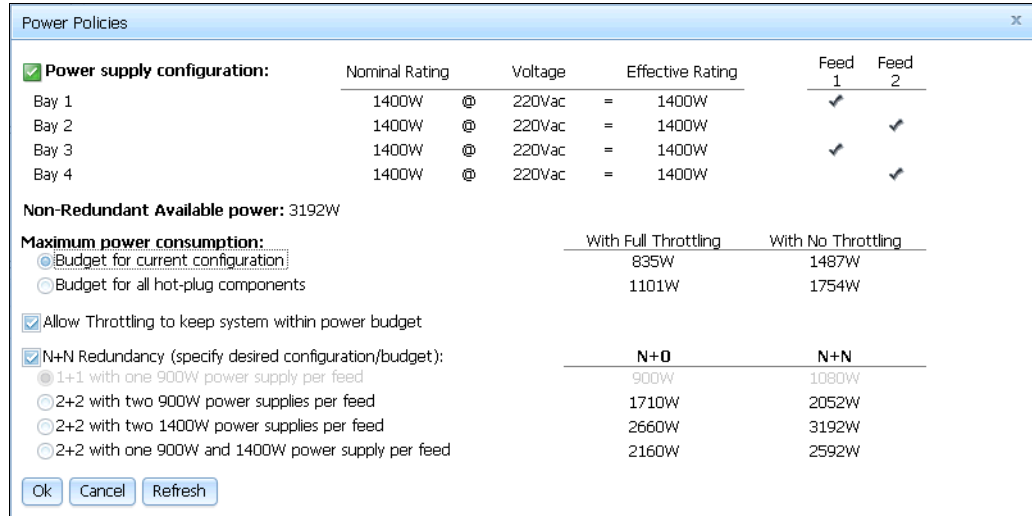


Figure 4-7 Power Policy options

- Select the policy that you want to implement on your system and click **OK** to implement that policy.

**Warning:** Select only a Power Policy that meets the hardware configuration of your system. For example, you cannot select 2+2 redundancy with only two power supplies installed. This selection can result in your system not starting.



## 4.5.6 More power settings in the IMM2

In this section, we describe the power capping and power allocation information you can set and view in the Server Power Management page via the IMM2 to help you manage your systems power usage.

### Power capping

From the Policies tab, you can also set power capping. To set the overall power limit on your server, click **Change** under the Power Limiting/Capping Policy section. Figure 4-8 shows the power capping window in the IMM2.

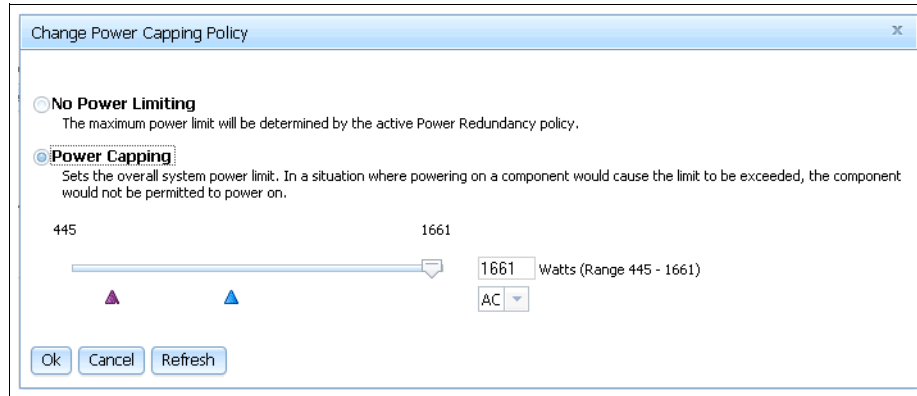


Figure 4-8 Power capping policy

Without capping enabled, the maximum power limit is determined by the active power redundancy policy that is enabled.

**Warning:** With power capping enabled, the component is not permitted to power on in a situation where powering on a component exceeds the limit.

### Power allocation and consumption

You can view your power supply usage and power consumption information via the Power Allocation tab in the Server Power Management window in the IMM2.

The Power Supply Utilization graph that is shown in Figure 4-9 on page 144 shows the theoretical maximum amount of power that all components together might use and the remaining available capacity.

The Current DC Power Consumption graph displays the theoretical power consumption of the individual components in your system, including memory, CPU, and Others.

Figure 4-9 shows the power allocation page in the IMM2.

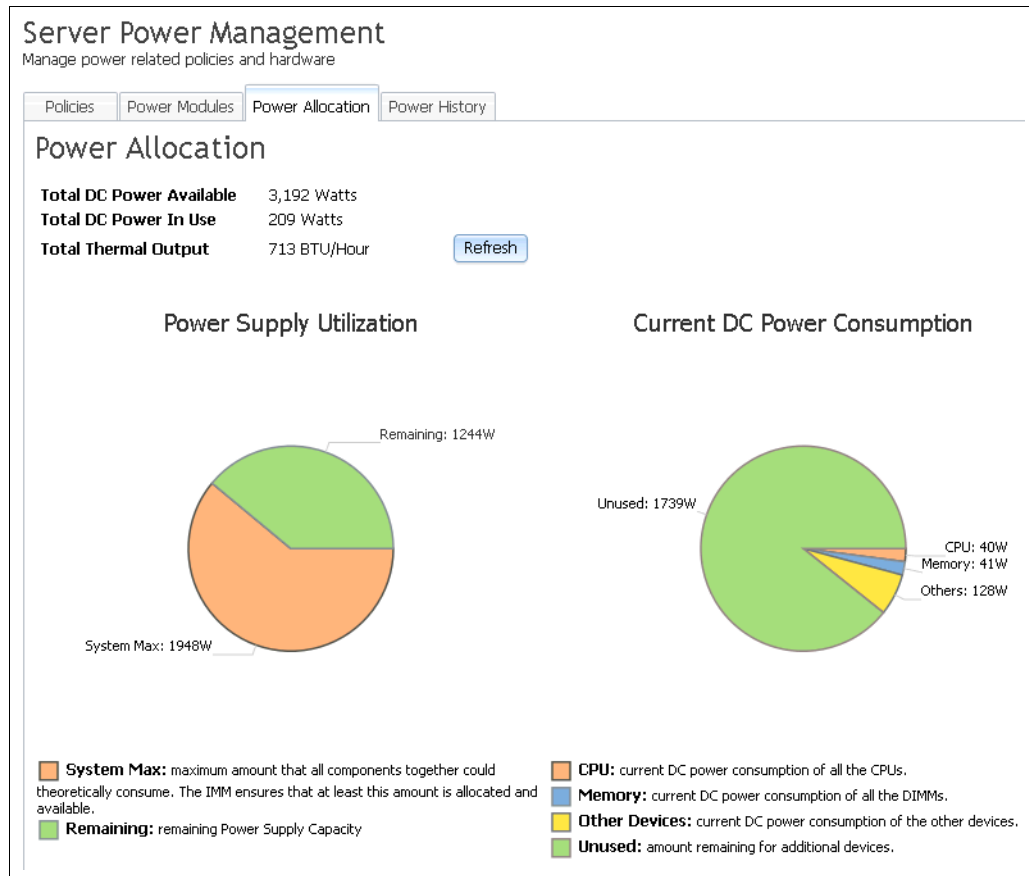


Figure 4-9 Power allocation in the Server Power Management window

The Power History tab also shows a representation of your server's power consumption over a selected period, as seen in Figure 4-10.




Figure 4-10 History of power consumption

## 4.5.7 Examples of power connections

The *System x3850 X6 and x3950 X6 Data Center Planning* book features numerous examples of the connections from single-phase and three-phase PDUs. The guide is available for download at this web page:

<http://support.lenovo.com/documents/LNVO-POWINF>

## 4.6 Cooling considerations

It is important to maintain sufficient temperature and humidity control in the data center environment to prevent thermal failures. Consider the following points to help plan for the necessary cooling of the X6 servers and rack components:

- ▶ The airflow is from the front to the back. Ensure that the airflow to the server is always unobstructed.
- ▶ Install the server and devices in a rack that has perforated front and rear doors only with sufficient openings to allow for proper airflow.
- ▶ Always install rack filler panels in empty bays to help ensure proper air circulation.
- ▶ Install side covers on single racks to help prevent air recirculation.
- ▶ Multiple racks that are placed side by side must have side covers or be joined as a suite to prevent air recirculation between the racks.
- ▶ Install a skirt or the recirculation prevention bracket on the front of the rack to prevent air circulation below the rack.

If you plan to install the optional Rear Door Heat eXchanger V2 on the 42U 1100 mm Enterprise V2 Dynamic Rack, consider the following requirements:

- ▶ Chilled, conditioned water must be provided that meets the heat exchanger specifications.
- ▶ A water supply system must be installed that is suitable for your data center.
- ▶ A redundant secondary cooling loop water supply, or enough room air conditioning to handle the heat load must be provided if the function of one or more of the heat exchangers is compromised. For example, if the rear door is opened for rack maintenance or the conditioned water supply to the door is stopped, the rack heat load is sent out into the room and must be handled by room air conditioning until the conditioned water supply is restored.
- ▶ Protective coverings must be installed to avoid tripping hazards on non-raised floors as part of the water hose management.
- ▶ The rack must have front and rear doors with at least 60% open area for airflow.
- ▶ The rack must have front to back airflow.
- ▶ Hot air recirculation from the back of the rack to the front must be prevented. An air sealing plane must be established at the front of the rack at the vertical EIA flanges. All U spaces must be occupied at the front EIA air sealing plane by a device or a U space filler.

## 4.7 Uninterruptible power supply units

There are several rack-mounted uninterruptible power supply (UPS) units that can be used with the x3850 X6 and x3950 X6 servers. The size and type of the UPS you select depends on your country's voltage and type, the number of available outlets, and your system and peripheral device power load. You can use the Power Configurator to help determine what UPS system is adequate for supporting your load.

You can download the Power Configurator and information about UPS runtimes for your power load from this web page:

<http://support.lenovo.com/documents/LNVO-PWRCONF>

The x3850 X6 and x3950 X6 support the attachment to UPS units that are listed in Table 4-5.

*Table 4-5 Uninterruptible power supply units*

Part number	Description
<b>Rack-mounted or tower UPS units</b>	
55941AX	RT1.5kVA 2U Rack or Tower UPS (100-125VAC)
55941KX	RT1.5kVA 2U Rack or Tower UPS (200-240VAC)
55942AX	RT2.2kVA 2U Rack or Tower UPS (100-125VAC)
55942KX	RT2.2kVA 2U Rack or Tower UPS (200-240VAC)
55943AX	RT3kVA 2U Rack or Tower UPS (100-125VAC)
55943KX	RT3kVA 2U Rack or Tower UPS (200-240VAC)
55945KX	RT5kVA 3U Rack or Tower UPS (200-240VAC)
55946KX	RT6kVA 3U Rack or Tower UPS (200-240VAC)
55948KX	RT8kVA 6U Rack or Tower UPS (200-240VAC)

Part number	Description
55949KX	RT11kVA 6U Rack or Tower UPS (200-240VAC)
55948PX	RT8kVA 6U 3:1 Phase Rack or Tower UPS (380-415VAC)
55949PX	RT11kVA 6U 3:1 Phase Rack or Tower UPS (380-415VAC)

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

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

## 4.8 PDU and line cord selection

There are several PDUs that can be used to power the X6 system. The size and type of the PDUs that you select depend on your country's voltage and type, number of required outlets, and the system and peripheral device power load. You can use the Power Configurator to help determine your total power load, which is available at this website:

<http://support.lenovo.com/documents/LNVO-PWRCONF>

This section describes the power supply unit to PDU line cord options that are available to connect your X6 power supplies to your PDU source.

### 4.8.1 Server-to-PDU power cord options

The power supplies that are installed in the X6 server have C14 connectors, which means that the line cord you use requires C13 plugs on the server side. One identical power cable is required for each installed power supply. The default power cable is a 2.8 m, 10 A/100 - 250 V, C13 to IEC 320-C14 Rack Power Cable (46M2982). Table 4-6 on page 148 lists power cables options for the 900 W and 1400 W power supplies.

Table 4-6 Server to PDU power cord options

Power cable part numbers	Power cable feature codes	Description
39M5375	6201	1.5m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable
39M5377	6311	2.8m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable
39M5392	6204	2.8m, 10A/100-250V, C13 to IEC 320-C20 Rack Power Cable
39M5378	6263	4.3m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable
47C2487 <sup>a</sup>	A3SS	1.2m, 10A/100-250V, 2 Short C13s to Short C14 Rack Power Cable
47C2488 <sup>a</sup>	A3ST	2.5m, 10A/100-250V, 2 Long C13s to Short C14 Rack Power Cable
47C2489 <sup>a</sup>	A3SU	2.8m, 10A/100-250V, 2 Short C13s to Long C14 Rack Power Cable
47C2490 <sup>a</sup>	A3SV	4.1m, 10A/100-250V, 2 Long C13s to Long C14 Rack Power Cable
47C2491 <sup>a</sup>	A3SW	1.2m, 16A/100-250V, 2 Short C13s to Short C20 Rack Power Cable
47C2492 <sup>a</sup>	A3SX	2.5m, 16A/100-250V, 2 Long C13s to Short C20 Rack Power Cable
47C2493 <sup>a</sup>	A3SY	2.8m, 16A/100-250V, 2 Short C13s to Long C20 Rack Power Cable
47C2494 <sup>a</sup>	A3SZ	4.1m, 16A/100-250V, 2 Long C13s to Long C20 Rack Power Cable

a. This cable is a Y cable that connects to two identical power supply units (PSUs). Y cables can be used with 900 W PSUs only. All other power supplies are not compatible with Y cables.

Because of the location of the power supplies at the back of the server, power cables with right angle plugs do not physically fit into the power supply connector; therefore, right-angle plug connectors are not supported.

## 4.8.2 PDU and line cord options

Some PDUs feature attached line cords and others require a line cord to be ordered separately based on your requirement of Three-phase power or Single-phase power. Table 4-7 lists the PDU, line cord, and phase options for North America and International PDUs.

Table 4-7 Power distribution units and line cords

Part number	Description	Line cord part number	Phase	Volts (V)	Line cord rating (Derated)	Line cord plug	Number / Type of outlets
<b>Switched and Monitored PDUs: North America</b>							
46M4002	1U 9 C19/3 C13 AEM DPI PDU	40K9614	1ph	200 - 240 V	30 A (24 A)	NEMA L6 30P	9 / C19 3 / C13
		40K9615	1ph	200 -240 V	60 A (48 A)	IEC 309 2P+G	
46M4003	1U 9 C19/3 C13 AEM 60A 3-Phase PDU	Attached	3ph	208 V	60 A (27.7 A/ph)	IEC 309 3P+G	9 / C19 3 / C13
46M4004	1U 12 C13 AEM DPI PDU	40K9614	1ph	200 - 240 V	30 A (24 A)	NEMA L6 30P	12 / C13
		40K9615	1ph	200 - 240 V	60 A (48 A)	IEC 309 2P+G	
46M4005	1U 12 C13 AEM 60A 3-Phase PDU	Attached	3ph	208 V	60 A (27.7 A/ph)	IEC 309 3P+G	12 / C13

Part number	Description	Line cord part number	Phase	Volts (V)	Line cord rating (Derated)	Line cord plug	Number / Type of outlets
46M4167	1U 9 C19/3 C13 Switched and Monitored 30A 3-Phase PDU	Attached	3ph	208 V	30 A (13.85 A/ph)	NEMA L21-30P	9 / C19 3 / C13
46M4116	0U 24 C13 Switched and Monitored 30A PDU	Attached	1ph	200 - 240 V	30 A (24 A)	NEMA L6 30P	24 / C13
46M4134	0U 12 C19/12 C13 Switched and Monitored 50A 3-Phase PDU	Attached	3ph	208 V	50 A (23.09 A/ph)	CS8365L	12 / C19 12 / C13
00YJ781	0U 20 C13/4 C19 Switched and Monitored 24A 1 Phase PDU	Attached	1ph	200-240V	24 A	L6-30P	4 / C19 20 / C13
00YJ783	0U 12 C13/12 C19 Switched and Monitored 48A 3 Phase PDU	Attached	3ph	200-240V	48 A	IEC60309 460P9	12 / C19 12 / C13
<b>Switched and Monitored PDUs: International</b>							
46M4137	0U 12 C19/12 C13 Switched and Monitored 32A 3-Phase PDU	Attached	3ph Y	380 - 415 V	32 A (32A/ph)	IEC 309 3P+N+G	12 / C19 12 / C13
46M4002	1U 9 C19/3 C13 AEM DPI PDU	40K9612	1ph	220 -240 V	32 A	IEC 309 P+N+G	9 / C19 3 / C13
		40K9613	1ph	220 - 240 V	63 A	IEC 309 P+N+G	
		40K9617	1ph	230 - 240 V	32 A	AUS/NZ 3112	
		40K9618	1ph	220 V	30 A	KSC 8305	
		40K9611	3ph Y	380 - 415 V	32 A (32 A/ph)	IEC 309 3P+N+G	
		47C2495	3ph Y	380 - 415 V	16 A (16 A/ph)	IEC 309 3P+N+G	
46M4004	1U 12 C13 AEM DPI PDU	40K9612	1ph	220 - 240 V	32 A	IEC 309 P+N+G	12 / C13
		40K9613	1ph	220 - 240 V	63 A	IEC 309 P+N+G	
		40K9617	1ph	230 - 240 V	32 A	AUS/NZ 3112	
		40K9618	1ph	220 V	30 A	KSC 8305	
		40K9611	3ph Y	380 - 415 V	32 A (32 A/ph)	IEC 309 3P+N+G	
		47C2495	3ph Y	380 - 415 V	16 A (16 A/ph)	IEC 309 3P+N+G	
00YJ780	0U 20 C13/4 C19 Switched and Monitored 32A 1 Phase PDU	Attached	1ph	200-240V	32 A	IEC60309 332P6	4 / C19 20 / C13
00YJ782	0U 18 C13/6 C19 Switched and Monitored 32A 3 Phase PDU	Attached	3ph Y	200-240V, 340-415V	32 A (32 A/ph)	IEC60309 532P6	6 / C19 18 / C13

Part number	Description	Line cord part number	Phase	Volts (V)	Line cord rating (Derated)	Line cord plug	Number / Type of outlets
<b>Enterprise PDUs - North America</b>							
71762NX	Ultra Density Enterprise PDU C19 PDU	40K9614	1ph	200 - 240 V	30 A (24 A)	NEMA L6-30P	9 / C19 3 / C13
		40K9615	1ph	200 - 240 V	60 A (48 A)	IEC 309 2P+G	
71763MU	Ultra Density Enterprise PDU C19 3-Phase 60A PDU+ Monitored	Attached	3ph	208 V	60 A (27.7 A/ph)	IEC 309 2P+G	9 / C19 3 / C13
71763NU	Ultra Density Enterprise PDU C19 3-Phase 60A PDU Basic	Attached	3ph	208 V	60 A (27.7 A/ph)	IEC 309 2P+G	9 / C19 3 / C13
39M2816	DPI C13 Enterprise PDU+ without line cord Monitored	40K9614	1ph	200 - 240 V	30 A (24 A)	NEMA L6-30P	12 / C13
		40K9615	1ph	200 - 240 V	60 A (48 A)	IEC 309 2P+G	
39Y8941	DPI Single Phase C13 Enterprise PDU without line cord	40K9614	1ph	200 - 240 V	30 A (24 A)	NEMA L6-30P	12 / C13
		40K9615	1ph	200 - 240 V	60 A (48 A)	IEC 309 2P+G	
39Y8948	DPI Single Phase C19 Enterprise PDU without line cord	40K9614	1ph	200 - 240 V	30 A (24 A)	NEMA L6-30P	6 / C19
		40K9615	1ph	200 - 240 V	60 A (48 A)	IEC 309 2P+G	
39Y8923	DPI 60A Three-Phase C19 Enterprise PDU with IEC309 3P+G (208 V) fixed line cord	Attached	3ph	208 V	60 A (27.7 A/ph)	IEC 309 3P+G	6 / C19
<b>Enterprise PDUs: International</b>							
71762NX	Ultra Density Enterprise PDU C19 PDU (WW)	40K9612	1ph	220 - 240 V	32 A	IEC 309 P+N+G	9 / C19 3 / C13
		40K9613	1ph	220 - 240 V	63 A	IEC 309 P+N+G	
		40K9617	1ph	230 - 240 V	32 A	AUS/NZ 3112	
		40K9618	1ph	220 V	30 A	KSC 8305	
		40K9611	3ph Y	380 - 415 V	32 A (32 A/ph)	IEC 309 3P+N+G	
		47C2495	3ph Y	380 - 415 V	16 A (16 A/ph)	IEC 309 3P+N+G	
71762MX	Ultra Density Enterprise PDU C19 PDU+ (WW)	40K9612	1ph	220 - 240 V	32 A	IEC 309 P+N+G	9 / C19 3 / C13
		40K9613	1ph	220 - 240 V	63 A	IEC 309 P+N+G	
		40K9617	1ph	230 - 240 V	32 A	AUS/NZ 3112	
		40K9618	1ph	220 V	30 A	KSC 8305	
		40K9611	3ph Y	380 - 415 V	32 A (32 A/ph)	IEC 309 3P+N+G	
		47C2495	3ph Y	380 - 415 V	16 A (16 A/ph)	IEC 309 3P+N+G	



Part number	Description	Line cord part number	Phase	Volts (V)	Line cord rating (Derated)	Line cord plug	Number / Type of outlets
39M2816	DPI C13 Enterprise PDU without line cord Monitored	40K9612	1ph	220 - 240 V	32 A	IEC 309 P+N+G	12 / C13
		40K9613	1ph	220 - 240 V	63 A	IEC 309 P+N+G	
		40K9617	1ph	230 - 240 V	32 A	AUS/NZ 3112	
		40K9618	1ph	220V	30 A	KSC 8305	
		40K9611	3ph Y	380 - 415 V	32 A (32 A/ph)	IEC 309 3P+N+G	
		47C2495	3ph Y	380 - 415 V	16 A (16 A/ph)	IEC 309 3P+N+G	
39Y8941	DPI Single Phase C13 Enterprise PDU without line cord	40K9612	1ph	220 - 240 V	32 A	IEC 309 P+N+G	12 / C13
		40K9613	1ph	220 - 240 V	63 A	IEC 309 P+N+G	
		40K9617	1ph	230 - 240 V	32 A	AUS/NZ 3112	
		40K9618	1ph	220 V	30 A	KSC 8305	
		40K9611	3ph Y	380 - 415 V	32 A (32 A/ph)	IEC 309 3P+N+G	
		47C2495	3ph Y	380 - 415 V	16 A (16 A/ph)	IEC 309 3P+N+G	
39Y8948	DPI Single Phase C19 Enterprise PDU without line cord	40K9612	1ph	220 - 240 V	32 A	IEC 309 P+N+G	6 / C19
		40K9613	1ph	220 - 240 V	63 A	IEC 309 P+N+G	
		40K9617	1ph	230 - 240 V	32 A	AUS/NZ 3112	
		40K9618	1ph	220 V	30 A	KSC 8305	
		40K9611	3ph Y	380 - 415 V	32 A (32 A/ph)	IEC 309 3P+N+G	
		47C2495	3ph Y	380 - 415 V	16 A (16 A/ph)	IEC 309 3P+N+G	
<b>Front-end PDUs: North America</b>							
39Y8938	30 amp/125V Front-end PDU	Included	1ph	125 V	30 A (24 A)	NEMA L5-30P	3 / C19
39Y8939	30 amp/240V Front-end PDU	Included	1ph	200 - 240 V	30 A (24 A)	NEMA L6-30P	3 / C19
39Y8940	60 amp Front-end PDU	Included	1ph	200 - 240 V	60 A (48 A)	IEC 309 2P+G	3 / C19
<b>Front-end PDUs: International</b>							
39Y8934	DPI 32 amp Front-end PDU	Included	1ph	220 - 240 V	32 A	IEC 309 P+N+G	3 / C19
39Y8935	DPI 63 amp Front-end PDU	Included	1ph	220 - 240 V	63 A	IEC 309 P+N+G	3 / C19
<b>Universal Rack PDUs</b>							
00YE443	DPI Universal Rack PDU	39M5389	1ph	200 - 240 V	20 A	2.5m IEC 320 C19 to C20	7 / C13
		39M5388				2.0m IEC 320 C19 to C20	
		41Y9222		100-127V	16A	4.3m NEMA L5-20 20A 16A / 100-127V Single Phase (LV)	

Part number	Description	Line cord part number	Phase	Volts (V)	Line cord rating (Derated)	Line cord plug	Number / Type of outlets
00YE443 (con't)	DPI Universal Rack PDU (continued)	40K9772	1ph (con't)	200-240V	16A	4.3m NEMA L6-20 20A 16A / 200-240V Single Phase (HV)	7 / C13 (con't)
		40K9770		230V	16A	4.3m SABS 164 single phase 16A / 230V	
		40K9766		230V	16A	4.3m CEE7-VII single phase 16A / 230V	
		40K9767		230V	13A	4.3m BS 1363 / A single phase 13A / 230V	
		40K9773		230V	15A	4.3m AS / NZ 3112 single phase 15A / 230V	
		40K9774		220V	16A	4.3m GB 2099.1 single phase 16A / 220V	
		69Y1989		250V	16A	4.3m 16A/250V C19 to NBR 14136 line cord	
		40K9776		240V	16A	4.3m IS6538 single phase 16A / 240V line cord	
		40K9777		220V	16A	4.3m IRAM 2037 single phase 16A / 220V line cord	
		40K9771		230V	16A	4.3m SI 32 single phase 16A / 230V line cord	
		40K9768		230V	16A	4.3m CEI 23-16 single phase 16A / 230V line cord	
		40K9769		230V	16A	4.3 meter Dmk/Swiss C19 - Dmk/Swiss line cord	
		81Y2391		250V	16A	4.3m 16A / 250V C19/SEV 1011 T23 line cord	
		41Y9229		125V	16A	4.3m, 16A/125V, C19 to CNS 10917-3 line cord	
		41Y9230		250V	16A	4.3m, 16A/250V, C19 to CNS 10917-3 line cord	
		41Y9232		100V	15A	4.3m, 15A/100V, C19 to JIS C-8303 line cord	

Part number	Description	Line cord part number	Phase	Volts (V)	Line cord rating (Derated)	Line cord plug	Number / Type of outlets
00YE443 (con't)	DPI Universal Rack PDU (con't)	41Y9233	1ph (con't)	200V	15A	4.3m 15A/200V, C19 to JIS C-8303 line cord	7 / C13 (con't)
<b>0U Basic PDUs: North America</b>							
00YJ776	0U 36 C13/6 C19 30A 1 Phase PDU	Attached	1ph	200-240V	24A	L6-30P	6 / C19 36 / C13
00YJ779	0U 21 C13/12 C19 60A 3 Phase PDU	Attached	3ph	200-240V	48A	IEC 603-309 460P9	12 / C19 21 / C13
<b>0U Basic PDUs: International</b>							
46M4143	0U 12 C19/12 C13 32A 3-Phase PDU	Attached	3ph Y	380 - 415 V	32 A (32 A/ph)	IEC 309 3P+N+G	12 / C19 12 / C13
00YJ777	0U 36 C13/6 C19 32A 1 Phase PDU	Attached	1ph	200-240V	32A A	IEC60309 332P6	6 / C19 36 / C13
00YJ778	0U 21 C13/12 C19 32A 3 Phase PDU	Attached	3ph Y	200-240V, 350-450V	32 A (32 A/ph)	IEC60309 532P6	12 / C19 21 / C13

For more information about PDUs and cables, see the North America, Japan, and International PDU planning guides that are available at this website:

<https://support.lenovo.com/documents/LNVO-POWINF>



# Preparing the hardware

This chapter provides assistance for making configuration, monitoring, and maintenance decisions when an x3850 X6 or x3950 X6 server is implemented. Some of the topics that are described in this section are only guidelines. Individual requirements and settings can vary from those requirements and settings that are outlined here.

This chapter includes the following topics:

- ▶ 5.1, “Configuring the IMM2 settings” on page 156
- ▶ 5.2, “Unified Extensible Firmware Interface settings” on page 165
- ▶ 5.3, “UEFI common settings” on page 169
- ▶ 5.4, “Use of NVMe drives and adapters” on page 184
- ▶ 5.5, “PCIe adapter placement advice” on page 198
- ▶ 5.6, “Hot-swap procedures” on page 201
- ▶ 5.7, “Partitioning the x3950 X6” on page 205
- ▶ 5.8, “Updating firmware” on page 211
- ▶ 5.9, “Troubleshooting” on page 221

## 5.1 Configuring the IMM2 settings

For any successful server implementation, you provide access to perform troubleshooting or routine maintenance. The x3850 X6 and x3950 X6 servers ship standard with the Integrated Management Module (IMM2). The IMM2 is a separate, independent service processor that is running if the server is plugged into a power source. The IMM2 monitors the hardware components of the server for potential hardware faults.

This section includes the following topics:

- ▶ 5.1.1, “IMM2 virtual presence” on page 156
- ▶ 5.1.2, “IMM2 network access” on page 157
- ▶ 5.1.3, “Configuring the IMM2 network interface” on page 158
- ▶ 5.1.4, “IMM2 dedicated versus shared ML2 Ethernet port” on page 159
- ▶ 5.1.5, “x3950 X6 IMM2 communication” on page 160
- ▶ 5.1.6, “Troubleshooting IMM2 communications” on page 161
- ▶ 5.1.7, “IMM2 functions to diagnose and manage the server” on page 161

### 5.1.1 IMM2 virtual presence

The IMM2 provides a virtual presence capability for remote server management capabilities. It also provides remote server management through the following industry-standard interfaces:

- ▶ Intelligent Platform Management Interface (IPMI) Version 2.0
- ▶ Simple Network Management Protocol (SNMP) Version 3
- ▶ Common Information Model (CIM)
- ▶ Web browser (supported browsers include Internet Explorer 8 and higher, Firefox 3.6 and higher, and Chrome 13 and higher)

The remote presence provides the following functions:

- ▶ Remotely viewing video with graphics resolutions up to 1600 x 1200 at 75 Hz with up to 23 bits per pixel colors, regardless of the system state.
- ▶ Remotely accessing the server by using the keyboard and mouse from a remote client.
- ▶ Mapping the CD or DVD drive, diskette drive, and USB flash drive on a remote client, and mapping ISO and diskette image files as virtual drives that can be used by the server.
- ▶ Uploading a diskette image to the IMM2 memory and mapping it to the server as a virtual drive.

Part of the information that is stored in the IMM2 can be accessed with F1-Setup by selecting **System Settings** → **Integrated Management Module**. Figure 5-1 shows the first panel of the IMM2 configuration panel.




Figure 5-1 Integrated Management Module panel

**Tip:** If you have a server but you do not know the logon credentials, you can reset the credentials by going to the panel that is shown in Figure 5-1. From F1-Setup, you can restore the IMM2 configuration to the factory defaults by selecting **Reset IMM to Defaults**.

The next section describes how to connect to the IMM2, set up your network, and gain a remote presence.

## 5.1.2 IMM2 network access

The greatest strength of the IMM2 is the ability to completely monitor and manage the server from over the network. How much functionality you have through this remote access depends entirely on your configuration of the IMM2.

### IMM2 default configuration

The default network connection for the IMM2 on the x3850 X6 is through the System Management port on the back of the server (see Figure 5-3 on page 159). The IMM2 features the following the default settings from the factory:

- ▶ Network IP: DHCP. If DHCP fails, use the following static address:
  - IP Address: 192.168.70.125
  - Subnet Mask: 255:255:255:0
  - Gateway: 0.0.0.0

- ▶ Default user ID: USERID
- ▶ Default password: PASSWORD where the 0 is a zero.

For more information about how to access the IMM2, see 7.2, “Integrated Management Module II” on page 244.

### 5.1.3 Configuring the IMM2 network interface

The IMM2 provides two paths to establish a network connection between you and the IMM2 by setting Dedicated or Shared for the Network Interface Port in the Network Configuration panel of F1-Setup. In F1-Setup, you can access this panel by selecting **System Setting** → **Integrated Management Module** → **Network Configuration**, as shown in Figure 5-2.

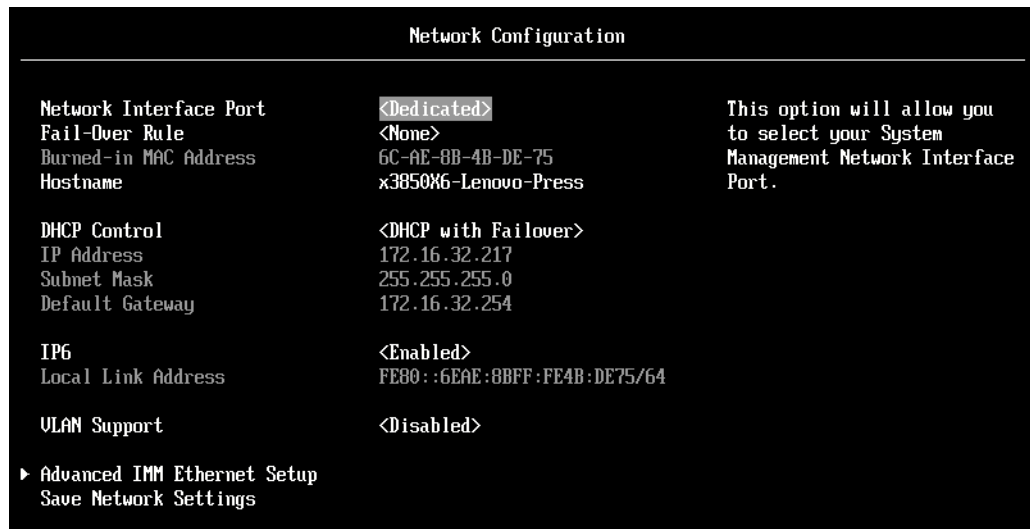


Figure 5-2 Integrated Management Module, Network Configuration

In 5.1.4, “IMM2 dedicated versus shared ML2 Ethernet port” on page 159, the reasons for the use of only the dedicated Ethernet port for IMM2 access or sharing the access to an ML2 Ethernet port are described. Also described are the steps to take to enable shared access.

You can also change the IMM default IP address in the Network Configuration panel, as shown in Figure 5-2.



## 5.1.4 IMM2 dedicated versus shared ML2 Ethernet port

When configured as *Dedicated*, you are connecting to the network via the system management port. As shown in Figure 5-3, the system management port is on the rear of the server to the right side of the USB ports.




Figure 5-3 Dedicated Gigabit IMM2 system management port

The use of this port allows for easier separation of public and management network traffic. Separating the traffic is done when you connect your public network port to switch ports that belong to a public access virtual LAN (VLAN). The management port is connected to a switch port defined by a separate management VLAN.

When configured as *Shared*, you are sharing network traffic on an ML2 Ethernet adapter, which is the adapter slot closest to the rear fans, as shown in Figure 5-4. The default port that is used on the ML2 adapter is port 1.




Figure 5-4 Location of the ML2 adapter slot that can share IMM2 access

Although this configuration eliminates a physical switch port and patch cable configuration, the media access control (MAC) address for the shared Ethernet port and the MAC address for the IMM2 address through this single network port. This situation means that there are at least two separate IP addresses for the same physical port, which prevents you from configuring the ML2 adapter's Ethernet port in a network team by using 802.3ad load balancing if the particular ML2 adapter that you use supports this function.

To maintain separation between public and management traffic, network teaming software must be used to establish a VLAN to be used by the server to send public-tagged traffic to the network switch. The switch port must be configured as a trunk port to support the public-tagged VLAN traffic and the untagged traffic for the management. The management VLAN must be defined as the native VLAN on the switch port so that its untagged traffic from the switch is accepted by the IMM2 MAC and dropped by the second Ethernet port's MAC.

Although the IMM2 uses a dedicated 32-bit RISC processor, there are limitations as to the amount of network traffic that the IMM2 can be exposed to before complex functions, such as starting from a remote DVD, or USB storage becomes unreliable because of timing issues. Although the operating system has all of the necessary drivers in place to deal with these timing issues, the Unified Extensible Firmware Interface (UEFI) is not as tolerant. For this reason (maintaining secured access), the IMM2 should be kept on a separate management network.

### 5.1.5 x3950 X6 IMM2 communication

The x3950 X6 has two IMM2 management processors. The two IMM2s in the x3950 X6 communicate with each other internally through the mid-plane. Because the IMM2s are internally linked, discovery, communication, sharing information that is related to partitioning, critical events, and management can be achieved with only one connection to the primary IMM2.

The primary IMM2 is on the lower half of the x3950 X6, as seen in Figure 5-5.



Figure 5-5 Primary IMM2 port at the back of an x3950 X6

For each IMM2 in the x3950 X6 server, you assign its own IP address. The Primary IMM2 is at the bottom half of the server and is used to connect.

**XClarity and the x3950 X6:** To properly manage the x3950 X6, both IMM2 interfaces must be connected with valid IP addresses.

## 5.1.6 Troubleshooting IMM2 communications

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

<https://support.lenovo.com/docs/UM103336>

Most communication errors are because of network switch configuration options, such as blocked ports or VLAN mismatches. The following procedure shows you how to determine this type of problem by connecting directly to the IMM2 port with a notebook and Ethernet patch cable, pinging, and then starting a web session.

If you can ping the IMM, you have a good direct network link. If the web session fails, complete the following steps:

1. Try another web browser (see supported browsers 5.1.1, “IMM2 virtual presence” on page 156).
2. Directly access the IMM configuration panel and reset the IMM2 in F1-Setup by selecting **System Settings** → **Integrated Management Module** → **Reset IMM**. You must wait approximately 5 minutes for the IMM2 to complete enough of its restart to allow you to ping it. This IMM2 reset has no affect on the operating system that is running on the server.
3. Try clearing the web browser cache.
4. Load the factory default settings back on the IMM2 through F1-Setup by selecting **System Settings** → **Integrated Management Module** → **Reset IMM2 to Defaults**. The IMM2 must be reset again after the defaults are loaded.
5. Contact Lenovo support.

## 5.1.7 IMM2 functions to diagnose and manage the server

This section provides more problem determination tips for the IMM. This section includes the following topics:

- ▶ “System Status” on page 162
- ▶ “Virtual light path diagnostics” on page 163
- ▶ “Hardware event log” on page 163
- ▶ “Remote control” on page 164

## System Status

The first panel that you see after completing the login credentials and setting your timeout period is the System Status and Hardware Health page, as shown in Figure 5-6. This panel provides a quick summary review of the hardware status of the server.

The screenshot displays the 'System Status' page in the Lenovo Integrated Management Module II. At the top, there is a navigation bar with a green checkmark icon and menu items: System Status, Events, Service and Support, Server Management, and IMM Management. Below the navigation bar, the system is identified as 'System x3850 X6' with a link to 'Add System Descriptive Name...'. The host name is 'IMM2-x3850X6-Brdwell' with a 'Rename...' link. A brief description states: '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.'

The 'System Status' section shows 'Power: On' and 'System state: System running in setup'. Below this are four buttons: 'System Information', 'Power Actions', 'Remote Control...', and 'Latest OS Failure Screen'.

The 'Active Events' section features a table with columns for Severity, Source, Date, and Message.

The 'Hardware Health' section contains a table with two columns: Component Type and Status. All components are listed as 'Normal'.

Component Type	Status
Cooling Devices	Normal
Power Modules	Normal
Local Storage	Normal
Processors	Normal
Memory	Normal
System	Normal

Figure 5-6 IMM2 System Status panel

A green tick at the top menu bar indicates that all is working from a strict hardware perspective. The IMM2 can check on the status of server components, the ServeRAID controllers, and PCIe interfaces to most PCIe adapters.

The IMM2 does not check the functional status of most PCIe adapters regarding their hardware connections to external devices. You must refer to the system event log from within the operating system or the switch logs of the network and fiber switches to which the server is connected to resolve connectivity issues.

When a hardware error is detected in the server, a red X replaces the green tick. The Hardware Health summary at the bottom of the page provides information about any hardware errors that are unresolved in the server. This information also is represented by a red X.

## Virtual light path diagnostics

If you are physically in front of the server, it is easy to track hardware problems by noticing the first tier of light path diagnostics by using the LED panel at the front of the server, as shown in Figure 5-7.



Figure 5-7 LED light path diagnostic panel

Scrolling through the LED panel by using the up and down arrows that are shown in Figure 5-7 shows the hardware subsystem that is experiencing an error.

Most servers are not physically near the people who manage them. To help you see the event from a remote location, you can use the IMM2 to review all tiers of the light path diagnostics via the System Status page under the Hardware Health menu, as seen in Figure 5-6 on page 162. Clicking a hardware component shows the hardware component that is causing the error. A Hardware Status example of Local Storage is shown in Figure 5-8.

### Local Storage

Display storage devices physical structure and storage configuration. You can refresh to get latest status.

[Refresh](#)

Physical Resource | **Storage RAID Configuration** | RAID Logs

Click on a device to see active events and properties.

#### RAID Controllers and Physical Drives

Name	Health Status	Capacity	Serial No
<input type="checkbox"/> ServeRAID M5210(PCI Slot 12)			SV34417128
Drive 0	<input checked="" type="checkbox"/> Normal	745.213GB	BTWL334203S4800RGNSSDS C2BB800G
Drive 1	<input checked="" type="checkbox"/> Normal	745.213GB	BTWL334201S7800RGNSSDS C2BB800G

Figure 5-8 Hardware Status example

## Hardware event log

For more information, including the events that led up to a failure, you can access the hardware event log. Although not every event in the hardware event log is an event that needs attention, the event log can provide insight into the cause or conditions that led up to the failure. The event log can be saved to a text file that can be sent to Lenovo support.

Figure 5-9 shows the IMM Event Log for various errors.

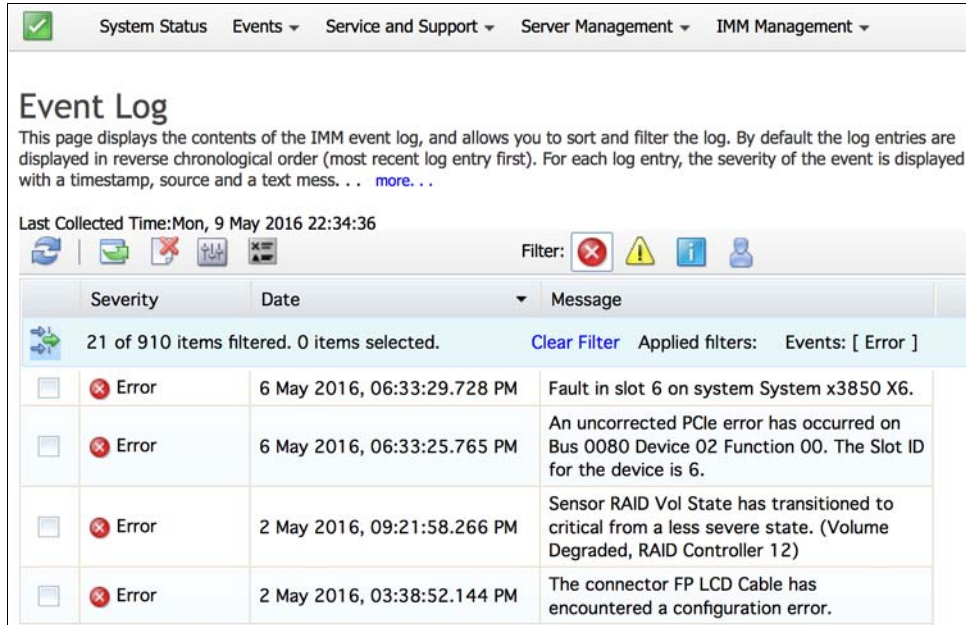


Figure 5-9 IMM Event Log for various errors

## Remote control

Certain problems require that you enter the operating system or F1-Setup to detect or fix the problems. For remotely managed servers, you can use the Remote Control feature of the IMM, which is accessed via the top menu by selecting **Server Management** → **Remote Control**.

Figure 5-10 shows the available options for starting a remote control session.

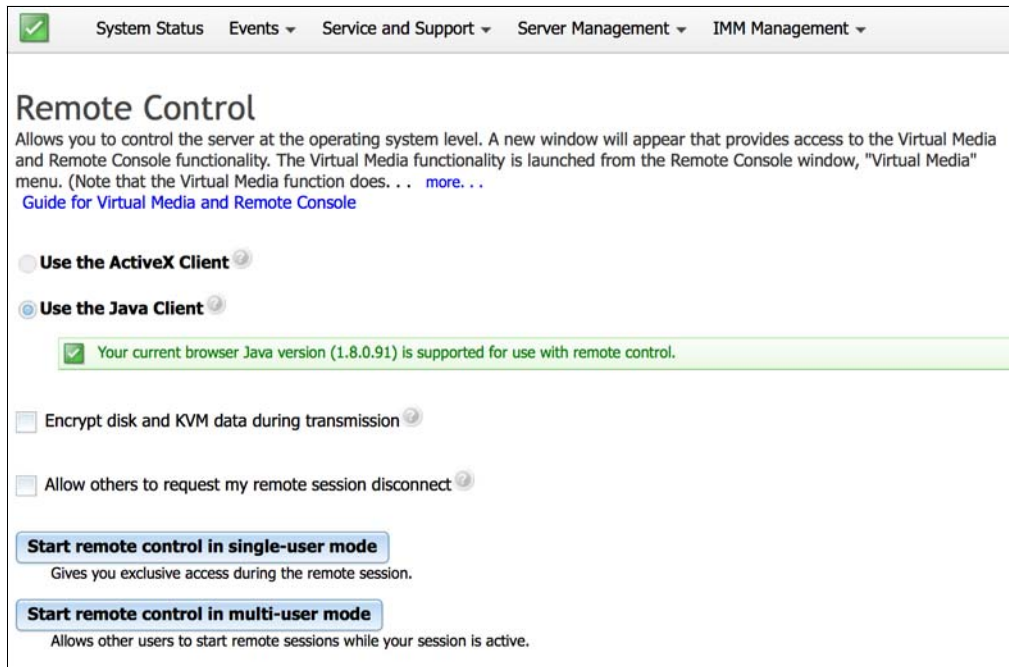


Figure 5-10 Integrated Management Module Remote Control session start panel

The IMM2 Remote Control provides the following features:

- ▶ The remote control provides you with the same capability that you have with a keyboard, mouse, and video panel that is directly connected to the server.
- ▶ You can encrypt the session when it is used over public networks.
- ▶ You can use local storage or ISO files as mounted storage resources on the remote server that you are using. These storage resources can be unmounted, changed, and remounted throughout the session, as needed.

When combined with the Power/Restart functions of the IMM, you can power down, restart, or power on the server while maintaining the same remote control session. For more information about the remote control feature on the X6 servers, see 7.3, “Remote control” on page 248.

## 5.2 Unified Extensible Firmware Interface settings

The UEFI is the interface between the operating system (OS) and platform firmware. UEFI provides a modern, well-defined environment for starting an OS and running pre-start applications.

UEFI is the replacement for BIOS. BIOS was available for many years but was not designed to handle the amount of hardware that can be added to a server today. New System x models implement UEFI to use its advanced features. For more information about UEFI, see this website:

<http://www.uefi.org/home/>

The UEFI page is accessed by pressing F1 during the system initializing process, as shown in Figure 5-11.

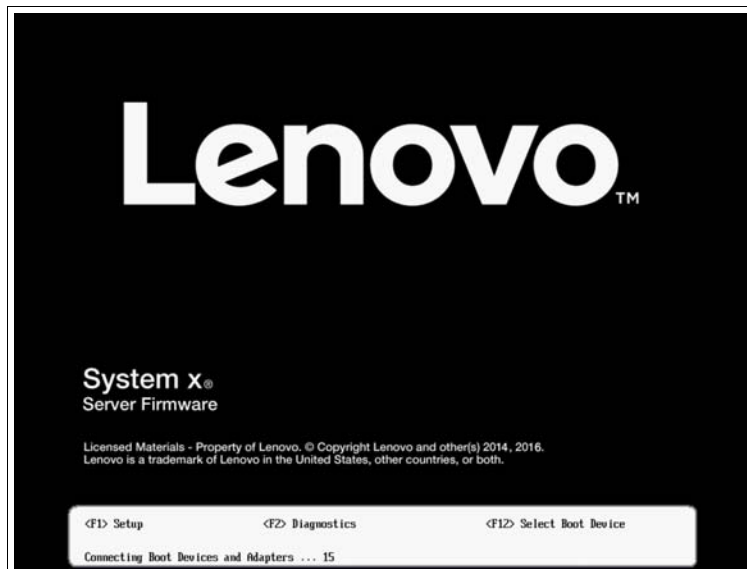


Figure 5-11 UEFI window on system start

You can change UEFI settings to meet your system requirements. In this section, we provide an overview of the UEFI settings for tuning your system for performance using different power saving and RAS options. The UEFI home page is shown in Figure 5-12.

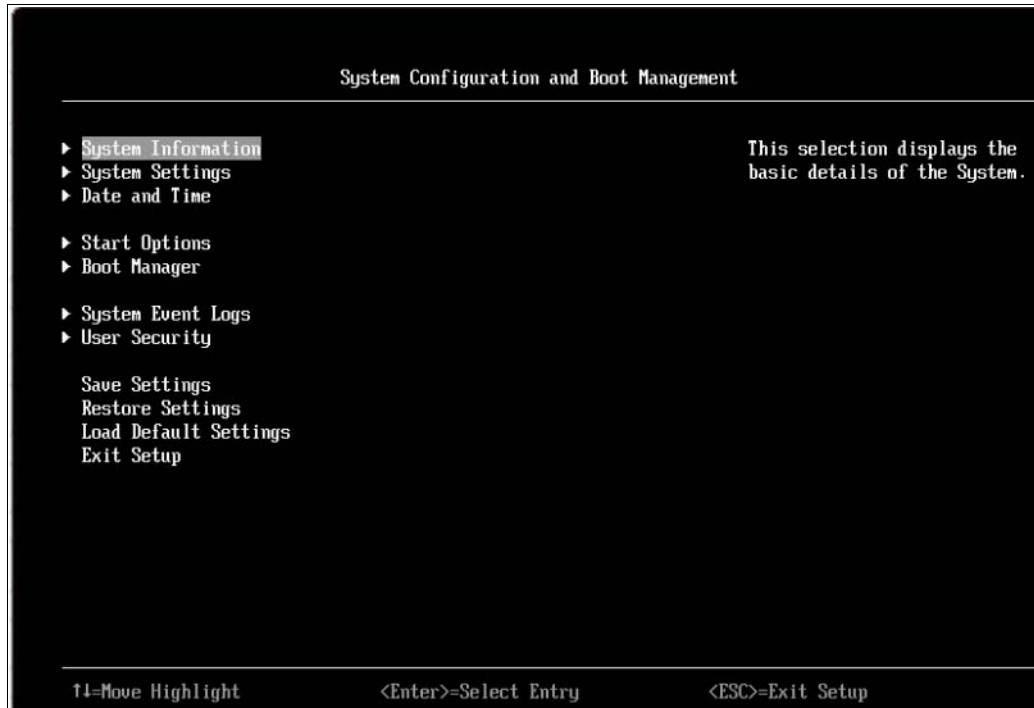


Figure 5-12 UEFI settings main panel

Choose **System Settings** to access the system settings options that are described here, as shown in Figure 5-13.




Figure 5-13 UEFI System Settings panel

You can use the Advanced Settings Utility (ASU) tool to change the UEFI settings values. ASU makes more settings available than the settings that are accessed by using the F1-Setup panel. For more information about ASU, see 7.8, “Advanced Settings Utility” on page 289.



Table 5-1 lists the most commonly used UEFI settings and their default values.

Table 5-1 UEFI default settings

UEFI value	Default setting
<b>Processor settings</b>	
Turbo Mode (Intel Turbo Boost)	Enabled
CPU P-State Control	Autonomous
CPU C-States	Autonomous
C1 Enhanced Mode	Enabled
Hyper-Threading	Enabled
Execute Disable Bit	Enabled
Intel Virtualization Technology	Enabled
Enable SMX Disabled	Disabled
Hardware Prefetcher	Enabled
Adjacent Cache Prefetch	Enabled
DCU Streamer Prefetcher	Enabled
DCU IP Prefetcher	Enabled
Direct Cache Access (DCA)	Enabled
Cores in CPU Package	All
QPI Link Frequency	Max Performance
Energy Efficient Turbo	Enabled
Uncore Frequency Scaling	Enabled
CPU Frequency Scaling	Full turbo uplift
AES-NI	Enabled
<b>Memory Settings</b>	
Memory Mode	Lockstep
Memory Speed	Max Performance
Memory Power Management	Disabled
Socket Interleave	NUMA
Patrol Scrub	Enabled
Memory Data Scrambling	Enabled
Mirroring	Disabled
Sparing	Disabled
<b>Power</b>	
Power/Performance Bias	Platform Controlled
Platform Controlled Type	Efficiency - Favor Performance

UEFI value	Default setting
Workload Configuration	Balanced
<b>Advanced RAS</b>	
Machine Check Recovery	Enabled
Hardware Memory Dump	Disabled
Dump Partition	Not Present
Post Process	Spinloop

## 5.2.1 Operating modes

X6 servers provide optimal performance with reasonable power consumption, which depends on the operating frequency and voltage of the processors and memory subsystem. The operating frequency and voltage of the processors and memory subsystem affect the system fan speed that adjusts to the current cooling requirement of the server.

In most operating conditions, the default settings provide the best performance possible without wasting energy during off-peak usage. However, for certain workloads, it might be appropriate to change these settings to meet specific power to performance requirements.

The UEFI provides several predefined setups for commonly wanted operation conditions. These predefined values are referred to as *operating modes*. Access the menu in UEFI by selecting **System Settings** → **Operating Modes** → **Choose Operating Mode**. You can see the five operating modes from which to choose, as shown in Figure 5-14. When a mode is chosen, the affected settings change to the appropriate predetermined values.

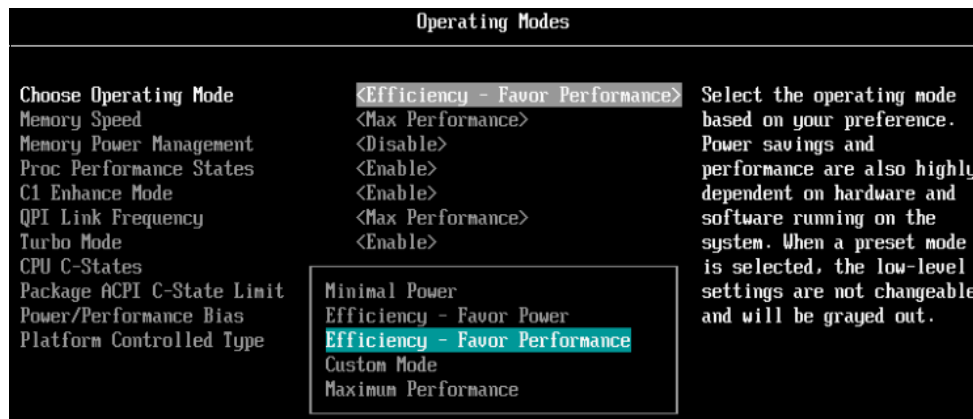


Figure 5-14 Choose Operating Mode in UEFI System Settings

The following Operating Modes are available (the default mode is Efficiency - Favor Performance):

- ▶ Minimal Power

Select this mode to minimize the absolute power consumption of the system during operation. Server performance in this mode might be reduced, depending on the application that is running.

- ▶ **Efficiency - Favor Power**

Select this mode to configure the server to draw the minimum amount of power and generate the least noise. Server performance might be degraded, depending on the application that you are running. This mode provides the best features for reducing power and increasing performance in applications where the maximum bus speeds are not critical.
- ▶ **Efficiency - Favor Performance**

Select this mode to maintain the optimal balance between performance and power consumption. The server generally produces the best performance per watt while it is in this mode. No bus speeds are derated in this mode. This is the default setting.
- ▶ **Maximum Performance**

Select this mode for the maximum performance for most server applications. The power consumption in this mode is often higher than in the Efficiency - Favor Power or Efficiency - Favor Performance mode.

Power saving and performance are also highly dependent on the hardware and software that is running on the system.
- ▶ **Custom Mode**

Select this mode only if you understand the function of the low-level IMM2 settings. This mode is the only choice with which you can change the low-level IMM2 settings that affect the performance and power consumption of the server.

## 5.3 UEFI common settings

The UEFI default settings are configured to provide optimal performance with reasonable power consumption. Other operating modes as described in 5.2.1, “Operating modes” on page 168 are also available to meet various power and performance requirements. However, individual system settings enable users to fine-tune the characteristics of the X6 servers.

This section describes the following individual system settings:

- ▶ 5.3.1, “System power settings”
- ▶ 5.3.2, “Processor settings” on page 171
- ▶ 5.3.3, “Memory settings” on page 173
- ▶ 5.3.4, “UEFI settings for different workload types” on page 174
- ▶ 5.3.5, “ServeRAID M5210 RAID controller configuration” on page 176
- ▶ 5.3.6, “Speeding up the boot process” on page 181

### 5.3.1 System power settings

Figure 5-15 shows the available options in the UEFI system Power settings.

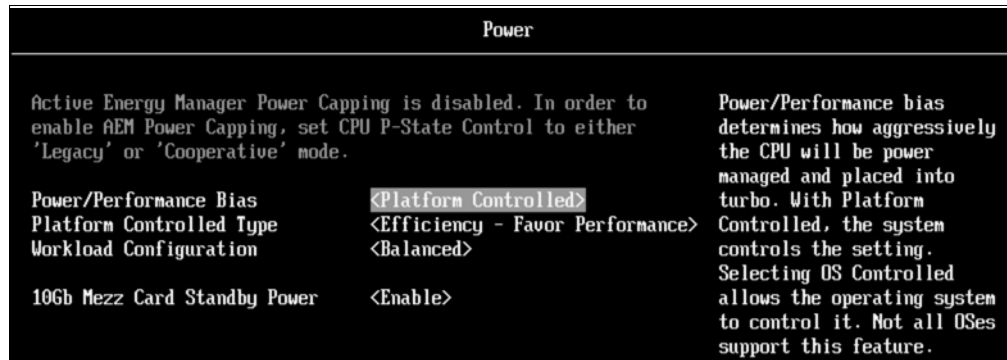


Figure 5-15 UEFI Power settings panel

Power settings include the following options:

- ▶ IBM Systems Director Active Energy Manager (Default: Capping Disabled)

The Active Energy Manager option enables the server to use the power capping feature of Active Energy Manager, which is an extension of IBM Systems Director.

Active Energy Manager measures, monitors, and manages the energy and thermal components of Lenovo systems, which enables a cross-platform management solution and simplifies the energy management of Lenovo servers, storage, and networking equipment. In addition, Active Energy Manager extends the scope of energy management to include non Lenovo systems, facility providers, facility management applications, protocol data units (PDUs), and equipment supporting the IPv6 protocol.

With Active Energy Manager, you can accurately understand the effect of the power and cooling infrastructure on servers, storage, and networking equipment. One of its features is to set caps on how much power the server can draw.

In order to enable Active Energy Manager Power Capping, set CPU P-State Control in Processors section to either **Legacy** or **Cooperative mode**.

For more information about IBM Systems Director Active Energy Manager, see the following web page:

<http://ibm.com/systems/director/aem/>

- ▶ Power/Performance Bias

This option determines how aggressively the CPU is power managed and places into turbo. This option includes the following choices:

- Platform Controlled: The system controls the settings.
- OS Controlled: The operating system controls it.

You can change the Power/Performance Bias option in the Operating Modes menu when you select **Custom Mode**. For more information about operating modes, see 5.2.1, “Operating modes” on page 168.

- ▶ Platform Controlled Type

This option refers to the operating mode that your system is set to run in. For more information about operating modes, see 5.2.1, “Operating modes” on page 168.

► Workload Configuration

The default selection for workload configuration is Balanced. You can change the workload configuration to I/O Sensitive. However, I/O Sensitive is used with expansion cards that require high I/O bandwidth when the CPU cores are idle to allow enough frequency for the workload.

Choose 'I/O sensitive' mode for I/O throughput, bandwidth or latency sensitive workloads.

► 10Gb Mezz Card Standby Power (Default: Enabled)

When this option is enabled, system fans keep running when the server is off. If disabled, the system saves more power when off, but loses network connection on a 10Gb Mezz card, which affects Wake on LAN and shared System Management Ethernet port functions.

It is recommended to disable this option, if you need to have maximum performance.

### 5.3.2 Processor settings

Processor settings control the various performance and power features that are available on the installed Xeon processor.

Figure 5-16 shows the UEFI Processor system settings window with the default values.

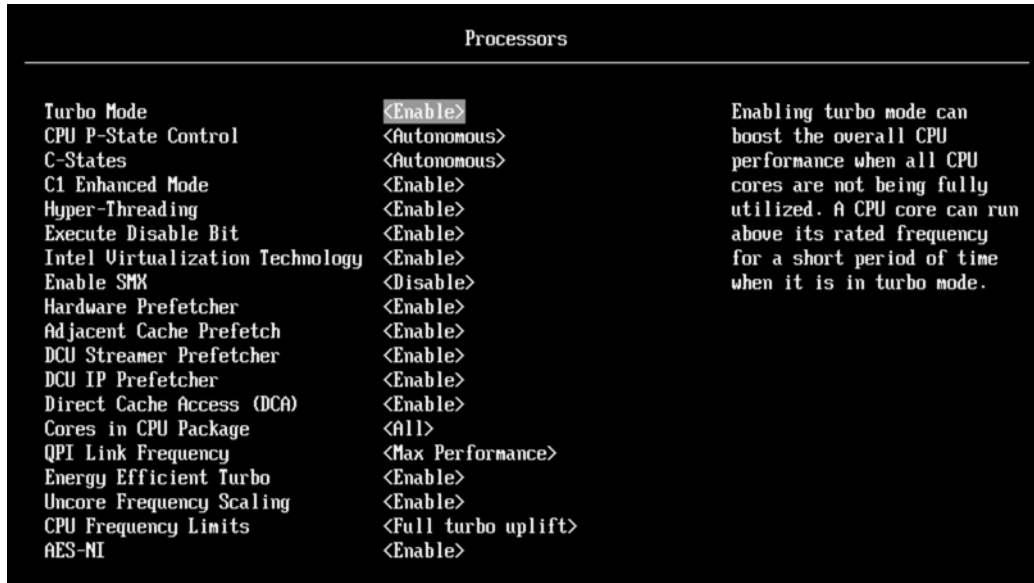


Figure 5-16 UEFI Processor system settings panel

The following processor feature options are available:

► Turbo Mode (Default: Enabled)

This mode enables the processor to increase its clock speed dynamically if the CPU does not exceed the Thermal Design Power (TDP) for which it was designed. For more information about Intel Turbo Boost, see "Turbo Boost Technology 2.0" on page 28.

► CPU P-State Control (Default: Autonomous)

This parameter defines the method used to control CPU performance states. Available options:

- None: All P-states are disabled and the CPUs run at either their rated frequency or in turbo mode (if turbo is enabled).

- Legacy: The CPU P-states will be presented to the OS and the OS power management will directly control which P-state is selected.
  - Autonomous: The P-states are controlled fully by system hardware. No P-state support is required in the OS or VM.
  - Cooperative: This mode is a combination of Legacy and Autonomous modes. The P-states are still controlled in hardware, but the OS can provide hints to the hardware for P-state limits and the desired settings.
- ▶ C-States (Default: Autonomous)
 

This option enables dynamic processor frequency and voltage changes in the idle state, which provides potentially better power savings. Available options:

    - Disabled: All C-states are disabled.
    - Legacy: The OS initiates the C-state transitions by using ACPI C-state tables.
    - Autonomous: ACPI C-state tables are not used and CPU HALT and C1 idle requests are automatically converted to CPU C6 idle requests in hardware.
  - ▶ C1 Enhanced Mode (Default: Enabled)
 

This option enables processor cores to enter an enhanced halt state to lower the voltage requirement and provides better power savings.
  - ▶ Hyper-Threading (Default: Enabled)
 

This option enables logical multithreading in the processor so that the operating system can run two threads simultaneously for each physical core. For more information, see “Hyper-Threading Technology” on page 28.
  - ▶ Execute Disable Bit (Default: Enabled)
 

This option enables the processor to disable the running of certain memory areas, which prevents buffer overflow attacks.
  - ▶ Intel Virtualization Technology (Default: Enable)
 

This option enables the processor hardware acceleration feature for virtualization. For more information, see “Intel Virtualization Technology” on page 28.
  - ▶ Enable SMX (Default: Disabled)
 

This option enables Safer Mode Extensions.
  - ▶ Hardware Prefetcher (Default: Enabled)
 

This option enables Hardware Prefetcher. Lightly threaded applications and some benchmarks can benefit from having the Hardware Prefetcher enabled.
  - ▶ Adjacent Cache Prefetcher (Default: Enabled)
 

This option enables Adjacent Cache Prefetcher. Lightly threaded applications and some benchmarks can benefit from having the Adjacent Cache Prefetcher enabled.
  - ▶ DCU Streamer Prefetcher (Default: Enabled)
 

This option enables DCU Streamer Prefetcher. Lightly threaded applications and some benchmarks can benefit from having the DCU Streamer Prefetcher enabled.
  - ▶ DCU IP Prefetcher (Default: Enabled)
 

This option enables the DCU IP Prefetcher. This option often is best left enabled for most environments. Some environments can benefit from having it disabled (for example, Java).

- ▶ Cores in CPU Package (Default: All)
 

This option sets the number of processor cores to be activated within each CPU package. You might want to change your CPU cores to lower your power consumption or to meet software licensing requirements.
- ▶ QPI Link Frequency (Default: Max Performance)
 

This option sets the operating frequency of the processor's QPI link:

  - Minimal Power provides less performance for better power savings.
  - Power Efficiency provides best performance per watt ratio.
  - Max performance provides the best system performance.
- ▶ Energy Efficient Turbo (Default: Enabled)
 

When Energy Efficient Turbo is enabled, the CPU's optimal turbo frequency is tuned dynamically based on the CPU usage. The power and performance bias setting also influences Energy Efficient Turbo.
- ▶ Uncore Frequency scaling (Default: Enabled)
 

When enabled, the CPU uncore dynamically changes speed based on the workload. All miscellaneous logic inside the CPU package is considered the uncore.
- ▶ CPU Frequency Limits (Default: Full turbo uplift)
 

The maximum Turbo Boost frequency can be restricted with turbo limiting to a frequency that is between the maximum turbo frequency and the rated frequency for the CPU installed. This can be useful for synchronizing CPU tasks. Note that the maximum turbo frequency for N+1 cores cannot be higher than for N cores. If the turbo limits are being controlled through application software, leave this option at the default value.

Available options:

  - Full turbo uplift
  - Restrict maximum frequency.
- ▶ AES-NI (Default: Enabled)
 

This option enables the support of Advanced Encryption Instructions to improve the encryption performance.

### 5.3.3 Memory settings

The Memory settings panel provides the available memory operation options, as shown in Figure 5-17.



Figure 5-17 UEFI Memory system settings panel

The memory settings panel includes the following options:

▶ **Memory Mode (Default: Lockstep)**

This option enables Independent or Lockstep Memory Mode. In independent mode, all channels operate independently with SMI links that are operating at twice the DDR3 speed. In Lockstep mode, two channels operate synchronously, which allows advanced memory protection. SMI link operations occur at the DDR3 speed.

▶ **Memory Speed (Default: Max Performance)**

This option sets the following operating frequency of the installed DIMMs:

- Minimal Power provides less performance for better power savings. The memory operates at the lowest supported frequency, which, is 1333 MHz in the X6 system.
- Power Efficiency provides the best performance per watt ratio. The memory operates one step under the rated frequency.

Max Performance provides the best system performance. The memory operates at the rated frequency; that is, 1600 MHz for DIMMs rated at 1600 MHz or higher.

▶ **Socket Interleave (Default: NUMA)**

This option sets Socket Interleave to Non Unified Memory Architecture (NUMA) or Non-NUMA. In NUMA mode, memory is not interleaved across processors. In Non-NUMA, memory is interleaved across processors.

▶ **Patrol Scrub (Default: Enable)**

This option enables scheduled background memory scrubbing before any error is reported, instead of default demand scrubbing on an error event. This option provides better memory subsystem resiliency at the expense of a small performance loss.

▶ **Memory Data Scrambling (Default: Enable)**

This option enables a memory data scrambling feature to further minimize bit-data errors.

▶ **Mirroring (Default: Disabled)**

This option enables or disables memory mirroring, as described in “Memory mirroring” on page 38.

**Memory mirror mode:** Memory mirror mode cannot be used with Memory spare mode.

▶ **Sparing (Default: Disabled)**

This option enables or disables memory rank sparing, as described in “Rank sparing” on page 40.

▶ **Memory Power Management (Default: Disabled):**

This option enables or disables Memory Power Management.

### 5.3.4 UEFI settings for different workload types

Tuning the X6 servers for performance is a complicated topic because it depends on which application you installed or which workload this application generates. For different types of workload you need to find a balance between performance, reliability, and power efficiency. For example, for database workload high performance and reliability are essential requirements. For virtualization the main focus is on reliability and power efficiency. For web servers, power efficiency and performance are more important.



In this section, we provide general settings for the x3850 X6 that can be a good starting point for further tuning. Table 5-2 lists the UEFI settings that are recommended for specific workload requirements.

Table 5-2 Overview of UEFI settings for specific workload

UEFI Setting	Web server	Virtualization <sup>a</sup>	Low latency	Database	HPC
TurboMode <sup>b</sup>	Enabled	Enabled	Enabled	Enabled	Enabled
CPU P-States	Enabled	Enabled	Disabled	Enabled	Enabled
C-states	Enabled	Enabled	Disabled	Enabled	Disabled
C1 Enhanced Mode	Enabled	Enabled	Disabled	Enabled	Disabled
Prefetcher	Enabled	Enabled	Enabled	Enabled	Enabled
Hyper-Threading	Enabled	Enabled	Disabled	Enabled	Disabled
Virtualization extensions	Disabled	Enabled	Disabled	Disabled	Disabled
QPI Link speed	Maximum performance	Maximum performance	Maximum performance	Maximum performance	Maximum performance
Power Workload Configuration	Balanced	Balanced	I/O Sensitive	I/O Sensitive	I/O Sensitive
Memory Mode	Independent	Lockstep	Independent	Lockstep	Independent
Memory speed	Maximum performance	Maximum performance	Maximum performance	Maximum performance	Maximum performance
Patrol scrub	Enabled	Enabled	Disabled	Enabled	Disabled

- a. These Virtualization settings are advised for a stand-alone host only. For multiple virtualized hosts in clustered workloads, use the Maximum performance settings instead.
- b. Depending on the processor workload, enabling TurboMode might also increase power consumption. The actual processing performance boost that you get from TurboMode depends on the environment of the server, in terms of temperature and humidity, because the processor boosts performance only up to the environmental limits set by the processor.

For low latency, HPC workloads, and other applications where you need maximum performance, make sure that active power policy has disabled throttling and enabled N+N power redundancy. You should also disable power capping to avoid performance degradation. For more information about power policies, see 4.5.5, “Power policy” on page 141.

To avoid installation problems using VMware 5.1 and later, we recommend you disable PCI 64-Bit Resource allocation and set parameter MMConfigBase to 3 GB. You can set these UEFI parameters using ASU utility, as shown in Figure 5-18.

```
# asu64 set DevicesandIOPorts.PCI64-BitResourceAllocation Disable
# asu64 set DevicesandIOPorts.MMConfigBase "3GB"
```

Figure 5-18 VMware specific UEFI settings

You can also use UEFI setting presets like Maximum performance or Minimum Power, which we cover in 5.2.1, “Operating modes” on page 168.

### 5.3.5 ServeRAID M5210 RAID controller configuration

The UEFI provides access to the LSI MegaRAID Configuration Utility as used by the ServeRAID M5210 SAS/SATA Controller. Access the menu in UEFI by selecting **System Settings** → **Storage** → **LSI MegaRAID <ServeRAID M5210> Configuration Utility**, as shown in Figure 5-19.




Figure 5-19 LSI MegaRAID Utility panel for ServeRAID M5210

The LSI MegaRAID Configuration Utility has several options for managing your internal disks. Figure 5-20 shows the options that are available in the LSI MegaRAID Configuration Utility panel.



Figure 5-20 LSI MegaRAID Configuration Utility, Main Menu panel

This menu features the following settings:

- ▶ **Configuration Management**

Displays configuration options. Some options appear only if the controller supports them. For example, create virtual drive, create CacheCade virtual drive, make JBOD, make Unconfigured Good, clear configuration, manage foreign configuration, view drive group properties, and view global hot spare drives.
- ▶ **Controller Management**

Displays the controller status and basic properties of the controller, such as product name, serial number, PCI ID, firmware version, NVDATA Version, set factory defaults, and clear configuration. You can also use the advanced link to view more properties and perform other tasks, such as changing the security key and saving the TTY log.
- ▶ **Virtual Drive Management**

Manages the virtual drive properties and enables you to view the basic virtual drive properties and perform operations, such as background initialization and consistency check. You can also view other properties by using the Advanced link.

- ▶ Drive Management
 

Displays the basic drive properties and performs operations, such as assign or unassign a hot spare drive, locate drives, place drive offline or online, and rebuild a drive. You can also view other properties by using the Advanced link.
- ▶ Hardware Components
 

Displays the battery status and the status of temperature sensors, fans, and power supplies. You can also view more properties and perform other operations by using the Advanced link. Some options appear only if the controller supports them.

The UEFI supports the configuration of your RAID array with a supported controller. Complete the following steps to create a RAID-1 array that is configured in the UEFI:

1. Enter UEFI setup by selecting F1 when prompted at start time. Access the LSI Configuration Utility from the menu by selecting **System Settings** → **Storage** → **LSI MegaRAID <ServeRAID M5210> Configuration Utility**, as shown in Figure 5-21.




Figure 5-21 LSI MegaRAID Utility panel for ServeRAID M5210

2. Figure 5-22 shows the Main Menu for the Configuration Utility. Enter the Configuration Manager by selecting **Configuration Management**.



Figure 5-22 LSI MegaRAID Configuration Utility, Main Menu panel

3. Access the Controller Management panel, as shown in Figure 5-23.

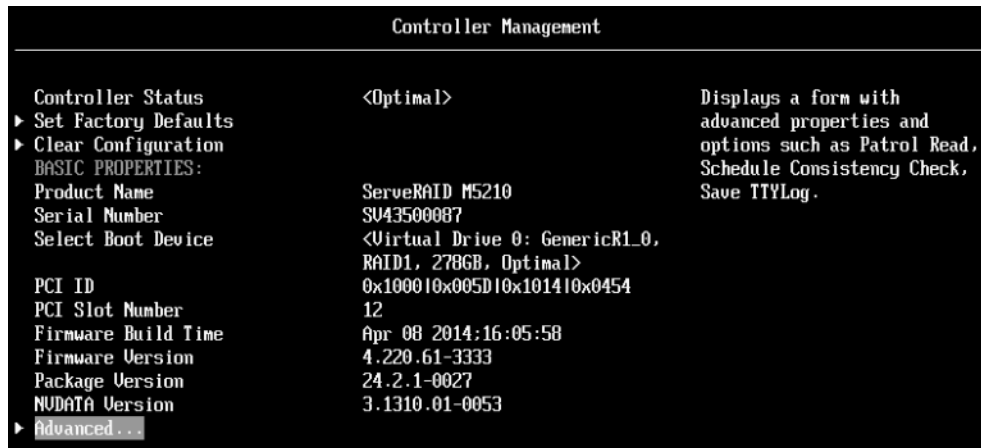


Figure 5-23 Controller Management Main Menu panel

4. From the Controller Management main menu panel, select **Advanced**, as seen in Figure 5-24.

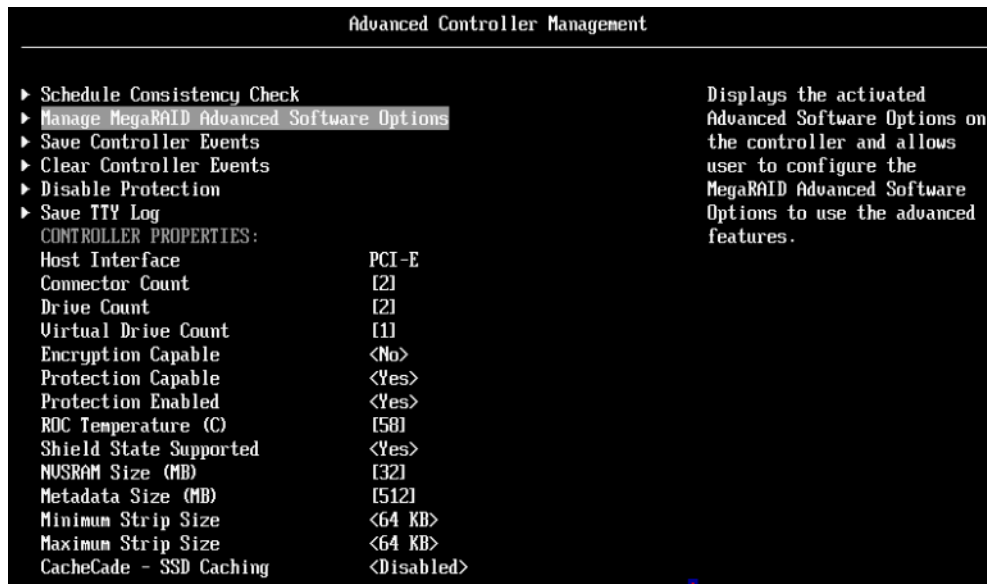


Figure 5-24 Advanced Controller Management menu panel

5. From the Advanced Controller Management menu, select **Manage MegaRAID Advanced Software Options**, as shown in Figure 5-25.



Figure 5-25 Manage MegaRAID Advanced Software Options

6. Ensure that the MegaRAID Advanced Software Option is activated. Exit this menu by pressing the Esc key and return to the Main Menu, as shown in Figure 5-22 on page 177.
7. From this menu, select **Drive Management**. Ensure that your disks are in an *Unconfigured Good* state, as shown in Figure 5-26.



Figure 5-26 Drive Management panel

8. Select each drive individually. From here, you can find a drive and initialize a drive from the menu, as shown in Figure 5-27.

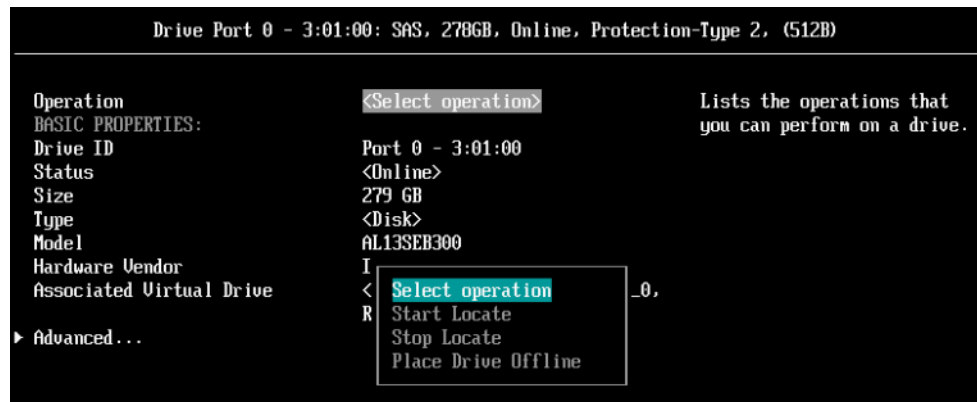


Figure 5-27 Drive operation selection

9. Return to the Main Menu, as shown in Figure 5-22 on page 177. From here, select **Configuration Management** panel (see Figure 5-28).



Figure 5-28 Configuration Management panel

10. From the Configuration Management panel, select **Create Virtual Drive**, as shown in Figure 5-29.



Figure 5-29 Creating a Virtual Drive

11. Select the wanted RAID level for your drives. For this example, we selected RAID 1, as shown in Figure 5-30.

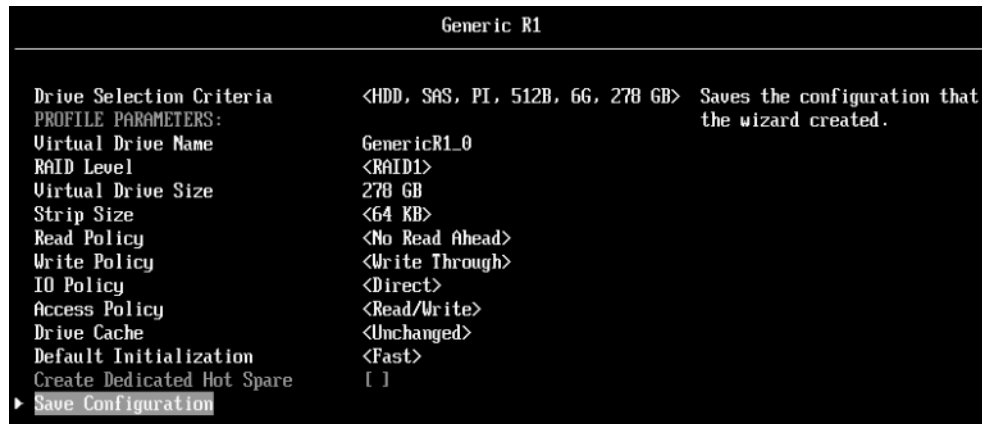


Figure 5-30 RAID 1 array panel

12. Select **Save Configuration** and confirm that you want to create your RAID array, as shown in Figure 5-31.

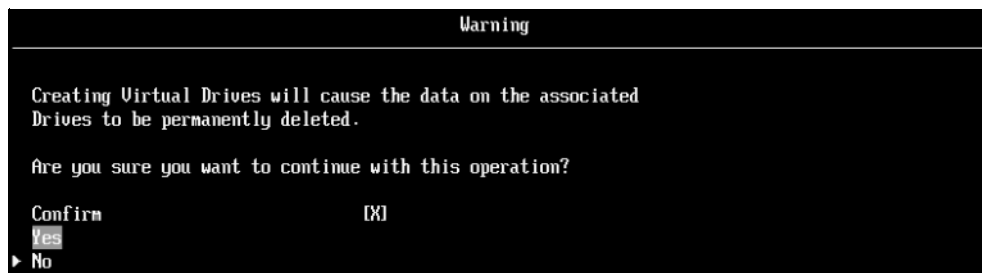


Figure 5-31 Warning and confirmation of RAID array creation

13. Select **Confirm** to mark an **X** in the box and then select **Yes**. After you confirm the array that you created, you see a Success message, as shown in Figure 5-32.

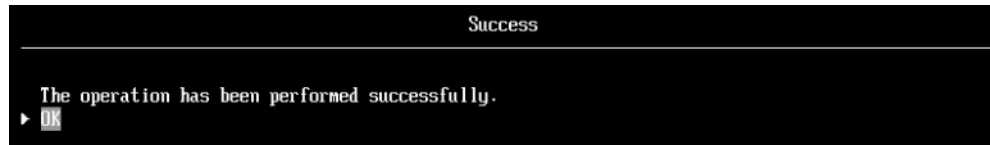


Figure 5-32 Successful RAID array creation

14. Select **OK**, and press Esc to return to the main menu. Your array is now active.

There are more utilities for configuring and managing your disks in the x3850 X6 and x3950 X6 server. For more information, see 6.2, “ServerGuide” on page 233 and 7.9, “MegaRAID Storage Manager” on page 292.

The IMM2 displays the RAID configuration in the Local Storage pane, as shown in Figure 5-33.

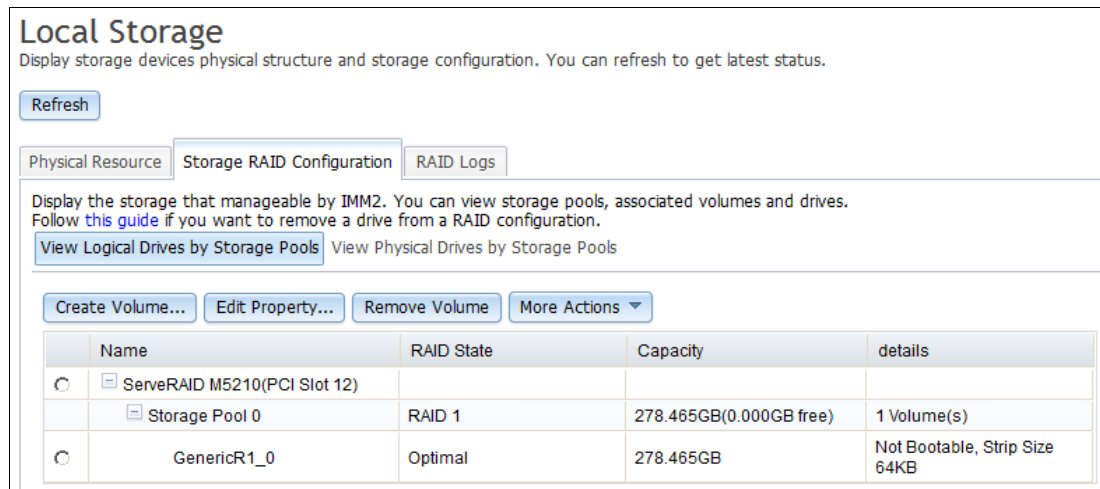


Figure 5-33 IMM Local storage RAID information

### 5.3.6 Speeding up the boot process

In this section we describe some common UEFI options that allow you to reduce the time needed for the server restart. During system restart the server passes several stages including hardware check, memory initialization, and booting devices initialization. The following considerations may help you to optimize your server configuration in order to speed up the booting process:

- ▶ To reduce the time needed for boot device initialization, remove from the boot order any devices that will not be used. (Every device that must be polled for a boot record increases the boot time. )

For example, remove boot options for Floppy disk, CD/DVD ROM, and PXE Network if you don't use them, as shown in Figure 5-34 on page 182.

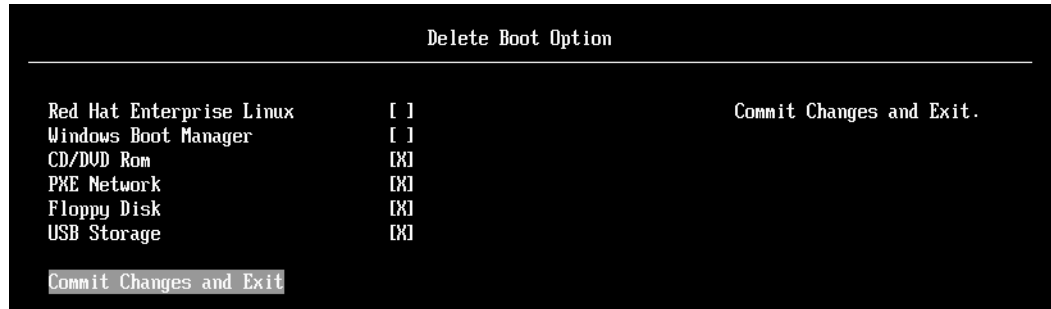


Figure 5-34 Delete Boot Option section

- Disable Legacy Option ROMs on devices not used as boot targets

Go to **System Settings** → **Devices and I/O Ports** → **Enable / Disable Adapter Option ROM Support** and disable ROM options for Legacy devices, as shown in Figure 5-35 on page 182:

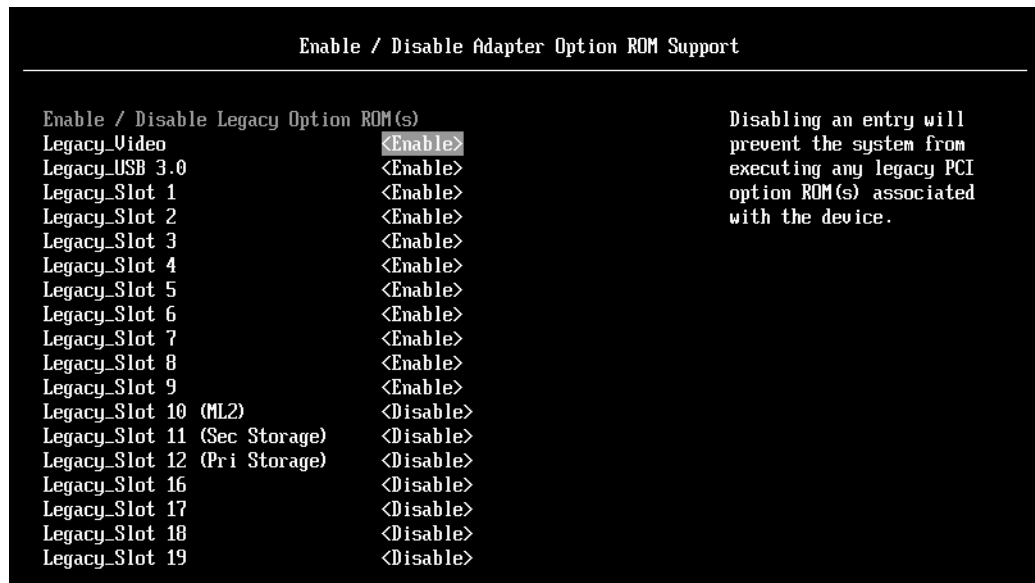


Figure 5-35 Disable Legacy Option ROM

- For PXE, disable PXE-boot on ports that will not be used for booting

Enable PXE-boot only on the appropriate network port, which is used for network booting. Avoid wasted time while all available network ports try to reach DHCP and TFTP servers. Go to **System Settings** → **Network** → **Network Boot Configuration**, choose unused network ports and disable PXE-boot, as shown in Figure 5-36:





Figure 5-36 Disable PXE-boot

- The more memory you have, the longer the system takes in the “System initializing memory” state.

You cannot reduce the amount of time needed for memory initializing, but you can temporarily enable Quick boot mode. In this mode the server initializes only a minimum amount of memory (only one or two DIMMs depending on the number of Compute Books installed), which allows you to reboot the system relatively quickly. Quick boot mode may be useful for any maintenance or OS deployment, when you need to install an OS, or boot a rescue image or a BOMC image for firmware updates.

**Important:** You should not use Quick boot mode for production system.

To enable the Quick boot mode, you can use the Front operator panel, located on the Storage book. Go to the **Actions** section and choose **Quick boot mode**.

You can also use ASU utility to enable or disable the Quick boot mode. Use UEFI option IMM.DeploymentBoot with the following parameters:

- **Disable** to disable the Quick boot mode
- **NextBoot** to enable it during the next system reboot
- **NextAC** to enable this mode on the next power cycle.

Figure 5-37 shows the ASU usage examples on Linux:

```

[root@localhost ~]# /opt/lenovo/toolscenter/asu/asu64 set IMM.DeploymentBoot NextBoot
Lenovo Advanced Settings Utility version 10.1.90F
Licensed Materials - Property of Lenovo
(C) Copyright Lenovo Corp. 2007-2015 All Rights Reserved
Connected via IPMI device driver (KCS interface)
IMM.DeploymentBoot=NextBoot
Waiting for command completion status.
Command completed successfully.
[root@localhost ~]# /opt/lenovo/toolscenter/asu/asu64 set IMM.DeploymentBoot NextAC
Lenovo Advanced Settings Utility version 10.1.90F
Licensed Materials - Property of Lenovo
(C) Copyright Lenovo Corp. 2007-2015 All Rights Reserved
Connected via IPMI device driver (KCS interface)
IMM.DeploymentBoot=NextAC
Waiting for command completion status.
Command completed successfully.
[root@localhost ~]# /opt/lenovo/toolscenter/asu/asu64 set IMM.DeploymentBoot Disabled
Lenovo Advanced Settings Utility version 10.1.90F
Licensed Materials - Property of Lenovo
(C) Copyright Lenovo Corp. 2007-2015 All Rights Reserved
Connected via IPMI device driver (KCS interface)
IMM.DeploymentBoot=Disabled
Waiting for command completion status.
Command completed successfully.
[root@localhost ~]#

```

Figure 5-37 Enable Quick boot mode using ASU

## 5.4 Use of NVMe drives and adapters

The x3850 X6 and x3950 X6 support PCIe Non-Volatile Memory Express (NVMe) drives and adapters that connect directly to the PCIe bus of the processors rather than via a SAS controller.

For more information about the technology, see 2.6.3, “NVMe SSD technology” on page 45. For more information about the components, see 3.11, “Storage subsystem” on page 94.

In this section, we describe the planning and use of these drives. This section includes the following topics:

- ▶ 5.4.1, “NVMe drive placement”
- ▶ 5.4.2, “NVMe PCIe SSD adapter placement” on page 188
- ▶ 5.4.3, “Using NVMe drives with Linux” on page 188
- ▶ 5.4.4, “Using NVMe drives with Microsoft Windows Server 2012 R2” on page 192
- ▶ 5.4.5, “Using NVMe drives with VMware ESXi server” on page 196

### 5.4.1 NVMe drive placement

The following rules apply for installing PCIe NVMe drives in the server:

- ▶ One NVMe backplane is supported in the x3850 X6 (as listed in Table 3-13 on page 96).
- ▶ Two NVMe backplanes are supported in the x3950 X6: one in the lower Storage Book and one in the upper Storage Book.
- ▶ The NVMe backplane may be only installed in the top backplane bay of the Storage book.

- ▶ You can install a SAS/SATA backplane in the lower backplane in the Storage Book.
- ▶ One backplane supports up to four NVMe drives.
- ▶ One Extender Adapter is required for every two NVMe drives installed (as listed in Table 3-15 on page 98).
- ▶ One NVMe cable is required between the adapter and backplane per drive installed.
- ▶ For 3 or 4 NVMe drives that are installed in the server, a second Extender Adapter must be installed. The Extender Adapters are installed in the following slots of the Storage Books:
  - For the x3850 X6, the slots are PCIe slots 11 and 12.
  - For the x3950 X6, the slots are PCIe slots 11, 12, 43, and 44.
- ▶ If four NVMe drives are installed in a Storage Book, two Extender Adapters are required to connect those drives. Because these two adapters occupy both PCIe slots in the Storage Book, no other controllers can be installed; therefore, the other four drive bays in the Storage Book must remain empty.

Table 5-3 shows the connectivity and slot installation ordering of the PCIe Extender Adapter, backplane, and solid-state drives (SSDs).

*Table 5-3 x3850 X6 NVMe slot and PCIe installation ordering*

<b>NVMe PCIe Extender Adapter</b>	<b>Extender Adapter location</b>	<b>PCIe signal cable connections</b>	<b>NVMe SSD drives population order, location in the Storage book, PCIe slot used by drive</b>
NVMe PCIe extender adapter 1	I/O book slot 11	Adapter port 0 to backplane port 3	Drive 1, bay 7, PCIe slot 19
		Adapter port 1 to backplane port 2	Drive 2, bay 6, PCIe slot 18
NVMe PCIe extender adapter 2	I/O book slot 12	Adapter port 0 to backplane port 1	Drive 3, bay 5, PCIe slot 17
		Adapter port 1 to backplane port 0	Drive 4, bay 4, PCIe slot 16

Figure 5-38 shows 2.5-inch NVMe SSD drives location in the Storage book, bays and PCIe slots used by them:




Figure 5-38 2.5-inch NVMe SSD drives location in the Storage Book

Figure 5-39 shows the ports of the Extender Adapter.




Figure 5-39 NVMe Extender Adapter port numbering

Figure 5-40 shows the ports of the NVMe SSD backplane (drive side of the backplane, as seen from the front of the server)




Figure 5-40 NVMe backplane port numbering (drive side of the backplane)

The operating system and UEFI report the NVMe drives attached to the 4x2.5-inch NVMe PCIe backplane as PCI devices, connected to PCIe slots 16-19. You can check connected NVMe SSD drives from the IMM2 web-interface at **Server Management** → **Adapters** page, as shown in Figure 5-41 on page 187:

**Adapters**

Display Adapters information. Click the link of each device to view more details. If you remove or replace adapters, the server needs to be powered on at least once after the removal/replacement to show the correct adapters information. Go to [Server Firmware page](#) if you need update firmware for any adapter.

Slot No.	Device Name	Device Type	Card Interface
OnBoard	<a href="#">Adapter 1B:00:00</a>	GPU	Onboard
3	<a href="#">Emulex VFA5 2x10 GbE SFP+ PCIe Adapter for IBM System x</a>		FlexSystem Mezzanine Connector
	... <a href="#">Emulex VFA5 2x10 GbE SFP+ PCIe Adapter for IBM System x C1:00:00</a>	Ethernet	
	... <a href="#">Emulex VFA5 2x10 GbE SFP+ PCIe Adapter for IBM System x C1:00:01</a>	Ethernet	
7	<a href="#">P3700 1.6TB NVMe Enterprise Performance Flash Adapter</a>	NVMe	PCI-E x4
8	<a href="#">P3700 1.6TB NVMe Enterprise Performance Flash Adapter</a>	NVMe	PCI-E x4
10	<a href="#">Adapter 01:00:01</a>		PCI-E Gen 3
	... <a href="#">Function 01:00:00</a>	Ethernet	
	... <a href="#">Function 01:00:01</a>	Ethernet	
12	<a href="#">ServeRAID M5210</a>	RAID	Unknown
18	<a href="#">P3700 1.6TB NVMe 2.5" Enterprise Performance PCIe SSD</a>	NVMe	PCI-E x4
19	<a href="#">P3700 1.6TB NVMe 2.5" Enterprise Performance PCIe SSD</a>	NVMe	PCI-E x4

Figure 5-41 PCIe slots used by NVMe SSD drives

**Know your PCIe slot numbers:** It's important to know the PCIe slots numbers used by NVMe drives: during the software RAID maintenance and NVMe SSD drives replacement, these PCIe slot numbers allows you to distinguish the appropriate drive in the set of similar NVMe drives.

For more information about NVMe drives, see 2.6.3, "NVMe SSD technology" on page 45.

## 5.4.2 NVMe PCIe SSD adapter placement

X6 servers also support P3700 NVMe Enterprise Performance Flash Adapters, which have a half-height, half-length, single-slot x4 PCIe adapter form-factor. x3850 X6 servers support up to nine NVMe PCIe SSD adapters, and x3950 X6 servers support up to 18 adapters.

To have the best performance using NVMe PCIe SSD adapters, follow the next PCIe placement rules as shown in Table 5-4 on page 188:

Table 5-4 NVMe PCIe SSD adapters population order

Server type	NVMe PCIe SSD adapters population order (PCIe slots)
x3850 X6	8, 7, 9, 4, 1, 5, 2, 6, 3
x3950 X6	8, 40, 7, 39, 9, 41, 4, 36, 1, 33, 5, 37, 2, 34, 6, 38, 3, 35

For more information about PCIe adapters placement rules, see 5.5, “PCIe adapter placement advice” on page 198.

## 5.4.3 Using NVMe drives with Linux

NVMe drives are supported on the following Linux distributions:

- ▶ Red Hat Enterprise Linux 6.5 and later
- ▶ Red Hat Enterprise Linux 7.0 and later
- ▶ SUSE Linux Enterprise Server 11 SP3
- ▶ SUSE Linux Enterprise Server 12

Other Linux distributions might have NVMe support, depending on the kernel version.

The RHEL and SLES distributions have NVMe kernel modules; therefore, no other drivers are required to use NVMe drives. NVMe drives are represented in the OS as block devices with device names, such as `/dev/nvmeXn1`, where X is a number that is associated with each NVMe drive that is installed in the server. For example, for two NVMe drives installed, the device names are `nvme0n1` and `nvme1n1` files in `/dev` directory, as shown in Figure 5-42.

```
[root@localhost ~]# ls /dev/nvme
nvme0  nvme0n1  nvme1  nvme1n1
```

Figure 5-42 NVMe drives device names in /dev directory

Device names `/dev/nvmeX` represent associated controllers.

Figure 5-43 shows other Linux commands that can show the NVMe drives.

```
[root@localhost ~]# lspci | egrep -i "ssd|953"
89:00.0 Non-Volatile memory controller: Intel Corporation PCIe Data Center SSD
(rev 01)
8e:00.0 Non-Volatile memory controller: Intel Corporation PCIe Data Center SSD
(rev 01)
[root@localhost ~]# lsblk
NAME        MAJ:MIN RM   SIZE RO TYPE MOUNTPOINT
loop0       7:0    0  53.8M 0 loop /media
sda         8:0    0 278.5G 0 disk
..sda1      8:1    0   50M 0 part /boot/efi
..sda2      8:2    0  256M 0 part /boot
..sda3      8:3    0    4G 0 part [SWAP]
..sda4      8:4    0 274.2G 0 part /
nvme0n1     252:0  0 372.6G 0 disk
nvme1n1     252:64 0 372.6G 0 disk
```

Figure 5-43 *lspci* and *lsblk* output

As shown in Figure 5-43, *lspci* and *lsblk* commands show that two NVMe controllers are connected to PCIe bus and two block devices *nvme0n1* and *nvme1n1* are available in the operating system. You can also run simple performance tests by using the *hdparm* utility, as shown in Figure 5-44.

```
[root@localhost ~]# hdparm -tT --direct /dev/nvme0n1

/dev/nvme0n1:
Timing O_DIRECT cached reads: 4594 MB in 2.00 seconds = 2298.38 MB/sec
Timing O_DIRECT disk reads: 8314 MB in 3.00 seconds = 2770.65 MB/sec
```

Figure 5-44 Performance test by using *hdparm* utility

As shown in Figure 5-44, direct read speed from one NVMe drive is 2.7 GBps and cached read speed is almost 2.3 GBps.

You can work with NVMe drives as with other block devices, such as SATA or SAS drives. You can use *fdisk* or *parted* utilities to manage disk partitions, create any supported file systems by using standard Linux commands, and mount these file systems.

For example, you can create partition with the parted utility, as shown in Figure 5-45.

```
[root@localhost ~]# parted /dev/nvme0n1
GNU Parted 2.1
Using /dev/nvme0n1
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) print
Model: Unknown (unknown)
Disk /dev/nvme0n1: 400GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt

Number  Start  End  Size  File system  Name  Flags

(parted) mkpart primary ext4 1M 400GB
(parted) print
Model: Unknown (unknown)
Disk /dev/nvme0n1: 400GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt

Number  Start  End  Size  File system  Name  Flags
1       1049kB 400GB 400GB                primary

(parted) quit
```

Figure 5-45 Partition creation with parted utility

When you create a partition on an NVMe drive, a new block device name appears in the /dev/ directory. For example, for /dev/nvme0n1 drive, /dev/nvme0n1p1 is created.



After that you can create the ext4 file system on that partition, as shown in Figure 5-46.

```
[root@localhost ~]# mkfs.ext4 /dev/nvme
nvme0      nvme0n1   nvme0n1p1 nvme1      nvme1n1
[root@localhost ~]# mkfs.ext4 /dev/nvme0n1p1
mke2fs 1.41.12 (17-May-2010)
Discarding device blocks: done
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
24420352 inodes, 97677568 blocks
4883878 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=4294967296
2981 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624, 11239424, 20480000, 23887872, 71663616, 78675968

Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 32 mounts or
180 days, whichever comes first. Use tune2fs -c or -i to override.
```

Figure 5-46 ext4 file system creation

Then, you can mount the new ext4 file system to the file system tree, as shown in Figure 5-47.

```
[root@localhost ~]# mkdir /media/nvme0
[root@localhost ~]# mount /dev/nvme0n1p1 /media/nvme0/
[root@localhost ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/sda4       270G  2.0G  255G   1% /
tmpfs           253G  224K  253G   1% /dev/shm
/dev/sda2       248M   39M  197M  17% /boot
/dev/sda1        50M  254K   50M   1% /boot/efi
/dev/nvme0n1p1  367G  195M  348G   1% /media/nvme0
```

Figure 5-47 File system mounting

You can manage other NVMe drives in the same way by using other Linux features, such as Logical Volume Manager (LVM) and software RAID, if needed.

## 5.4.4 Using NVMe drives with Microsoft Windows Server 2012 R2

NVMe drives are supported on following operation systems:

- ▶ Microsoft Windows Server 2008 R2
- ▶ Microsoft Windows Server 2012
- ▶ Microsoft Windows Server 2012 R2

Microsoft Windows Server 2012 R2 has native NVMe driver support and no other drivers are required to start use NVMe drivers. Other Windows version might require drivers.

Complete the following steps to check that NVMe drives are recognized by Windows:

1. Open Device Manager and the expand **Disk drives** section. All installed NVMe drives should present, as shown in Figure 5-48.

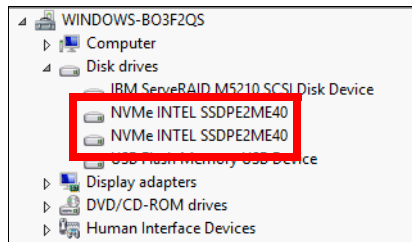


Figure 5-48 NVMe driver as shown in Device Manager

2. Open the Disk Management tool, you should see all installed NVMe drives. For example, both installed NVMe drives are presented as Disk 1 and Disk 2, as shown in Figure 5-49 on page 192.

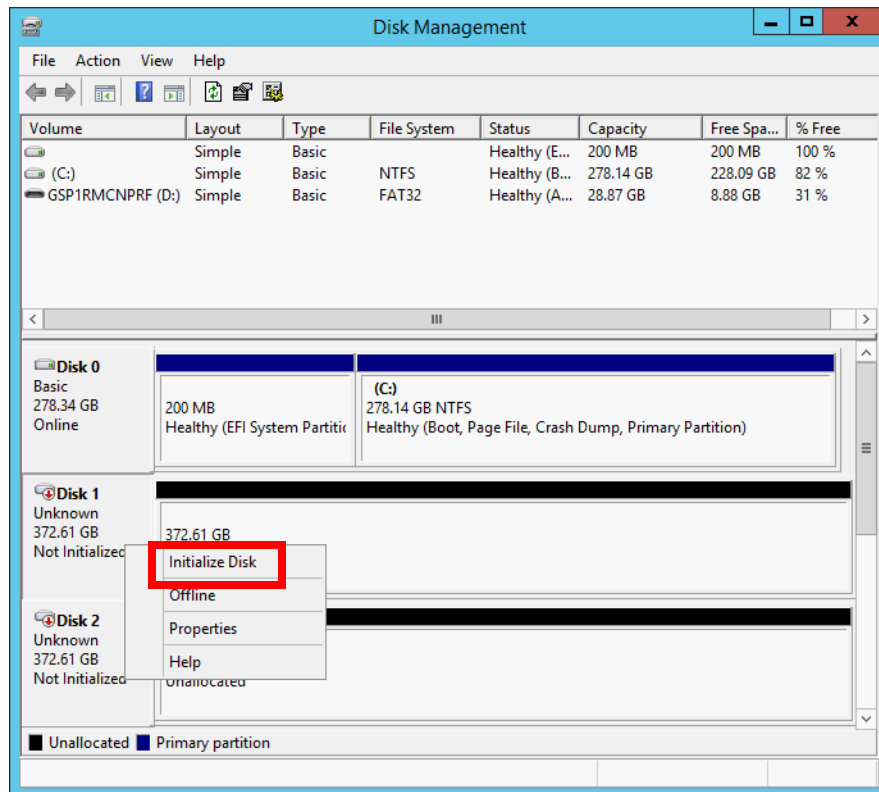


Figure 5-49 Disk Management tool and NVMe drives

- Both NVMe drives must be online and initialized. To initialize the drives, right-click the appropriate disk (Disk 1 or Disk 2 as shown in Figure 5-49) and select **Initialize Disk**. The Initialize Disk window opens, as shown in Figure 5-50.

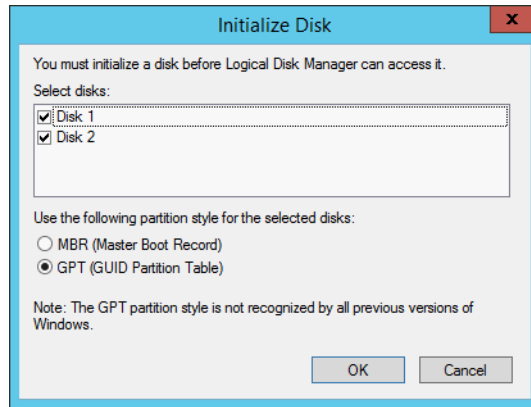


Figure 5-50 Disk initialization

- After the disks are initialized, you can create volumes. Right-click the NVMe drive and select the required volume to create, as shown in Figure 5-51 on page 193.

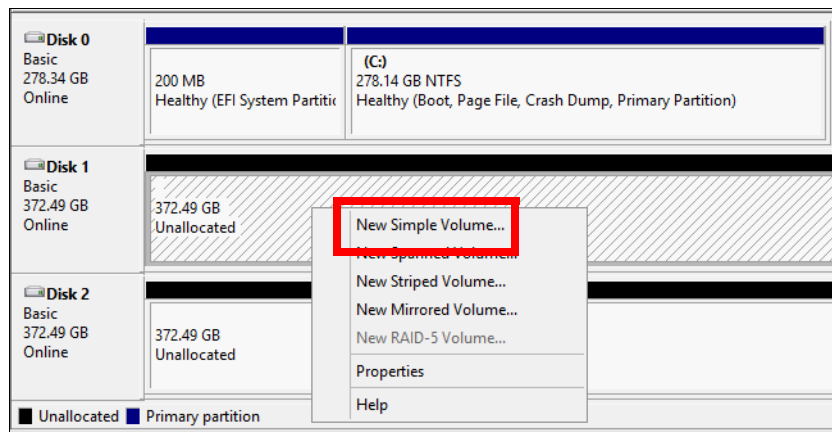


Figure 5-51 Creating a volume

- For example, choose **New Simple Volume**. The volume creation wizard opens. Click **Next** and specify a volume size, as shown in Figure 5-52.




Figure 5-52 Specify a volume size

6. You also must assign a drive letter or path for a new volume, as shown in Figure 5-53 on page 194.




Figure 5-53 Assign drive letter or path

7. You must format the new volume and specify the file system parameters, such as block size and volume label, as shown in Figure 5-54.




Figure 5-54 Format partition

8. Review all parameters and click **Finish**, as shown in Figure 5-55.




Figure 5-55 Completing the wizard

9. After the New Simple Volume wizard completes, you can see the new volume NVMe drive 1 (F:) by using the Disk Management tool, as shown in Figure 5-56.

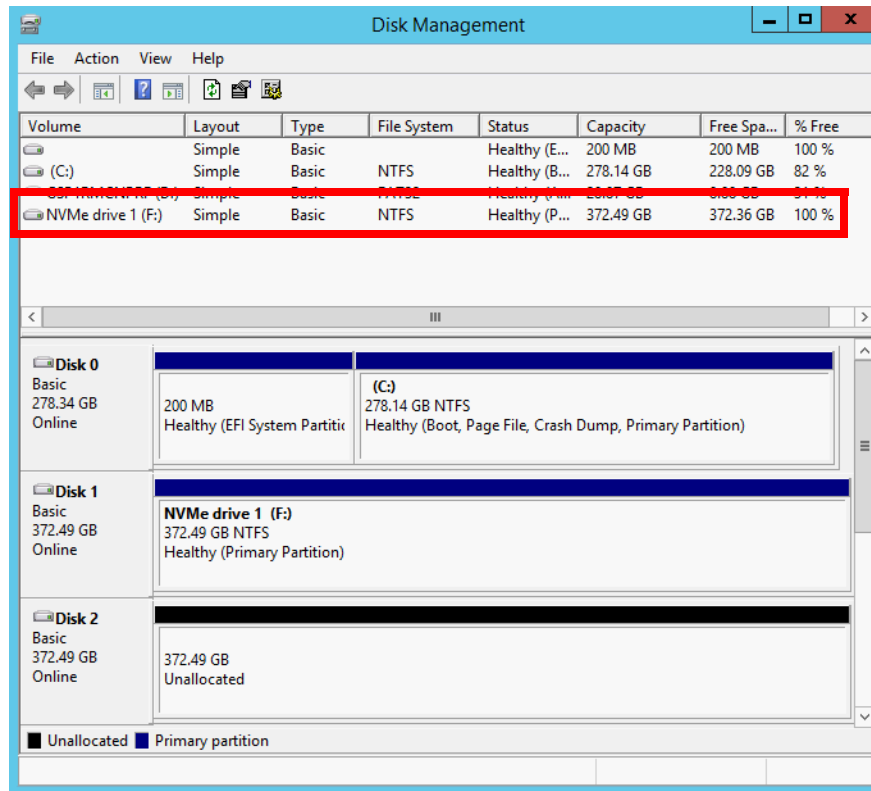


Figure 5-56 New NVMe volume

You now have NVMe drives that are available for storage. You can create software RAID arrays of different types by using two or more drives.

## 5.4.5 Using NVMe drives with VMware ESXi server

NVMe drives are supported by ESXi in the following configurations:

- ▶ VMware ESXi 6.0: Native support of NVMe drives
- ▶ VMware ESXi 5.5 Update 2: Requires another driver to install

The ESXi 5.5 driver for NVMe drives driver can be downloaded from the following VMware website:

<https://my.vmware.com/web/vmware/details?productId=353&downloadGroup=DT-ESXI55-VMWARE-NVME-10E030-1VMW>

Complete the following steps to install the NVMe driver on ESXi 5.5 Update 2:

1. Download VMware ESXi 5.5 NVMe driver from this website:

<https://my.vmware.com/web/vmware/details?productId=353&downloadGroup=DT-ESXI55-VMWARE-NVME-10E030-1VMW>

2. Enable SSH on the ESXi server.

3. Copy the VMware ESXi 5.5 NVMe driver to the ESXi server by using any SSH client, such as ISCP or WinSCP by using a command that is similar to the command that is shown in Figure 5-57.

```
scp VMW-ESX-5.5.0-nvme-1.0e.0.30-2284103.zip root@172.16.32.222:/tmp/
```

Figure 5-57 SCP usage for driver copying

4. Log in to the ESXi server by using SSH client and extracting the .zip file, as shown in Figure 5-58.

```
~ # cd /tmp
/tmp # unzip VMW-ESX-5.5.0-nvme-1.0e.0.30-2284103.zip
Archive:  VMW-ESX-5.5.0-nvme-1.0e.0.30-2284103.zip
  inflating: VMW-ESX-5.5.0-nvme-1.0e.0.30-offline_bundle-2284103.zip
  inflating: nvme-1.0e.0.30-1vmw.550.0.0.1391871.x86_64.vib
  inflating: doc/README.txt
  inflating: source/driver_source_nvme_1.0e.0.30-1vmw.550.0.0.1391871.tgz
  inflating:
doc/open_source_licenses_nvme_1.0e.0.30-1vmw.550.0.0.1391871.txt
  inflating: doc/release_note_nvme_1.0e.0.30-1vmw.550.0.0.1391871.txt
```

Figure 5-58 Extracting drivers from the archive

5. Install the extracted NVMe driver on the ESXi server, as shown in Figure 5-59.

```
/tmp # esxcli software vib install -d
/tmp/VMW-ESX-5.5.0-nvme-1.0e.0.30-2284103/VMW-ESX-5.5.0-nvme-1.0e.0.30-offli
ne_bundle-2284103.zip
Installation Result
  Message: The update completed successfully, but the system needs to be
rebooted for the changes to be effective.
  Reboot Required: true
  VIBs Installed: VMware_bootbank_nvme_1.0e.0.30-1vmw.550.0.0.1391871
  VIBs Removed:
  VIBs Skipped:
```

Figure 5-59 NVMe driver installation process

6. Reboot the ESXi server.

- To confirm that the installed NVMe SSDs are recognized by ESXi server after restart, connect to the ESXi server by using vSphere Client, go to the Configuration tab, choose **Storage Adapters** from the Hardware section. You should see one or more NVMe controllers in the Storage Adapters list, as shown in Figure 5-60.




Figure 5-60 Recognized NVMe SSDs in Storage Adapters list

In the Details panel, you should see more information about the selected NVMe drive, as shown in Figure 5-61.

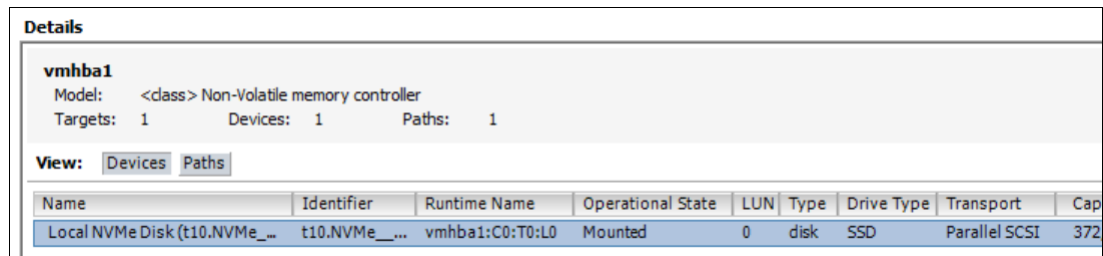


Figure 5-61 More information about the NVMe drive

The NVMe drives are now available for use. You can use NVMe drives as VMFS datastores or as Virtual Flash to improve I/O performance for all virtual machines or pass-through NVMe drives to the dedicated virtual machines.

## 5.4.6 Ongoing NVMe drive management

We discuss the ongoing management of NVMe drives, including how to correctly work with failed drives in a RAID array, in 7.10, “Managing NVMe drives and software RAID recovery” on page 295.

## 5.5 PCIe adapter placement advice

The x3850 X6 server supports up to 12 PCIe slots (including the dedicated ML2 slot) with up to three I/O Books and one Storage Book that is installed in the 4-socket chassis. The x3950 X6 offers up to 24 PCIe slots (including two ML2 slots) with up to six I/O Books and two Storage Books that are installed in the 8-socket chassis.

This section describes considerations to remember for determining how to use your PCIe slots, depending on the type of I/O Books and adapters that you installed.



## PCIe adapter slots

Table 5-5 shows the associated CPU with each PCIe adapter slot in the system. You cannot install an adapter in an I/O Book that does not have an associated Compute Book.

Table 5-5 Processor and PCIe associations

Processor	PCIe slot association
CPU 0 (Compute Book 1)	9, 10 <sup>a</sup> (Primary I/O Book), and Storage Book slot 12
CPU 1 (Compute Book 2)	7, 8 (Primary I/O Book), and Storage Book slot 11
CPU 2 (Compute Book 3)	4, 5, and 6 (I/O Book 2)
CPU 3 (Compute Book 4)	1, 2, and 3 (I/O Book 1)

a. Slot 10 is intended for ML2 network adapters only and cannot be used for regular PCIe adapters.

Figure 5-62 shows the PCIe slots numbering of the x3850 X6.




Figure 5-62 x3850 X6 PCIe slot numbering

Figure 5-63 shows the PCIe slots numbering of the x3950 X6.




Figure 5-63 x3950 X6 PCIe slot numbering

## PCIe performance

The x3850 X6 and x3950 X6 offer many PCIe slots. To ensure the best possible performance, it is important to understand the design of the servers and how the PCIe slots connect to each processor.

Figure 5-64 shows processor-to-slot PCIe connectivity for the x3850 X6.

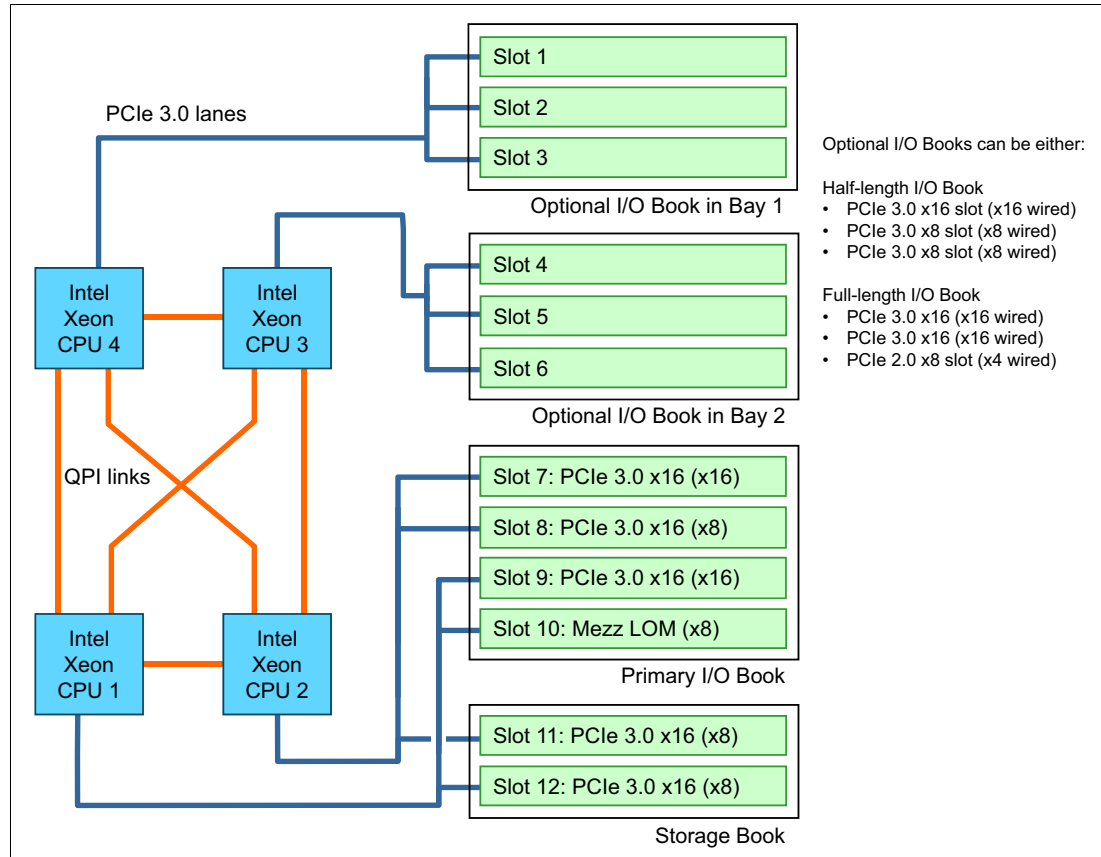


Figure 5-64 PCIe slot connectivity in the x3850 X6

For best performance, consider the following suggestions:

- ▶ Ensure that the slot that is used for an adapter has the necessary bandwidth. For example, slot 8 (in the Primary I/O Book) is an x8 wired slot, but it can physically accept an x16 adapter. Performance is negatively affected if an x16 adapter is installed in that slot.
- ▶ Balance the PCIe load across all processors. This balance ensures a distributed workload and reduces the chance of performance bottlenecks.
- ▶ Use processor affinity in the operating system to tie application threads that use PCIe resources the most to the processor that connects to that PCIe device.

## 5.6 Hot-swap procedures

The x3850 X6 and x3950 X6 hardware supports the ability to hot-swap certain components of the server. The term *hot-swap* refers to adding or removing certain hardware components while the server is running.

The following resources can be hot-swapped in the x3850 X6 and x3950 X6 server:

- ▶ All 2.5-inch and 1.8-inch drives
- ▶ All power supplies
- ▶ All system fans
- ▶ Optional I/O Books

Considerations for hot-swapping some of these components are described next.

### 5.6.1 Hot-swapping a power supply

Power supplies are independently powered by line cords and are hot-swappable from the rear of the server.

**Mixing power supplies:** The 900 W and 1400 W power supplies can be mixed in pairs; however, the 750 W DC power supply cannot be used with any AC supplies. For more information about power supply pairs, see 4.5.2, “Power supply redundancy” on page 137.

You can use the IMM2 to set and change the power supply Power Policy and System Power Configurations. You can set and change the policies and configurations by using the IMM2 web interface, CIM, or the ASU. You cannot set or change the Power Policy or System Power Configurations by using the UEFI Setup utility. The default configuration setting for AC and DC power supply models is non-redundant with throttling enabled.

For information about how to access the IMM, see 7.2, “Integrated Management Module II” on page 244. From the IMM2 web interface, select **Server Management** → **Power Management**. The Power Management panel is shown in Figure 5-65.

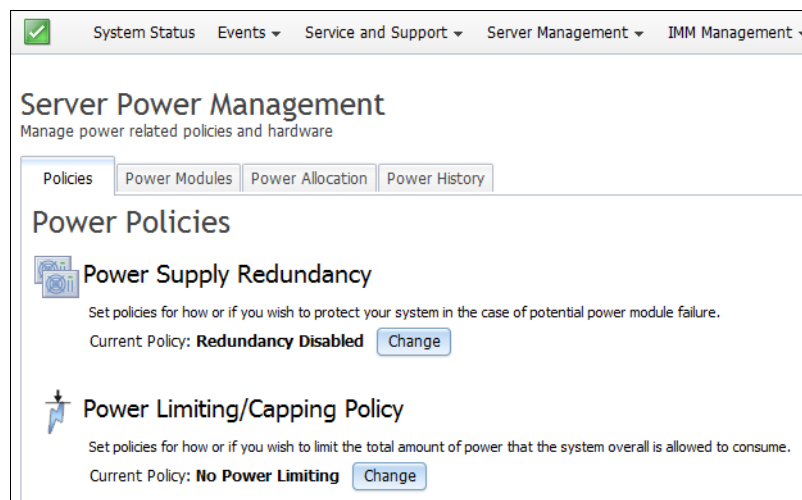


Figure 5-65 IMM2 Power Policy management panel

**Tip:** If you intend to upgrade your 900 W power supplies to 1400 W power supplies, ensure that your system's power is operating in a redundant configuration before removing a power supply for replacement.

## 5.6.2 Hot-swapping an I/O Book

As described in 3.15, “Hot-swap adapter support” on page 108, the Full-length I/O Books and Half-length I/O Books are hot-swap components.

**Primary I/O Book:** Hot-swapping of the Primary I/O Book is not supported.

If the Primary I/O Book or any non-hot swappable component must be added or removed, remove the AC power and wait for the LCD display and all LEDs to turn off.

Consider the following points regarding hot-swap:

- ▶ Only adapters in the Full-length I/O Books or Half-length I/O Books support hot-swap. The Primary I/O Book and Storage Book do not support hot-swap adapters.
- ▶ The operating system must support hot-swap. For more information, see “Hot-swap adapter support” on page 108.
- ▶ For an I/O Book to be hot-swappable, all adapters that are installed in the book must support hot-swap.
- ▶ Only certain adapters support hot-swap. Table 3-18 on page 109 lists the adapters that supported hot-swap at the time of this writing.
- ▶ When you remove the Full-length I/O Book or Half-length I/O Book, all adapters that are installed in that book also are removed.

The proper procedure to remove involves alerting the operating system that the adapters are being removed before physically removing the book. If the adapters are not properly brought offline before the I/O Book is removed, a Live Error Recovery (LER) results because a PCIe link goes offline unexpectedly (UEFI refers to this issue as “Surprise Link Down”).

PCIe Live Error Recovery occurs when errors are detected by the PCIe root port. The LER feature brings down the PCIe link that is associated with the affected root port within one cycle and then automatically recover the link. PCIe LER also protects against the transfer of associated corrupted data during this process.

**Warning:** The ability to hot-swap an I/O Book requires operating system support. If the OS does not support PCIe hot plug, removing or adding the I/O Book can cause an unrecoverable system error.

### Hot removal of an I/O Book

Complete the following steps to remove an I/O Book while the operating system is running:

1. Follow the instructions for your operating system to take your adapter offline.
2. Press the **Attention** button on the I/O Book to start a hot remove. The Attention button transitions power to offline for the I/O Book. For more information about the location of the Attention button, see Figure 5-66 on page 203.

The power indicator flashes, which indicates that a power is transitioning from on to off. This process takes approximately 5 seconds. You also can press the Attention button again while the indicator is flashing if you want to cancel the operation.




Figure 5-66 Location of the Attention button and LEDs on an Optional I/O Book

3. Wait for all three power indicators to turn off (see Figure 5-66). It is now safe to remove the I/O Book.
4. Open the I/O Book cam handle.
5. Slide the I/O Book out and away from the server.

Figure 5-67 shows adding and removing the Half-length I/O Book.




Figure 5-67 Adding and removing a Half-length I/O Book

## Hot add of an I/O Book

Complete the following steps to add an I/O Book while the operating system is running:

1. Install the adapters into the I/O Book. For the Full-length I/O Book, connect the supplied auxiliary power cables to the adapter as needed.
2. Open the I/O Book cam handle.
3. For the Full-length I/O Book, install the supporting brackets and top cover that are included with the book.
4. Align the I/O Book with the slot in the server and slide it into the server. Ensure that it is seated properly.
5. Rotate up the handle all the way and push it into the server until the handle locks in place.
6. The power indicator flashes, which indicates that a power is transitioning from on to off. This process takes approximately 5 seconds. You can also press the Attention button again while the indicator is flashing if you want to cancel the operation.

Figure 5-68 shows adding and removing the Full-length I/O Book.




Figure 5-68 Adding and removing a Full-length I/O Book

The Full-length I/O Book can be used in the I/O Book bays that are associated with CPU 2 or CPU 3. You can install up to two Full-length I/O Books.

**Note:** Because of the extended length of the full-length PCIe adapters, the Full-length I/O Book adds a 4-inch mechanical extension to the base length dimension of the server.

Table 5-6 lists the meaning of the I/O Book Attention LED. For more information about the location of this LED, see Figure 5-66 on page 203.

Table 5-6 Attention indicator for optional I/O Books

Appearance	Meaning
Off	Normal: Normal operation.
On	Attention: Operational problem at this slot.
Flashing	The slot is powering on or off. Do not remove the I/O Book when in flashing state.

Table 5-7 list the meaning of the I/O Book Power LED. For more information about the location of this LED, see Figure 5-66 on page 203.

Table 5-7 Power indicator for optional I/O Books

Appearance	Meaning
Off	Power Off: Insertion or removal of the adapter is permitted.
On	Power On: Insertion or removal of the adapter is not permitted.
Flashing	Power Transition: Hot-plug operation is in progress and insertion or removal of the adapter is not permitted.

## 5.7 Partitioning the x3950 X6

The x3950 X6 server can be electronically split into the following configurations by using a process called FlexNodePartitioning (or simply “partitioning”):

- ▶ An x3950 X6 server (8U chassis) with four processors can be split into two logical 2-socket systems.
- ▶ An x3950 X6 server (8U chassis) with eight processors can be split into two logical 4-socket systems.

Each half of the x3950 X6 is called a *node*.

**Note:** The x3850 X6 cannot be partitioned.

The following minimum number of components is required to support partitioning in the x3950 X6:

- ▶ Four Compute Books (two in each node)
- ▶ Two standard I/O Books (one in each node)
- ▶ Two Storage Books (one in each node)
- ▶ Two boot devices (one in each node), such as a local drive, hypervisor key, or external boot device
- ▶ Four power supplies (two power supplies in each node)

Consider the following points regarding partitioning an x3950 X6:

- ▶ By default, a single partition is created out of both halves of the x3950 X6 (that is, to form a single system image that uses all eight processors).
- ▶ Each half independently maintains its own set of hardware, such as storage backplanes, power supplies, and each scalable partition can support its own independent operating system installation. However, both halves must be at the same level of UEFI firmware and IMM firmware.
- ▶ Partitioning is performed at a whole node level of granularity. For example, if you have eight processors in your x3950 X6, you cannot have a partition of two processors and a partition of six processors. With eight processors installed, partitions can be four processors or eight processors only.
- ▶ The x3950 X6 server must be shut down before partitions can be created or changed.
- ▶ The bottom half of the x3950 X6 is considered the primary node in the complex. The primary node provides the information to the secondary node regarding partitioning.
- ▶ After being created, partitions cannot be expanded (for example, 4-socket partition to 8-socket partition). To create an 8-socket partition (the entire server), you must first remove the 4-socket partitions.
- ▶ Partition data remains consistent through reboots, shutdowns, and firmware upgrades.

Partitioning can be managed in the following ways:

- ▶ Via the IMM2 system management web interface
- ▶ Via the IMM2 telnet interface

The web interface is described next.

### 5.7.1 Partitioning an x3950 X6 via the IMM2 web interface

The x3950 X6 server cannot be dynamically partitioned (while the system is powered on). The system must be in powered off, but the power supplies must still be plugged into a power source. This process is done so you can still access the IMM2 to perform the partitioning functions.

The x3950 X6 server is made up of an upper node and a lower node, which correspond to the two physical halves of the x3950 X6. The nodes are joined internally via QPI links in the midplane to form an 8-socket server.

Nodes can be in one of the following states:

- ▶ Assigned node

These nodes are servers that are logically grouped into a partition. Assigned nodes are in one of the following modes:

- Partition mode: All processors are part of one (up to eight-socket) server. When you remove a partition, the partition grouping is dissolved and the nodes behave as independent systems (stand-alone).
- Stand-alone mode: Each half the processor is a separate server (each up to four-socket) and is functioning as two individual servers.



► Unassigned node

The servers are not a part of the assigned group and must be added. The processors cannot be accessed and the node does not function in an unassigned state. After it is added, the node can be assigned as a partition mode or assigned as stand-alone mode. An example of two nodes in an unassigned state is shown in Figure 5-69.

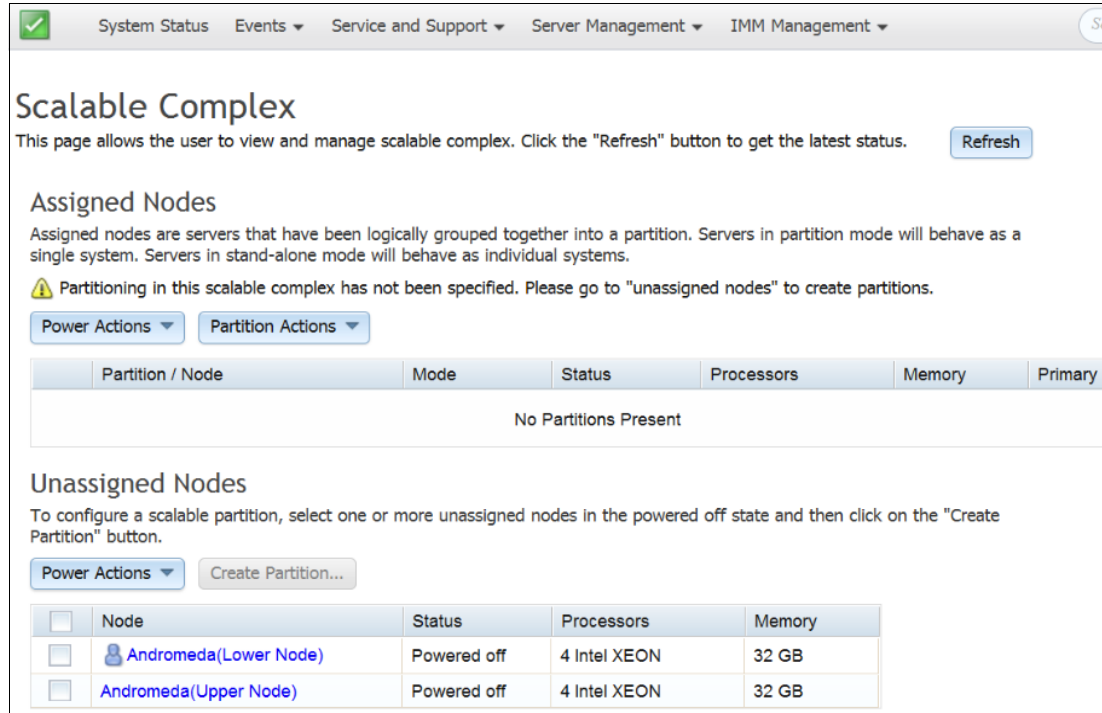


Figure 5-69 Scalable complex window with unassigned nodes

By using the IMM2 web interface scalable complex window, you can perform the following functions on assigned nodes:

► Power Actions

Power on or off the nodes immediately, shut down OS and power off, and restart the nodes, as shown in Figure 5-70.




Figure 5-70 Scalable complex window: Power Actions

► Partition Actions

As shown in Figure 5-71, the following Partition Actions are available:

- Remove Partition: Creates an unassigned node. An unassigned node can be redeployed back into an assigned node at any point. A node must be assigned to function.
- Activate Stand-alone Mode: Turns the server into two 4-socket servers that behave as individual systems.
- Restore Partition Mode: Turns the server back into one 8-socket server.




Figure 5-71 Scalable complex window: Partition Actions

Complete the following steps to remove the partition that is the x3950 X6 server and create two stand-alone servers:

1. Log on to the IMM2 web interface. For information about how to set up and access the IMM2, see 5.1.2, “IMM2 network access” on page 157.
2. Access the Scalable Complex window by selecting **Server Management** → **Scalable Complex** from the top menu, as seen in Figure 5-72.

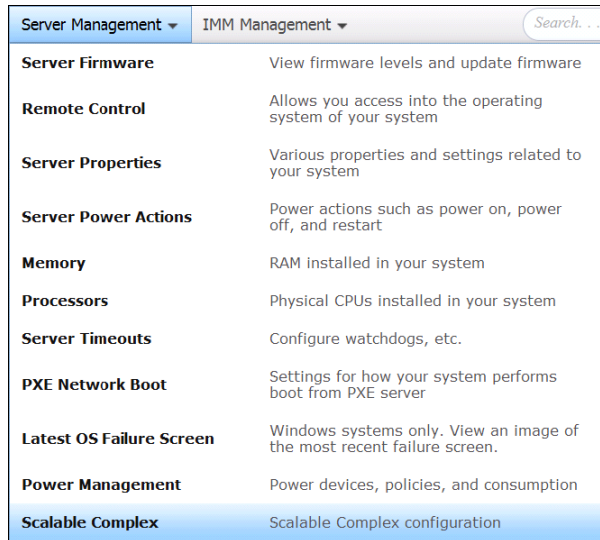


Figure 5-72 Accessing the Scalable Complex window via the IMM2 web interface

Figure 5-73 shows the scalable complex window pane with one x3950 X6 server, which contains two CPUs and 32 GB of RAM in each partition.

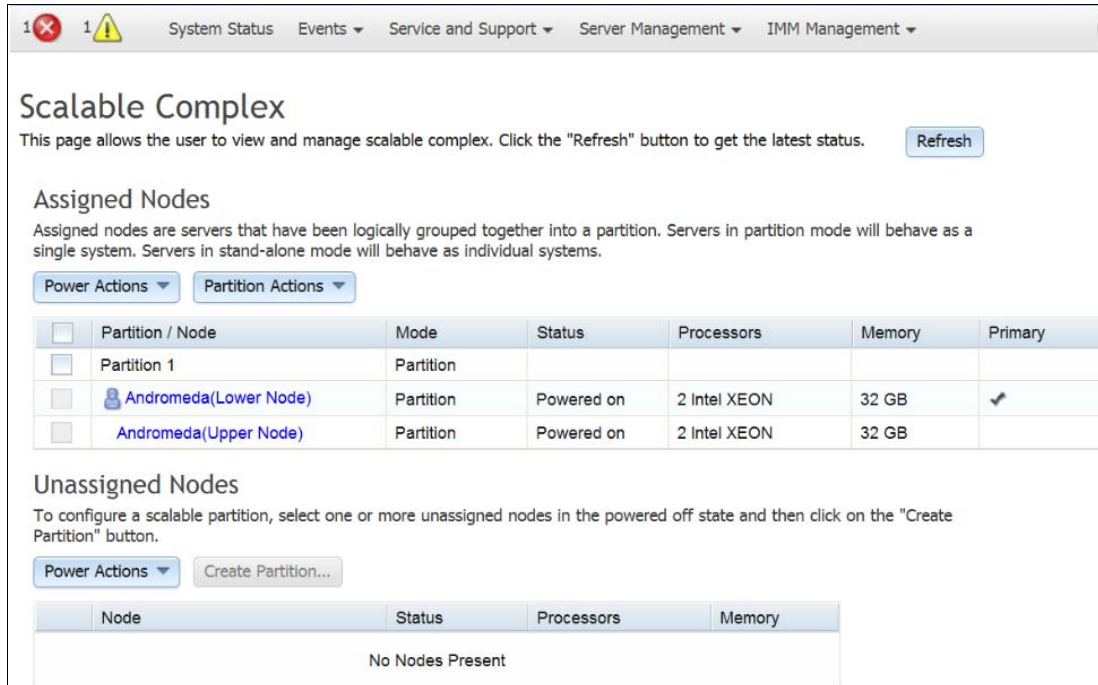


Figure 5-73 IMM2 web interface Scalable Complex window

- Before separating the x3950 X6 into two stand-alone servers, ensure that you turn off the server. Failing to turn off the server results in an error message, as shown in Figure 5-74.

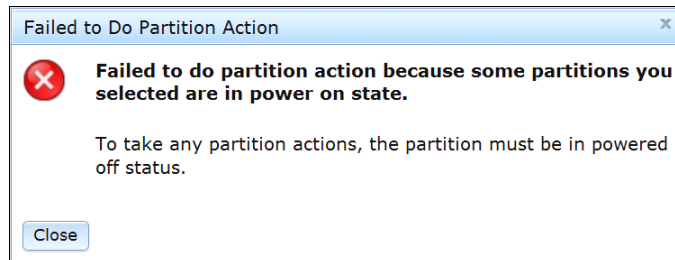


Figure 5-74 Partition failure message

- Check the option the left of the server that you intend to power off. When selected, you can turn off the server via the Power Actions menu, as shown in Figure 5-75.




Figure 5-75 Power Actions menu

- After the server is turned off, select **Partition Actions** → **Activate Stand-alone Mode**, as shown in Figure 5-76.




Figure 5-76 Partition options in the scalable complex window via the IMM2

- Click **Activate Stand-alone Mode** in the confirmation message, as shown in Figure 5-77.

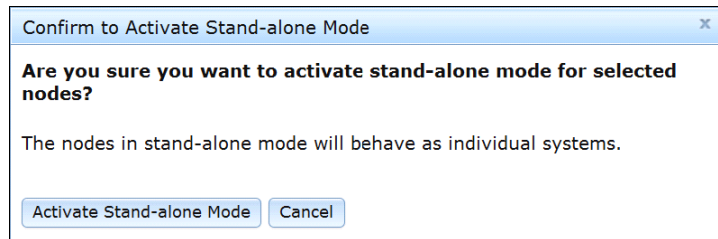


Figure 5-77 Partition confirmation message

- A progress window opens, as shown in Figure 5-78. You can track the progress of the partitioning by refreshing the page.




Figure 5-78 Partition progress panel

The servers are now listed as stand-alone and behave as two individual servers, as shown in Figure 5-79.



Figure 5-79 Stand-alone mode

You can restore partition mode on the servers by highlighting the server with the check box and selecting **Partition Actions** → **Restore Partition Mode**, as shown in Figure 5-80. By restoring the server to partition mode, you are restoring it to a full x3950 X6 server and it no longer functions as two separate servers.




Figure 5-80 Restoring partition mode

## 5.8 Updating firmware

Firmware updates often are released to provide clients with enhanced capabilities, extended functions, and problem resolutions. We advise performing a scheduled review of available updates to determine whether they are applicable to the systems that are used in your environment.

This section includes the following topics:

- ▶ 5.8.1, “Firmware tools” on page 211
- ▶ 5.8.2, “Updating firmware” on page 212

### 5.8.1 Firmware tools

Multiple methods are available for performing firmware updates. The preferred method to perform firmware updates is to use one of the following tools:

- ▶ XClarity Administrator
- ▶ UpdateXpress System Pack Installer (UXSPI)
- ▶ Bootable Media Creator™ (BoMC)
- ▶ IMM (for updating IMM, UEFI, and DSA only)

UXSPI can BoMC perform the following functions:

- ▶ Display an inventory of installed firmware and drivers.
- ▶ Download firmware and drivers.
- ▶ Download UXSP.
- ▶ Update all of the firmware and drivers in your system, including RAID, drives, network adapters, and Fibre Channel host bus adapters (HBAs).
- ▶ Apply updates in the correct order to update a system with the fewest reboots.
- ▶ Create a bootable CD/DVD/USB key/Preboot eXecution Environment (PXE) image to perform firmware updates (BoMC).

Firmware updates are provided by Lenovo and can be downloaded from the support site, including proven firmware from other manufacturers to be applied on Lenovo systems. In this section, we describe the methods of performing firmware updates by using UXSPI, BoMC, and IMM.

**Tip:** It is advised to update all System x firmware to the latest level before performing an OS or application installation.

You can also perform firmware updates by using Lenovo XClarity Administrator, IBM Systems Director, or Upward Integration Module (UIM) with a hypervisor.

### Lenovo XClarity Administrator

Within Lenovo XClarity Administrator, you can manage the firmware updates repository and apply and activate firmware updates for all managed endpoints. You can manage individual updates, set up updates, or define compliance policies with the firmware rules. For more information, see 7.4.5, “Firmware updates” on page 256.

### Upward Integration Module

The UIM integrates PFA and microcode management and diagnostic tests into standard hypervisors through added value modules for X6.

UIMs provide hardware visibility to the hypervisor for superior system and VM management with which you can perform the following tasks:

- ▶ **Concurrent Firmware Updates:** All system software can be concurrently updated in a virtualized environment with a single command.
- ▶ **Reliability, availability, serviceability (RAS) with UIM:** By using RAS, you can manage and set policies around all PFA in the system and evacuate, migrate, or manage VMs before an outage affects them.

For more information about the Upward Integration Module, see 7.5, “Lenovo XClarity integrators” on page 258.

## 5.8.2 Updating firmware

Before updates are applied, ensure that flashing is enabled on your system. To enable flashing, you must enable the LAN over USB interface in UEFI first.

For the x3850 X6 and x3950 X6 servers, in the UEFI menu, select **System Settings** → **Integrated Management Module** → **Commands on USB Interface Preference** and enable Commands on USB interface, as shown in Figure 5-81.

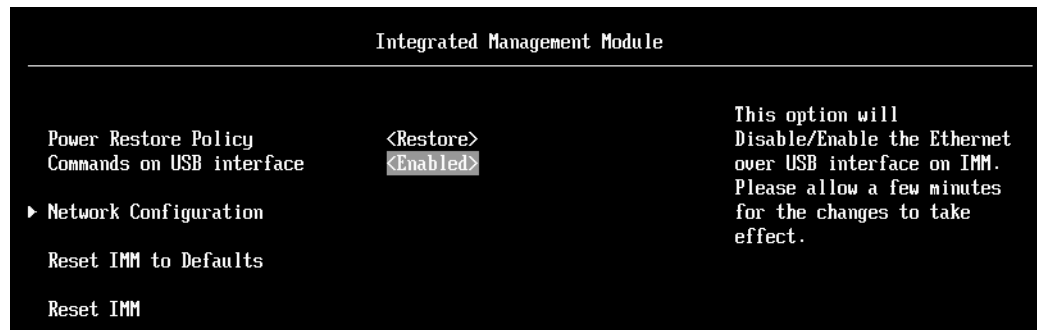


Figure 5-81 Enable flashing in the UEFI

## UpdateXpress System Pack Installer

UpdateXpress is a tool that allows the System x firmware and drivers to be updated via the OS.

By using the UpdateXpress System Pack Installer (UXSPI), you can update the firmware and device drivers of the system under an operating system. You also can deploy UpdateXpress System Packs™ (UXSPs) and the latest individual updates.

UXSPI uses the standard HTTP (port 80) and HTTPS (port 443) to get the updates from IBM. Your firewall must allow these ports. UXSPI is supported on Windows, Linux, and VMware operating systems. UXSPI is supported on 32-bit and 64-bit operating systems.

For more information about supported operating systems, see *UpdateXpress System Pack Installer User's Guide*, which is available at this website:

<https://support.lenovo.com/documents/LNVO-XPRESS>

Complete the following steps to use UXSPI:

1. Start the UXSPI setup utility that you downloaded from the website.
2. Accept the license agreement.

The Welcome window opens, as shown in Figure 5-82 on page 213. Click **Next**.

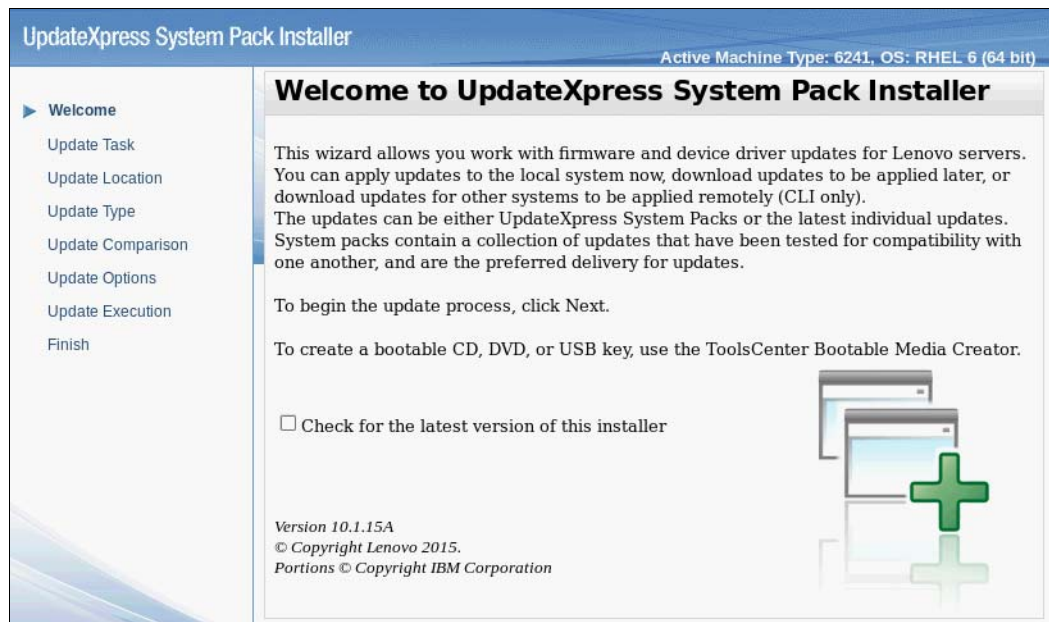


Figure 5-82 Welcome window

3. Accept the default **Update the local machine**, as shown in Figure 5-83. Click **Next**.

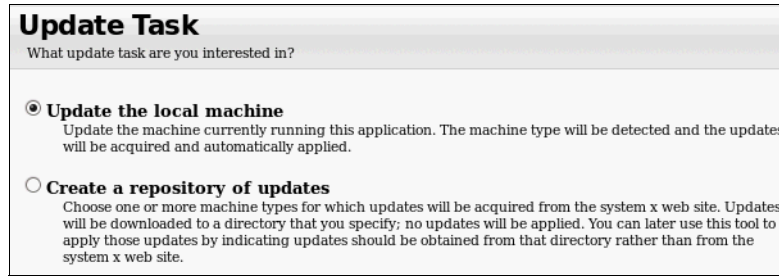


Figure 5-83 Selecting the update task

4. Accept the default **Check the system x web site**, as shown in Figure 5-84. Click **Next**.

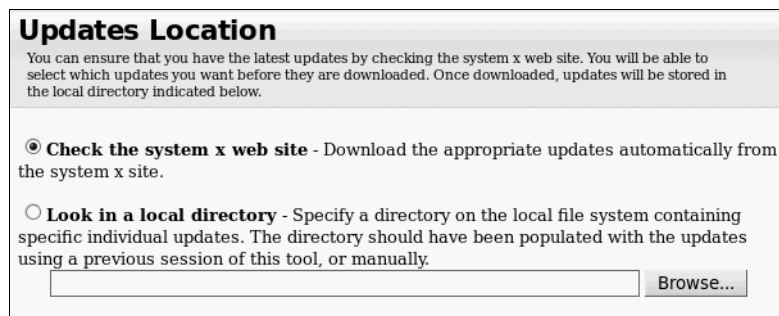


Figure 5-84 Selecting the update task

5. Select **Latest available individual updates**, as shown in Figure 5-85. Click **Next**.

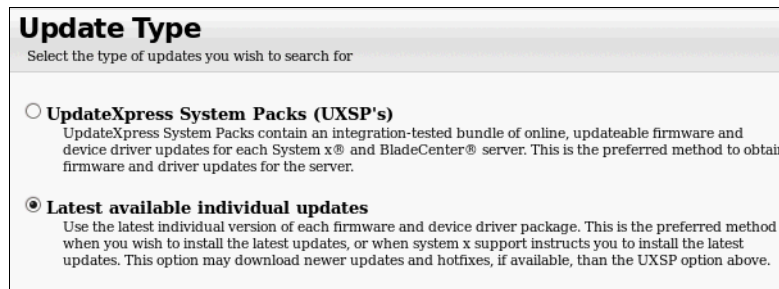


Figure 5-85 Selecting the type of updates

6. Enter the settings for an HTTP proxy server (if necessary) or leave the option cleared, as shown in Figure 5-87. Click **Next**.

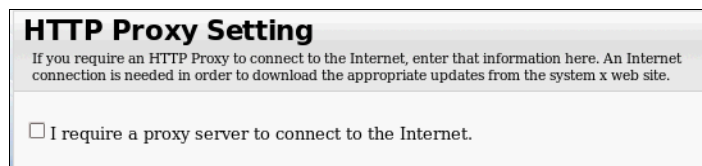


Figure 5-86 HTTP proxy setting



7. Select the directory in which you want to store the downloaded files, as shown in Figure 5-87. Click **Next**.




Figure 5-87 Selecting your target directory

8. A message appears that shows that the UXSPI acquired the updates for the machine, as shown in Figure 5-88. Click **Next**.




Figure 5-88 Successful completion of acquisition report

9. A message appears that shows that the download completed, as shown in Figure 5-89. Click **Next**.




Figure 5-89 Download process completes successfully

10. A component overview shows the components that need updating. By default, UXSPI selects the components to update. Accept these settings and click **Next**.
11. When the update is finished, a message appears that confirms the updates. Click **Next**.



Figure 5-90 Updates are successful

12. Click **Finish** to close the UXSPI.
13. Restart the system to complete the update process.

## Bootable Media Creator

The BoMC provides a tool for creating a bootable image for supported media (CD, DVD, ISO image, USB flash drive, or PXE files) to update the system firmware. Because BoMC runs in its own boot environment, you cannot update drivers. BoMC has a graphical and command-line interface. One bootable media image can include support for multiple systems. The tool uses standard HTTP (port 80) and HTTPS (port 443) to receive updates from IBM. Your firewall must allow these ports.

BoMC is supported on Windows, Linux, and VMware operating systems. BoMC supports 32-bit and 64-bit operating systems. For more information about supported operating systems, see *ToolsCenter BOMC Installation and User's Guide*, which is available at the BOMC web page:

<https://support.lenovo.com/documents/LNVO-BOMC>

Complete the following steps to create an updated media:

1. Create a folder named BoMC.

- Download the latest version of BoMC from the support page and save it in the BoMC folder:

<https://support.lenovo.com/documents/LNV0-BOMC>

From a command line, enter the command to start the BoMC. The command depends on the operating system, as listed in Table 5-8.

Table 5-8 Command for each supported operating system

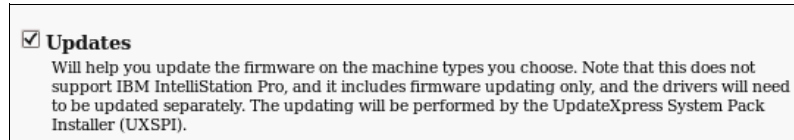
Operating system	Command
Windows	lnvgy_utl_bomc_v.r.m_windows_i386.exe
Red Hat Enterprise Linux 4.0 32-bit	lnvgy_utl_bomc_v.r.m_rhel4_i386.bin
Red Hat Enterprise Linux 4.0 64-bit	lnvgy_utl_bomc_v.r.m_rhel4_x86-64.bin
Red Hat Enterprise Linux 5.0 32-bit	lnvgy_utl_bomc_v.r.m_rhel5_i386.bin
Red Hat Enterprise Linux 5.0 64-bit	lnvgy_utl_bomc_v.r.m_rhel5_x86-64.bin
Red Hat Enterprise Linux 6.0 32-bit	lnvgy_utl_bomc_v.r.m_rhel6_i386.bin
Red Hat Enterprise Linux 6.0 64-bit	lnvgy_utl_bomc_v.r.m_rhel6_x86-64.bin
Red Hat Enterprise Linux 7.0 64-bit	lnvgy_utl_bomc_v.r.m_rhel7_x86-64.bin
SUSE Linux Enterprise Server 10 32-bit	lnvgy_utl_bomc_v.r.m_sles10_i386.bin
SUSE Linux Enterprise Server 10 64-bit	lnvgy_utl_bomc_v.r.m_sles10_x86-64.bin
SUSE Linux Enterprise Server 11 32-bit	lnvgy_utl_bomc_v.r.m_sles11_i386.bin
SUSE Linux Enterprise Server 11 64-bit	lnvgy_utl_bomc_v.r.m_sles11_x86-64.bin
SUSE Linux Enterprise Server 12 64-bit	lnvgy_utl_bomc_v.r.m_sles12_x86-64.bin

- Accept the license agreement.
- The Welcome window opens, as shown in Figure 5-91. Click **Next**.



Figure 5-91 Welcome page

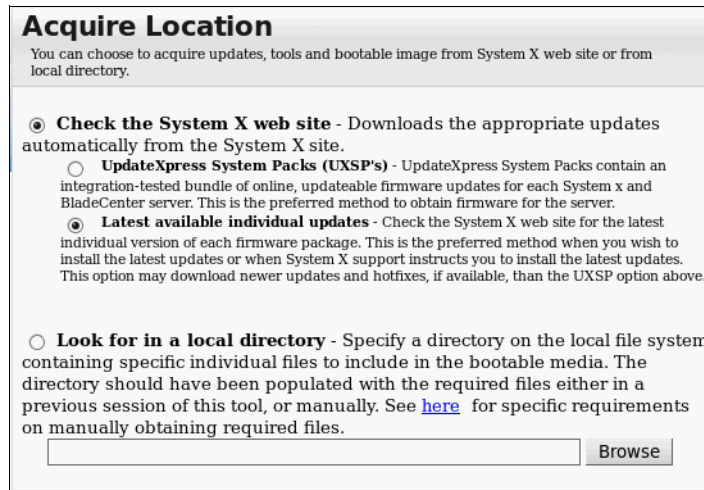
4. Select **Updates**, as shown in Figure 5-92. Click **Next**.



**Updates**  
Will help you update the firmware on the machine types you choose. Note that this does not support IBM IntelliStation Pro, and it includes firmware updating only, and the drivers will need to be updated separately. The updating will be performed by the UpdateXpress System Pack Installer (UXSPI).

Figure 5-92 Selected update task

5. Select **Latest available individual updates** and click **Next**, as shown in Figure 5-93.



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

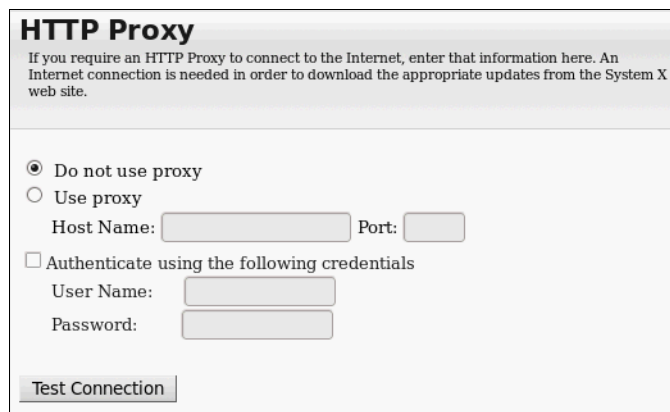
**Check the System X web site** - Downloads the appropriate updates automatically from the System X site.

- 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 System X 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 System X support instructs you to install the latest updates. This option may download newer updates and hotfixes, if available, than the UXSP option above.

**Look for 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 5-93 Select source of updates

6. Enter the settings for an HTTP proxy server (if necessary) or select **Do not use proxy**, as shown in Figure 5-94. Click **Next**.



**HTTP Proxy**  
If you require an HTTP Proxy to connect to the Internet, enter that information here. An Internet connection is needed in order to download the appropriate updates from the System X web site.

Do not use proxy

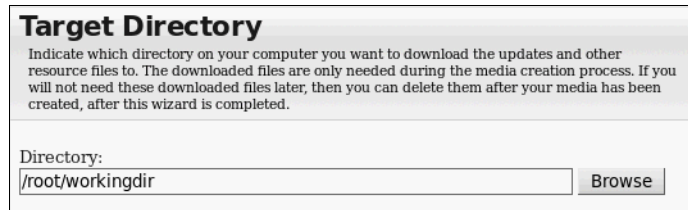
Use proxy  
Host Name:  Port:

Authenticate using the following credentials  
User Name:   
Password:

Figure 5-94 HTTP Proxy settings

7. Select one or more machine types that are on the bootable media and click **Next**.

8. Select the directory in which you want to store the downloaded files, as shown in Figure 5-95. Click **Next**.



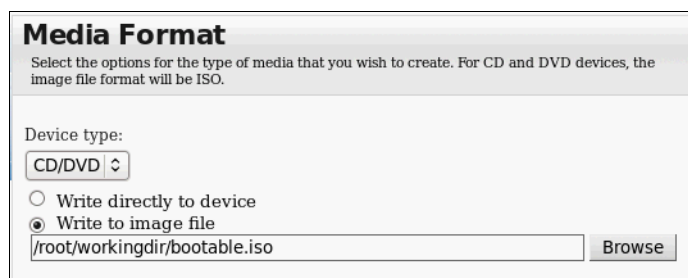
**Target Directory**

Indicate which directory on your computer you want to download the updates and other resource files to. The downloaded files are only needed during the media creation process. If you will not need these downloaded files later, then you can delete them after your media has been created, after this wizard is completed.

Directory: /root/workingdir

Figure 5-95 Example of a target directory

9. By default, BoMC creates an ISO file, as shown in Figure 5-96. You can choose another medium. Click **Next**.



**Media Format**

Select the options for the type of media that you wish to create. For CD and DVD devices, the image file format will be ISO.

Device type: CD/DVD

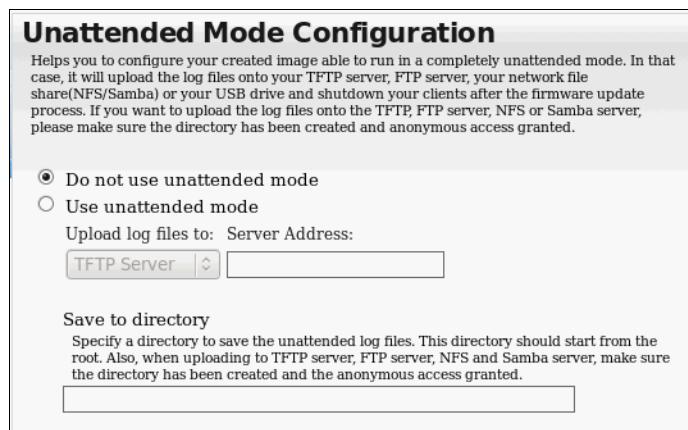
Write directly to device

Write to image file

/root/workingdir/bootable.iso

Figure 5-96 Example of a target media

10. Select **Do not use unattended mode**, as shown in Figure 5-97. Click **Next**.



**Unattended Mode Configuration**

Helps you to configure your created image able to run in a completely unattended mode. In that case, it will upload the log files onto your TFTP server, FTP server, your network file share(NFS/Samba) or your USB drive and shutdown your clients after the firmware update process. If you want to upload the log files onto the TFTP, FTP server, NFS or Samba server, please make sure the directory has been created and anonymous access granted.

Do not use unattended mode

Use unattended mode

Upload log files to: Server Address: TFTP Server

Save to directory:

Figure 5-97 Unattended Mode panel

11. Review the selections and confirm that they are correct. You can click **Save** to save this configuration information to a file. Click **Next**.

BoMC acquires the files. In the progress bar, you can see the progress of the updates, as shown in Figure 5-98.

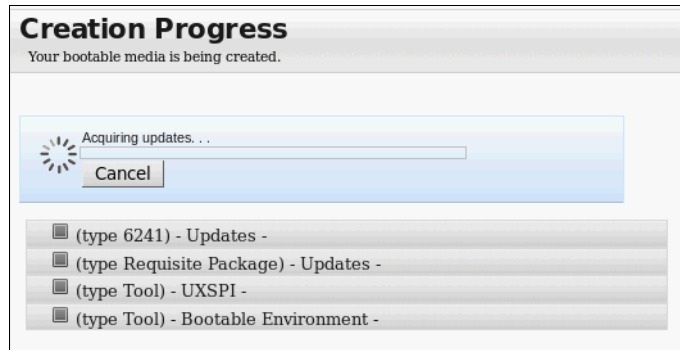


Figure 5-98 Downloading the files

12. After the update completes, click **Next**. Then, click **Finish**.
13. You now have a bootable image with the updates. You can mount the ISO image through the IMM2 and then boot the system with the medium.

## Integrated Management Module

You can perform certain firmware updates by using the IMM. For more information about how to connect to the IMM, see 7.2, “Integrated Management Module II” on page 244. By using the IMM2, you can update DSA, IMM2 (primary and backup images), and UEFI (primary and backup images).

**Note:** Updating server firmware via the IMM2 is intended for recovery purposes. The preferred method of updating firmware is to use USPI, BoMC, or XClarity Administrator as described in this section.

Complete the following steps to update the IMM2, UEFI, and DSA via the IMM2 page:

1. From the IMM2 web interface, select **Server Management** → **Server Firmware**, as shown in Figure 5-99.

Firmware Type	Status	Version	Build	Release Date
DSA	Active	10.1	DSALA6L	2015-03-17
IMM2				
IMM2 (Primary)	Active	1.62	TCO008S	2015-04-07
IMM2 (Backup)	Inactive	1.08	TCO006A	2015-01-27
UEFI				
UEFI (Primary)	Active	1.00	A9E122XUS	2015-04-10
UEFI (Backup)	Inactive	V.VV	YYYYYYY	

Figure 5-99 The IMM2 Server Firmware panel

2. Click **Update Firmware**. A wizard starts, which guides you through the update process, as shown in Figure 5-100.

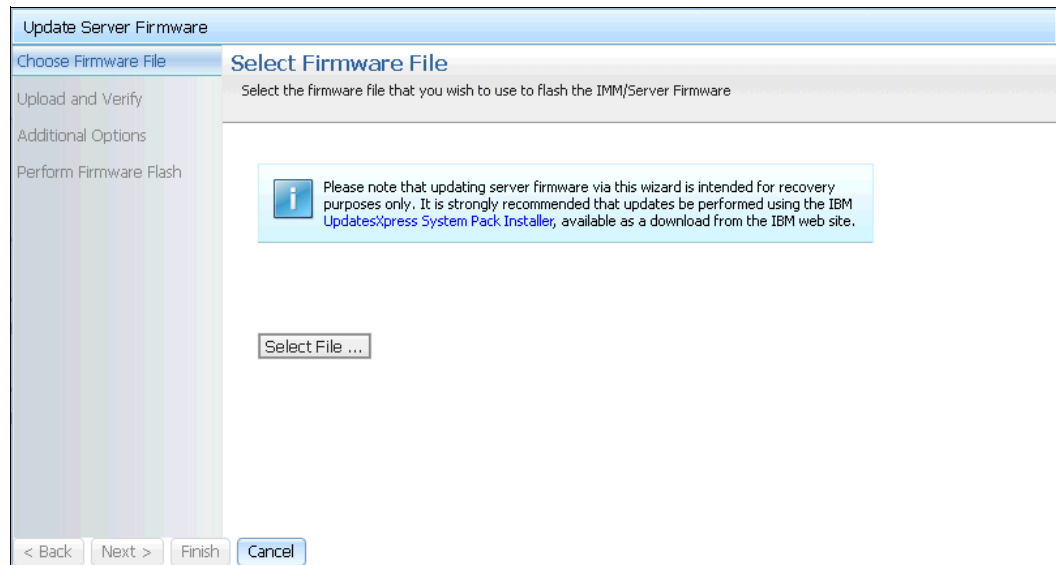


Figure 5-100 IMM2 firmware selection page

3. Ensure that you downloaded the appropriate firmware update from the Lenovo support page.
4. After you selected your file, you can perform the firmware flash.

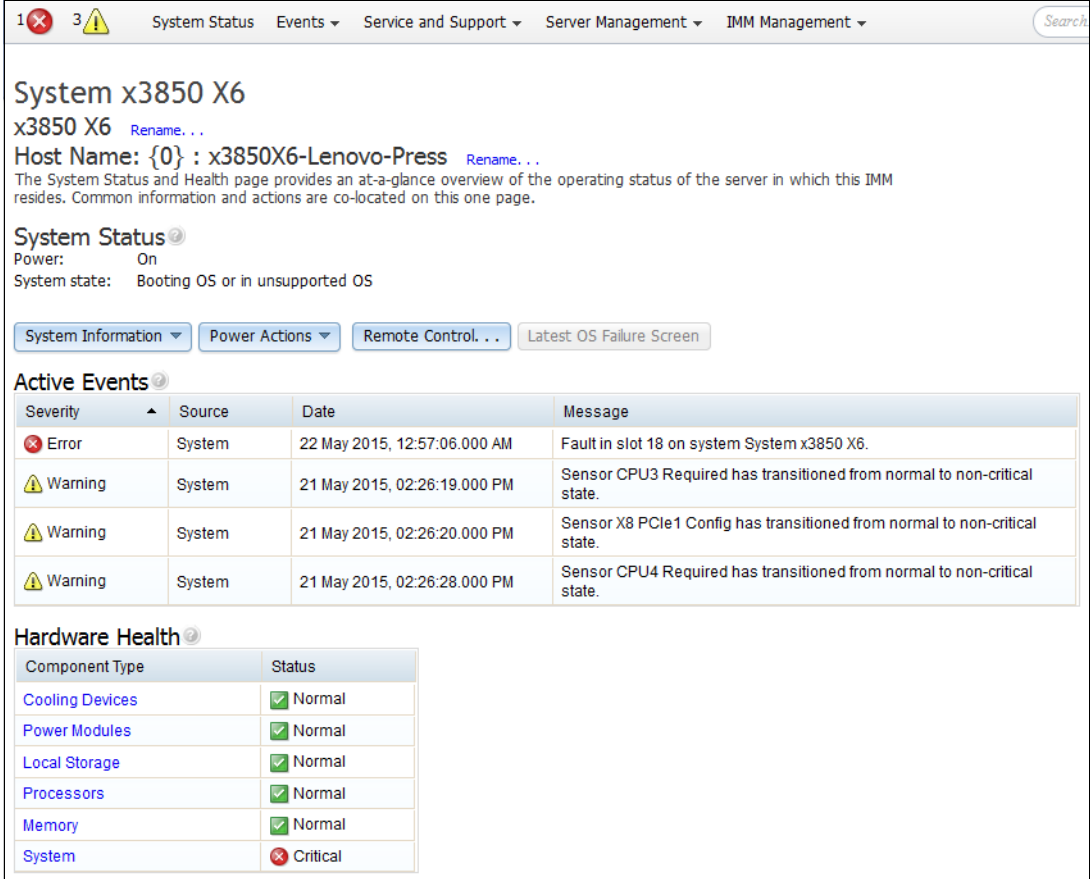
## 5.9 Troubleshooting

This section describes the tools that are available to assist with problem resolution for the X6 servers in any specific configuration. It also provides considerations for extended outages.

Use the following tools you are when troubleshooting problems on the X6 servers in any configuration.

### 5.9.1 Integrated Management Module

The first place to start troubleshooting the X6 servers often is the IMM. For more information about connecting to the IMM2, see to 7.2, “Integrated Management Module II” on page 244. After logging in, the first page is the System Status page, as shown in Figure 5-101.



The screenshot displays the IMM2 System Status page for a System x3850 X6. The page includes a navigation bar with tabs for System Status, Events, Service and Support, Server Management, and IMM Management. The main content area shows the system name, host name, and system status. Below this, there are buttons for System Information, Power Actions, Remote Control, and Latest OS Failure Screen. The Active Events section contains a table with columns for Severity, Source, Date, and Message. The Hardware Health section contains a table with columns for Component Type and Status.

Severity	Source	Date	Message
Error	System	22 May 2015, 12:57:06.000 AM	Fault in slot 18 on system System x3850 X6.
Warning	System	21 May 2015, 02:26:19.000 PM	Sensor CPU3 Required has transitioned from normal to non-critical state.
Warning	System	21 May 2015, 02:26:20.000 PM	Sensor X8 PCIe1 Config has transitioned from normal to non-critical state.
Warning	System	21 May 2015, 02:26:28.000 PM	Sensor CPU4 Required has transitioned from normal to non-critical state.

Component Type	Status
Cooling Devices	Normal
Power Modules	Normal
Local Storage	Normal
Processors	Normal
Memory	Normal
System	Critical

Figure 5-101 IMM2 System Status page

From this page, you can check the power status of the server and state of the OS.

You also can view the System Information, such as name, machine type, serial number, and state of the machine.

The Hardware Health of your system is also on this page, which monitors, fans, power, disk, processors, memory, and system.

By clicking the individual hardware, you can see more information about the selected component and read any errors that occurred. If a hardware error occurred, it is indicated by a red X in place of the green tick.

From the IMM, you also can access the hardware logs. From the main menu at the top of the panel, click **Events** → **Event log** to access a full log history of all events.

### 5.9.2 LCD system information panel

You can use the LCD information panel to diagnose system errors quickly. The LCD system information display panel is attached to the Storage Book on the front of the server. The LCD system information display panel enables you to have quick access to system status, firmware, network, and health information. Figure 5-102 shows the controls on the LCD system information display panel.




Figure 5-102 LCD System Information panel

The information that is displayed on the LCD panel is shown in Figure 5-103.

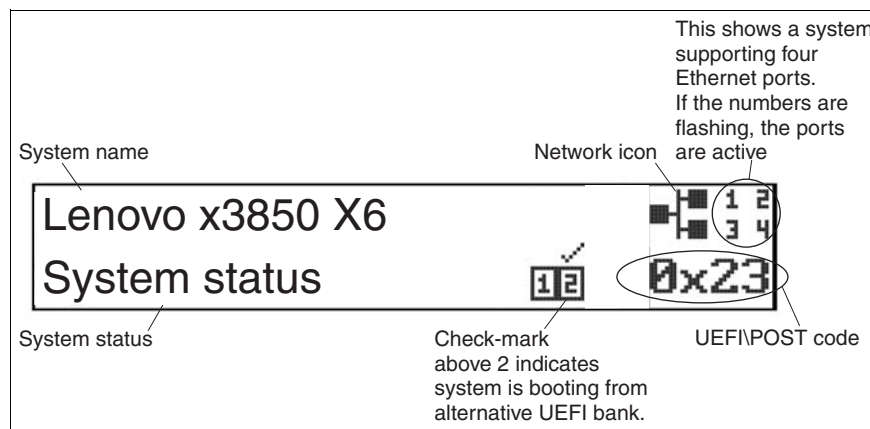


Figure 5-103 LCD System Information panel

For more information about the LCD panel and error checking, see the *Installation and Service Guide*.

### 5.9.3 System event log

This log contains POST and system management interrupt (SMI) events and all events that are generated by the Baseboard Management Controller (BMC) that is embedded in the IMM.



You can view the system event log through the UEFI by pressing F1 at system start and selecting **System Event Logs** → **System Event Log**.

#### 5.9.4 POST event log

This log contains the three most recent error codes and messages that were generated during POST. You can view the POST event log through the UEFI by pressing F1 at system start and selecting **System Event Logs** → **POST Event Viewer**.

#### 5.9.5 Installation and Service Guide

You can solve many problems without outside assistance by following the troubleshooting procedures that are described in the *Installation and Service Guide*, which is available at this website:

[http://publib.boulder.ibm.com/infocenter/systemx/documentation/index.jsp?topic=/com.lenovo.sysx.6241.doc/product\\_page.html](http://publib.boulder.ibm.com/infocenter/systemx/documentation/index.jsp?topic=/com.lenovo.sysx.6241.doc/product_page.html)

This document describes the diagnostic tests that you can perform and troubleshooting procedures and explains error messages and error codes.

If you completed the diagnostic procedure and the problem remains and you verified that all code is at the latest level and all hardware and software configurations are valid, contact Lenovo or an approved warranty service provider for assistance.



# Operating system installation

This chapter provides an overview of the options that you have when an operating system (OS) is installed on the X6 family servers.

This chapter includes the following topics:

- ▶ 6.1, “Installing an OS without a local optical drive” on page 226
- ▶ 6.2, “ServerGuide” on page 233
- ▶ 6.3, “ServerGuide Scripting Toolkit” on page 237
- ▶ 6.4, “Use of embedded VMware ESXi” on page 239
- ▶ 6.5, “Booting from SAN” on page 240

At the time of writing, the x3850 X6 and x3950 X6 servers with v4 processors are supported with the following operating systems:

- ▶ Microsoft Windows Server 2012
- ▶ Microsoft Windows Server 2012 R2
- ▶ Microsoft Windows Server 2016
- ▶ Microsoft Windows Server, version 1709
- ▶ Red Hat Enterprise Linux 6.10 x64
- ▶ Red Hat Enterprise Linux 6.7 x64
- ▶ Red Hat Enterprise Linux 6.8 x64
- ▶ Red Hat Enterprise Linux 7.2
- ▶ Red Hat Enterprise Linux 7.3
- ▶ Red Hat Enterprise Linux 7.4
- ▶ Red Hat Enterprise Linux 7.5
- ▶ SUSE Linux Enterprise Server 11 Xen x64 SP4
- ▶ SUSE Linux Enterprise Server 11 x64 SP4
- ▶ SUSE Linux Enterprise Server 12 SP1
- ▶ SUSE Linux Enterprise Server 12 SP2
- ▶ SUSE Linux Enterprise Server 12 SP3
- ▶ SUSE Linux Enterprise Server 12 Xen SP1
- ▶ SUSE Linux Enterprise Server 12 Xen SP2
- ▶ SUSE Linux Enterprise Server 12 Xen SP3
- ▶ SUSE Linux Enterprise Server 15
- ▶ SUSE Linux Enterprise Server 15 Xen
- ▶ VMware ESXi 6.0 U2

- ▶ VMware ESXi 6.0 U3
- ▶ VMware ESXi 6.5
- ▶ VMware ESXi 6.5 U1
- ▶ VMware ESXi 6.5 U2
- ▶ VMware ESXi 6.7
- ▶ VMware ESXi 6.7 U1

For specific OS support, see the Lenovo Operating System Interoperability Guide:

<https://lenovopress.com/osig#term=6241&support=all>

**vSphere 5.1 and 8-socket systems:** VMware vSphere 5.1 has a fixed upper limit of 160 concurrent threads. Therefore, if you use an 8-socket system with more than 10 cores per processor, you should disable Hyper-Threading.

Failing to disable Hyper-Threading in the Unified Extensible Firmware Interface (UEFI) with 12-core or 15-core processors and vSphere 5.1 affects performance.

## 6.1 Installing an OS without a local optical drive

The X6 family servers do not have a local optical drive. To install an OS, you can use any of the following methods:

- ▶ 6.1.1, “Integrated Management Module” on page 226
- ▶ 6.1.2, “Local USB port” on page 229
- ▶ 6.1.3, “Preboot eXecution Environment” on page 229
- ▶ 6.1.4, “Lenovo XClarity Administrator” on page 229

Each method is described next.

### 6.1.1 Integrated Management Module

A remote control feature is available through the Integrated Management Module 2 (IMM2) web interface. You must log in to the IMM2 with a user ID that includes Supervisor access. You can also assign to the server a CD or DVD drive, USB flash drive, or disk image that is on your computer.

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

<https://support.lenovo.com/docs/UM103336>

Complete the following steps to mount a drive through IMM2:

1. Connect to the IMM2 by using your web browser.

2. Click **Remote Control** on the main page of IMM2, as shown in Figure 6-1.



Figure 6-1 Remote Control option on the IMM2 main page

3. If you want to allow other users remote control access during your session, click **Start Remote Control in Multi-user Mode**. Otherwise, click **Start Remote Control in Single User Mode**.

The Java application window should open, as shown in Figure 6-2 on page 227.



Figure 6-2 Video Viewer window

4. To mount an image to the remote server as virtual media, you should ensure Activate is selected under the Virtual Media menu, as shown in Figure 6-3.




Figure 6-3 Virtual Media activating

5. After Activate is selected, click **Select Devices to Mount** to mount an image to the remote server as virtual media, as shown in Figure 6-4.




Figure 6-4 Starting virtual media

6. Click **Add Image** if you want to map an IMG or ISO image file, as shown in Figure 6-5.

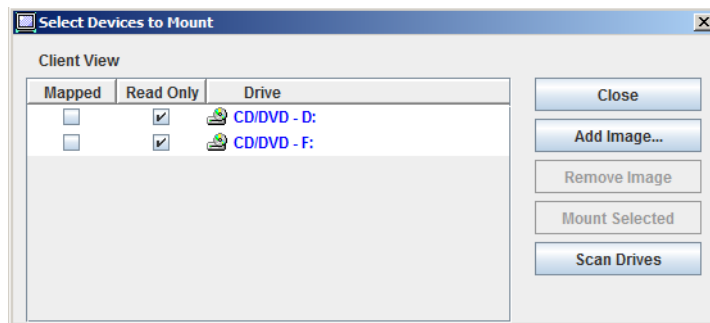


Figure 6-5 Adding an image

7. After adding an image, select the drive that you want to map and click **Mount Selected**, as shown in Figure 6-6.

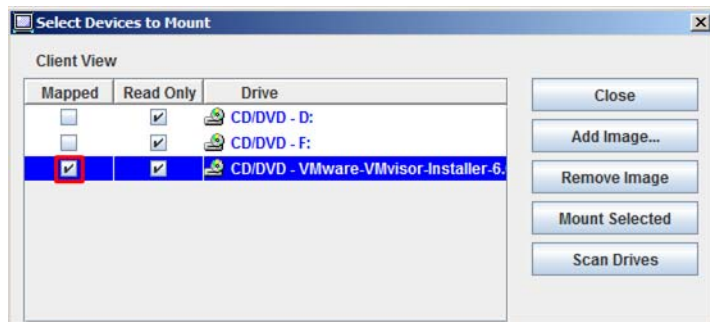


Figure 6-6 Selected image mounting

The image drive is now accessible by the system.

**Closing the session:** Closing the Virtual Media Session window when a remote disk is mapped to the machine causes the machine to lose access to the remote disk.

## 6.1.2 Local USB port

You can use the local USB port to attach a USB flash drive that contains the OS installation files. There are several methods available to create a bootable flash drive. For more information about the use of a USB key as an installation medium, see these websites:

- ▶ Installing Red Hat Linux from a USB flash drive:  
[https://access.redhat.com/site/documentation/en-US/Red\\_Hat\\_Enterprise\\_Linux/6/html-single/Installation\\_Guide/index.html#Making\\_Minimal\\_Boot\\_Media](https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux/6/html-single/Installation_Guide/index.html#Making_Minimal_Boot_Media)
- ▶ How to create a bootable USB drive to install SLES:  
<http://www.novell.com/support/kb/doc.php?id=3499891>
- ▶ Installing Windows from a USB flash drive:  
<http://technet.microsoft.com/en-us/library/dn293258.aspx>
- ▶ Formatting a USB flash drive to start the ESXi Installation or Upgrade:  
[https://pubs.vmware.com/vsphere-50/index.jsp?topic=%2Fcom.vmware.vsphere.install.doc\\_50%2FGUID-33C3E7D5-20D0-4F84-B2E3-5CD33D32EAA8.html](https://pubs.vmware.com/vsphere-50/index.jsp?topic=%2Fcom.vmware.vsphere.install.doc_50%2FGUID-33C3E7D5-20D0-4F84-B2E3-5CD33D32EAA8.html)

You can also use the ServerGuide Scripting Toolkit to create a bootable USB flash drive, as described in 6.1.3, “Preboot eXecution Environment” on page 229.

## 6.1.3 Preboot eXecution Environment

The Preboot eXecution Environment (PXE) is an environment to start computers by using a network interface for operating system deployment. All X6 systems support PXE.

For example, you can use xCAT software to deploy a broad set of operating systems by network. For more information, see this website:

<http://sourceforge.net/apps/mediawiki/xcat>

## 6.1.4 Lenovo XClarity Administrator

Lenovo XClarity Administrator can be used to manage the OS images repository and deploy operating system images to managed servers.

For information about Lenovo XClarity Administrator (including how to log on and start discovering), see 7.4, “Lenovo XClarity Administrator” on page 249.

Complete the following steps to mount and deploy an OS image from XClarity to a managed server:

1. Click **Provisioning** → **Deploy OS Images** under the Deploy Operating Systems section, as shown in Figure 6-7.

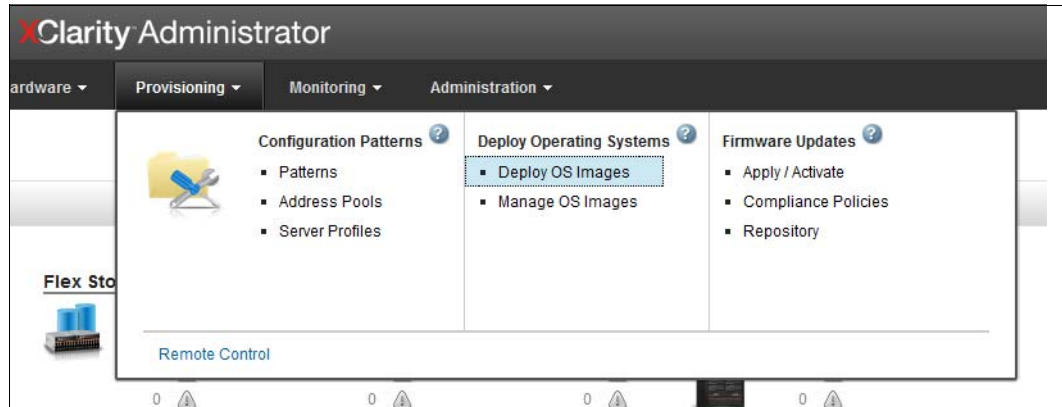


Figure 6-7 Deploy OS Images

2. On the Deploy Operating Systems: Deploy OS Images page, select the server and OS image to deploy and click **Deploy Images**, as shown in Figure 6-8 on page 230.



Figure 6-8 Server and OSS image selection

3. If you must change the OS image, select **Change Selected** → **Image to deploy**, as shown in Figure 6-9.

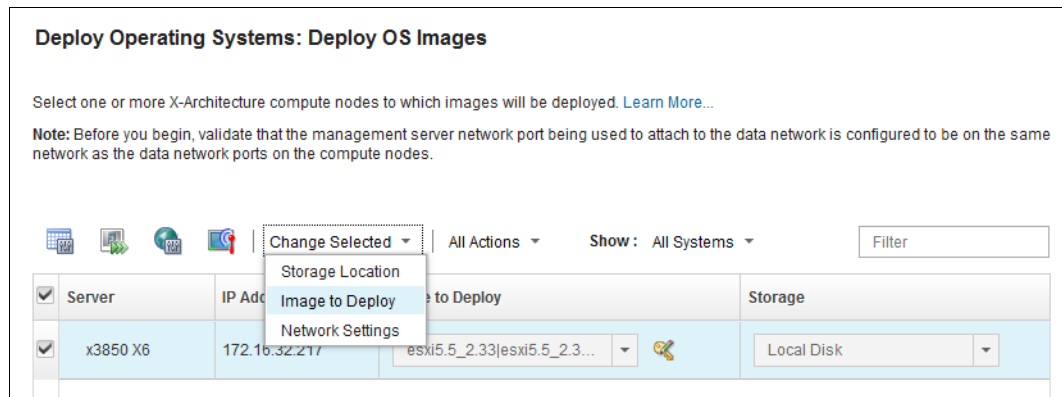


Figure 6-9 Change Selected drop-down menu



4. Choose the required OS image from list of available images, as shown in Figure 6-10.

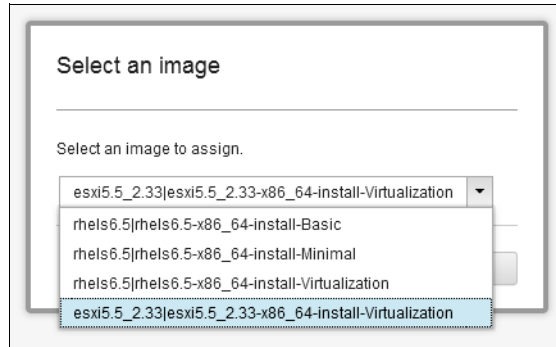


Figure 6-10 OS Image selection

5. You are prompted to set an administrator password when the OS is installed. Figure 6-11 on page 231 shows an example of deploying VMware ESXi 5.5.

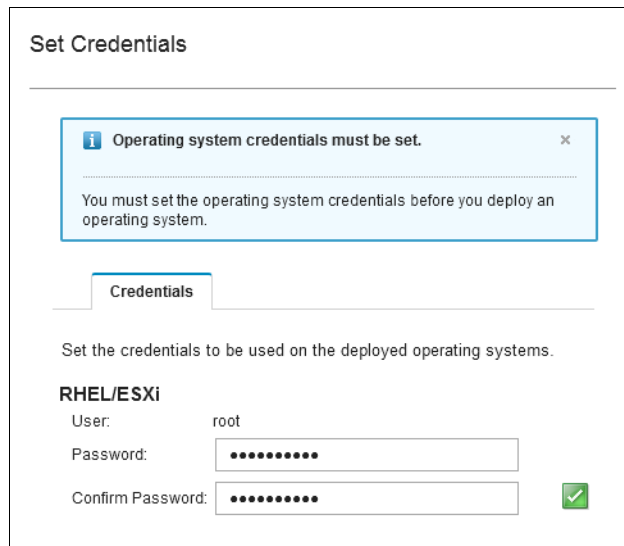


Figure 6-11 Root password request

6. When the administrator password is set, you must confirm the OS installation by clicking **Deploy**, as shown in Figure 6-12.





2. Mount the ServerGuide ISO image through the IMM2.
3. Start the server from ServerGuide.
4. After the files load, the Start window opens. Choose a language to continue.
5. Select your preferred keyboard layout and click **Next**.
6. Accept the license agreement and click **Next**.
7. The Welcome window opens and provides information about ServerGuide, which systems are supported, a readme file, and copyright and trademark information. Click **Next**.
8. In the next panel, you can import drivers that are not included in ServerGuide, as shown in Figure 6-15.

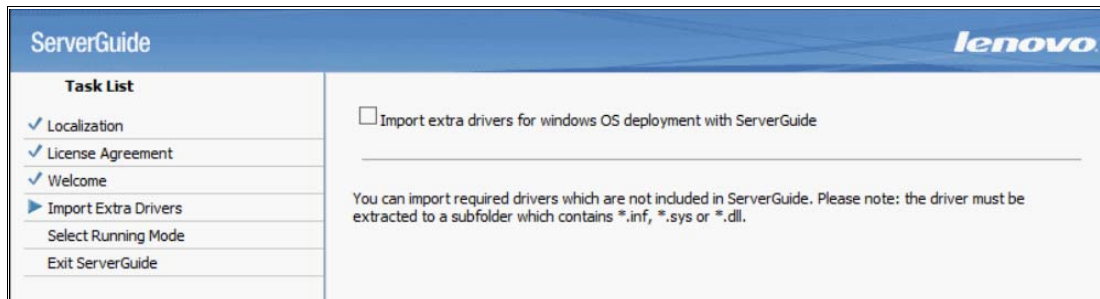


Figure 6-15 Import drivers

9. Select the operating system that you want to install and click **Next**, as shown in Figure 6-16.



Figure 6-16 Selecting the operating system

10. Enter the current date and time, as shown in Figure 6-17 and click **Next**.

<p><b>Task List</b></p> <ul style="list-style-type: none"> <li>✓ Localization</li> <li>✓ License Agreement</li> <li>✓ Welcome</li> <li>✓ Import Extra Drivers</li> <li>✓ Select Running Mode</li> <li>▶ Set Date and Time</li> <li>Configure RAID Adapters</li> <li>Create Partition</li> <li>Post Installation</li> <li>View Summary Report</li> <li>Operating System Setup</li> <li>Exit ServerGuide</li> </ul>	<div style="border: 1px solid black; padding: 5px;"> <p><b>Date</b></p> <p>Month: <input type="text" value="12"/> Day: <input type="text" value="4"/> Year: <input type="text" value="2013"/></p> </div> <div style="border: 1px solid black; padding: 5px; margin-top: 5px;"> <p><b>Time</b></p> <p><input checked="" type="radio"/> Twelve-hour clock  <input type="radio"/> Twenty-four hour clock</p> <p>Hour: <input type="text" value="7"/> Minute: <input type="text" value="27"/> Second: <input type="text" value="9"/> <input type="radio"/> AM <input checked="" type="radio"/> PM</p> </div> <p><b>Note:</b>  The service processor in your server has its own clock. After the NOS is installed, use your IBM systems-management software to set the service processor clock.</p>
---	--

Figure 6-17 Date and time settings

11. Create a RAID configuration. Select a RAID configuration and click **Next**, as shown in Figure 6-18.

<p><b>Task List</b></p> <ul style="list-style-type: none"> <li>✓ Localization</li> <li>✓ License Agreement</li> <li>✓ Welcome</li> <li>✓ Import Extra Drivers</li> <li>✓ Select Running Mode</li> <li>✓ Set Date and Time</li> <li>▶ Configure RAID Adapters</li> <li>▶ ServeRAID-M5210</li> <li>Create Partition</li> <li>Post Installation</li> <li>View Summary Report</li> <li>Operating System Setup</li> <li>Exit ServerGuide</li> </ul>	<p><b>Configure RAID Adapter</b></p> <p><b>Adapter Name:</b> ServeRAID-M5210  <b>Number of Drives:</b> 2 Physical Drive(s) (2 Good - 0 Bad)  <b>BIOS Version:</b> 6.12.00_4.12.05.00_0x06000500  <b>Firmware:</b> 4.200.21-2758</p> <hr/> <p><b>Select RAID Configuration:</b></p> <p><input type="radio"/> Keep Current Adapter Configuration  <b>Current Configuration:</b> [RAID 1 (Good) - 189781MB]  <b>Current Hotspare(s):</b> 0</p> <p><input type="radio"/> RAID Defaults</p> <p><input type="radio"/> RAID 0</p> <p><input checked="" type="radio"/> RAID 1</p> <p><input type="radio"/> RAID 10</p> <p><input type="radio"/> RAID 5</p> <p><input type="radio"/> RAID 1 + RAID 1</p> <p><input type="radio"/> RAID 1 + RAID 5</p> <p><input type="radio"/> RAID 6</p> <p><input type="radio"/> RAID 60</p> <p><input type="radio"/> RAID 50</p> <p><input type="checkbox"/> Add Hot Spare</p> <p style="color: red;"><b>Attention: If you click "Next", any existing partitions and data will be lost.</b></p> <p style="text-align: center;"> <input type="button" value="Advanced"/> <span style="margin-left: 200px;"><input type="button" value=" &lt;-Back"/></span> <span style="margin-left: 20px;"><input type="button" value=" Next-&gt;"/></span> <span style="margin-left: 20px;"><input type="button" value=" Exit"/></span> </p>
--	--

Figure 6-18 RAID selection

12.A Confirmation window opens that indicates that RAID configuration is complete, as shown in Figure 6-19. Click **Next**.

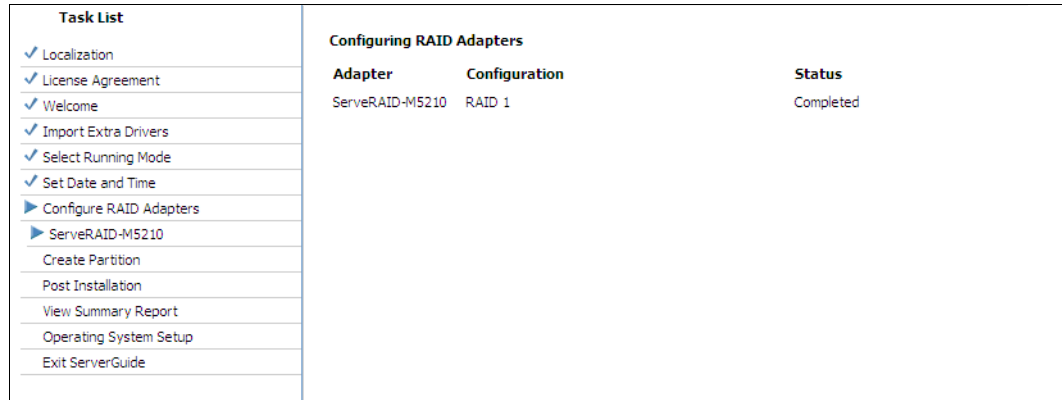


Figure 6-19 RAID confirmation panel

13.You must now create and format a partition. Choose your selection and click **Next** to start the process, as shown in Figure 6-20.

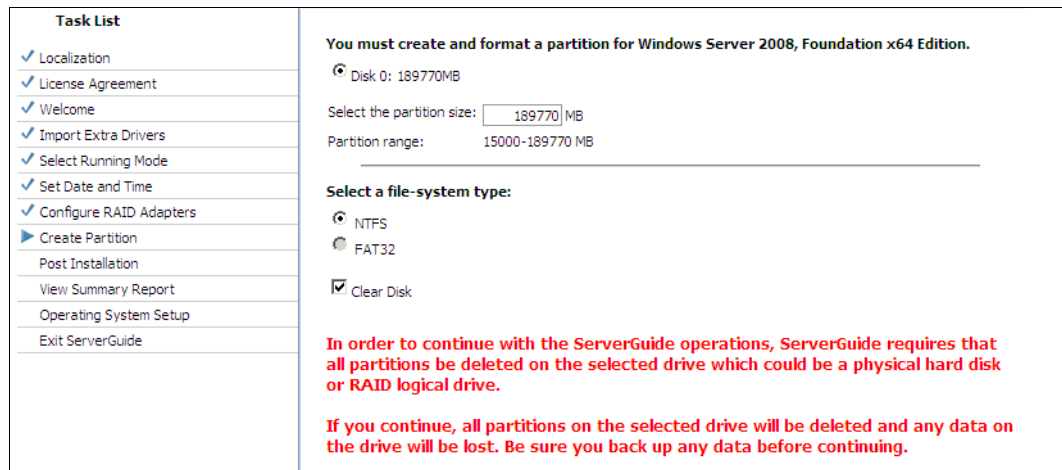


Figure 6-20 Selection for format and partition

14.When the process completes, click **Next**. You can select postinstallation options, as shown in Figure 6-21.

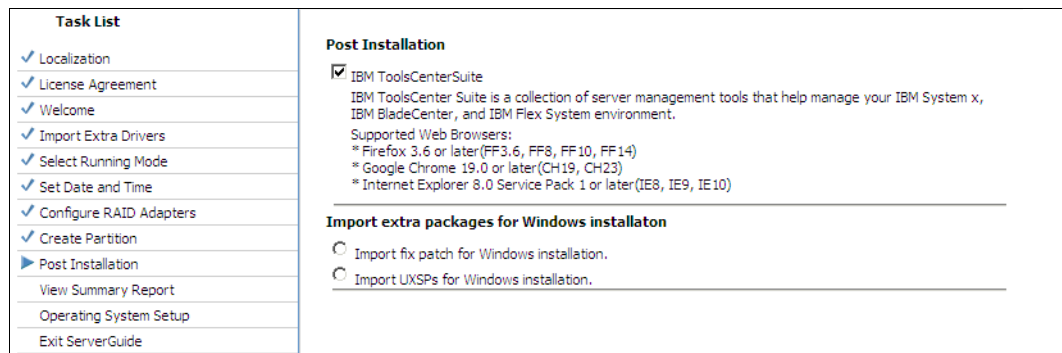


Figure 6-21 Postinstallation options

15. Review the configuration, as shown in Figure 6-22. Click **Next**.

Task List	Confirm your server configuration and hardware settings. Click "Next" to continue with your installation.	
	Item	Details
✓ Localization	Selected NOS	Windows Server 2008, Foundation x64 Edition
✓ License Agreement	System serial number	23Z1786
✓ Welcome	Machine type	3837
✓ Import Extra Drivers	System memory	1048576MB
✓ Select Running Mode	Processors, Processor slots	4,4
✓ Set Date and Time	Date and time	12/4/2013 7:30:36 PM
✓ Configure RAID Adapters	System BIOS version	1.00 (A8E103D) 11/12/2013
✓ Create Partition	NOS Partition	189770 MB - NTFS
✓ Post Installation		
▶ View Summary Report		
Operating System Setup		

Figure 6-22 Summary report

16. ServerGuide copies the necessary files to the disk in preparation for the operating system installation, as shown in Figure 6-23.

Task List	ServerGuide is copying files to your hard disk in preparation for the operating system installation. This may take up to five minutes...	
✓ Localization	✳ Copying Windows Server 2008, Foundation x64 Edition drivers...	
✓ License Agreement		
✓ Welcome		
✓ Import Extra Drivers		
✓ Select Running Mode		
✓ Set Date and Time		
✓ Configure RAID Adapters		

Figure 6-23 ServerGuide copying files

17. When the process is finished, click **Next**.

18. Mount your Windows ISO image via the IMM and click **Next**, as shown in Figure 6-24. ServerGuide searches for the disk.


Task List	Insert the Windows Server 2008, Foundation x64 Edition DVD and click "Next". (If you are using OS stored in the USB Key, just click "Next".)	
✓ Localization		
✓ License Agreement		
✓ Welcome		
✓ Import Extra Drivers		
✓ Select Running Mode		
✓ Set Date and Time		

Figure 6-24 Prompt to insert the OS DVD

19. The Windows setup installation procedure starts. Follow the Microsoft installation procedure to complete the installation of your OS.

## 6.3 ServerGuide Scripting Toolkit

You can use the ServerGuide Scripting Toolkit to create deployable images by using a collection of system-configuration tools and installation scripts. Versions of the ServerGuide Scripting Toolkit are available for the Windows Preinstallation Environment (PE) and Linux platforms.

By using the ServerGuide Scripting Toolkit, you can tailor and build custom hardware deployment solutions. It provides hardware configuration utilities and OS installation examples for System x and BladeCenter x86-based hardware.

If used with ServerGuide and UpdateXpress, the ServerGuide Scripting Toolkit provides a total solution for deploying System x and BladeCenter x86-based hardware in an unattended mode.

By using the ServerGuide Scripting Toolkit you can create a CD, DVD, or USB key that is used for starting that supports the following tasks and components:

- ▶ Network and mass storage devices
- ▶ Policy-based RAID configuration
- ▶ Configuration of system settings that uses Advanced Settings Utility (ASU)
- ▶ Configuration of fiber host bus adapters (HBAs)
- ▶ Local self-contained DVD deployment scenarios
- ▶ Local CD/DVD and network share-based deployment scenarios
- ▶ Remote Supervisor Adapter (RSA) II, IMM, and BladeCenter Management Module (MM)/Advanced Management Module (AMM) remote disk scenarios
- ▶ UpdateXpress System Packs installation that is integrated with scripted network operating system (NOS) deployment
- ▶ IBM Director Agent installation that is integrated with scripted NOS deployment

The ServerGuide Scripting Toolkit, Windows Edition supports the following versions of IBM Systems Director Agent:

- ▶ Common Agent 6.1 or later
- ▶ Core Services 5.20.31 or later
- ▶ Director Agent 5.1 or later

The Windows version of the ServerGuide Scripting Toolkit enables automated operating system support for the following Windows operating systems:

- ▶ Microsoft Windows Server 2012
- ▶ Microsoft Windows Server 2012 R2
- ▶ Windows Server 2008, Standard, Enterprise, Datacenter, and Web Editions
- ▶ Windows Server 2008 x64, Standard, Enterprise, Datacenter, and Web Editions
- ▶ Windows Server 2008, Standard, Enterprise, and Datacenter Editions without Hyper-V
- ▶ Windows Server 2008 x64, Standard, Enterprise, and Datacenter without Hyper-V
- ▶ Windows Server 2008 R2 x64, Standard, Enterprise, Datacenter, and Web Editions

The Linux version of the ServerGuide Scripting Toolkit enables automated operating system support for the following operating systems:

- ▶ SUSE Linux Enterprise Server 10 SP2 and later
- ▶ SUSE Linux Enterprise Server 11
- ▶ Red Hat Enterprise Linux 6 U1 and later
- ▶ Red Hat Enterprise Linux 5 U2 and later

Automated deployment of the following NOSes in Native uEFI mode is available:

- ▶ SUSE Linux Enterprise Server 11 SP1
- ▶ Red Hat Enterprise Linux 6 x64

To download the Scripting Toolkit or the *ServerGuide Scripting Toolkit User's Reference*, see this web page:

<https://support.lenovo.com/us/en/documents/LNV0-T00LKIT>



## 6.4 Use of embedded VMware ESXi

The x3850 X6 and x3950 X6 servers support a USB flash drive option that is preinstalled with VMware ESXi. VMware ESXi is fully contained on the flash drive, and does not require any disk space. The USB Memory Key for VMware Hypervisor plugs into the internal USB port that is on system board of Primary I/O Book.

For more information about supported options, see 3.22, “Integrated virtualization” on page 120.

VMware ESXi supports starting from the Unified Extensible Firmware Interface (UEFI). To ensure that you can start ESXi successfully, you must change the start order. The first start entry must be Embedded Hypervisor. Complete the following steps:

1. Press F1 for the UEFI Setup.
2. Select **Boot Manager** → **Add Boot Option**.
3. Select **Generic Boot Option**, as shown in Figure 6-25.




Figure 6-25 Boot options

4. Select **Embedded Hypervisor**, as shown in Figure 6-26. If either option is not listed, the option is in the boot list. When you finish, press Esc to go back one panel.



Figure 6-26 Add boot options

5. Select **Change Boot Order**.

6. Change the boot order to Embedded Hypervisor, as shown in Figure 6-27.




Figure 6-27

7. Select **Commit Changes** and press Enter to save the changes.

## 6.5 Booting from SAN

Boot from SAN (or SAN Boot) is a technique that is used when the node in the chassis has no local disk drives. It uses an external storage system logical unit number (LUN) to start the operating system. The OS and data are on the SAN. This technique is commonly used to provide higher availability and better usage of the systems storage (where the operating system is installed). Hot-spare nodes or “Rip-n-Replace” techniques also can be easily implemented by using boot from SAN.

To successfully implement SAN Boot, the following conditions must be met:

- ▶ Storage system supports SAN Boot
- ▶ Operating system supports SAN Boot
- ▶ FC HBAs or iSCSI initiators support SAN Boot

Consider the following useful guidelines for booting from SAN:

- ▶ Check if UEFI recognizes the adapter. Select **UEFI** → **System Settings** → **Adapters and UEFI Drivers**. The Adapters and UEFI Drivers panel opens, as shown in Figure 6-28. The “Card - HBA” option must be listed. If the option is not listed, reflash the UEFI, IMM, and Firmware of the HBA and check again.

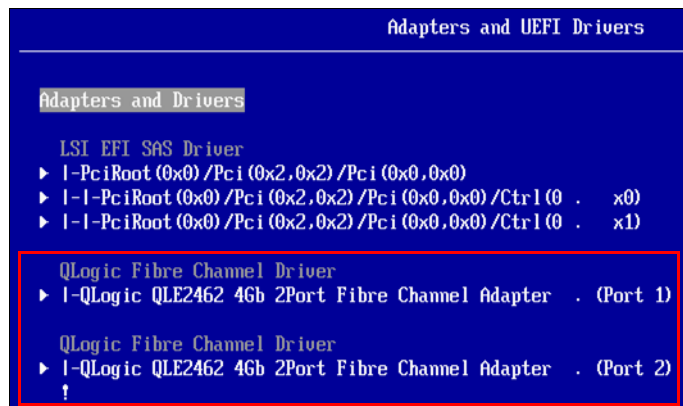


Figure 6-28 Adapters that are listed in UEFI

- ▶ If you do not have internal drives, disable the onboard SAS RAID Controller by selecting **System Settings** → **Devices and IO ports** → **Enable/Disable Onboard Devices** and disabling the SAS Controller or Planar SAS.
- ▶ Set the HBA as the first device in the Option ROM Execution Order by selecting **System Settings** → **Devices and IO Ports** → **Set Option ROM Execution Order**.
- ▶ For older operating systems that do not support UEFI, set **Legacy Only** as the first boot device.
- ▶ Remove all devices from the boot order that might not host an OS. The optimal minimum configuration is CD/DVD and Hard Disk 0. For older operating systems only, set **Legacy Only** as the first boot device.
- ▶ Enable the BIOS from your HBA.
- ▶ Verify that your HBA can see a LUN from your storage.
- ▶ For Microsoft Windows installations, ensure that the LUN is accessible through only one path (Zoning or LUN masking).
- ▶ After installation, remember to install the multipath driver *before* you set more than one path if you have more than one path to the LUN.

You can also check the documentation for the operating system that is used for Boot from SAN support and requirements and storage vendors. For more information about SAN boot, see the following resources:

- ▶ *Red Hat Enterprise Linux 7 Installation Guide:*  
[https://access.redhat.com/documentation/en-US/Red\\_Hat\\_Enterprise\\_Linux/7/html/Installation\\_Guide/sect-storage-devices-x86.html](https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Installation_Guide/sect-storage-devices-x86.html)
- ▶ *Red Hat Enterprise Linux 6 Installation Guide:*  
[https://access.redhat.com/site/documentation/en-US/Red\\_Hat\\_Enterprise\\_Linux/6/html/Installation\\_Guide/Storage\\_Devices-x86.html](https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Installation_Guide/Storage_Devices-x86.html)
- ▶ *Windows Boot from Fibre Channel SAN – Overview and Detailed Technical Instructions for the System Administrator:*  
<http://www.microsoft.com/download/en/details.aspx?displaylang=en&id=2815>
- ▶ *vSphere Storage* document from VMware:  
<http://pubs.vmware.com/vsphere-55/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcen-ter-server-55-storage-guide.pdf>
- ▶ IBM Redbooks publication, *SAN Boot Implementation and Best Practices Guide for IBM System Storage*, SG24-7958:  
<http://www.redbooks.ibm.com/abstracts/sg247958.html?Open>

For IBM System Storage compatibility information, see the IBM System Storage Interoperability Center at this website:

<http://www.ibm.com/systems/support/storage/config/ssic>



# Managing X6 hardware

In the information technology sector, server systems management received greater focus over recent years. Maintaining and managing systems efficiently is essential to business IT operations. In this chapter, we describe the embedded hardware and external applications that are available to manage and maintain the X6 systems. Several of these applications also are demonstrated.

This chapter includes the following topics:

- ▶ 7.1, “Introduction” on page 244
- ▶ 7.2, “Integrated Management Module II” on page 244
- ▶ 7.3, “Remote control” on page 248
- ▶ 7.4, “Lenovo XClarity Administrator” on page 249
- ▶ 7.5, “Lenovo XClarity integrators” on page 258
- ▶ 7.6, “Lenovo XClarity Energy Manager” on page 265
- ▶ 7.7, “Upward Integration Module for Zenoss” on page 288
- ▶ 7.8, “Advanced Settings Utility” on page 289
- ▶ 7.9, “MegaRAID Storage Manager” on page 292
- ▶ 7.10, “Managing NVMe drives and software RAID recovery” on page 295
- ▶ 7.11, “Serial over LAN” on page 312

## 7.1 Introduction

Lenovo provides several tools to successfully deploy, manage, and maintain the x3850 X6 and x3950 X6 servers. The collective name for these tools is Lenovo ToolsCenter™. These tools are available at the following website:

<http://support.lenovo.com/documents/LNVO-CENTER>

These tools are grouped in the following high-level categories:

- ▶ System deployment
- ▶ System configuration
- ▶ System updates
- ▶ System diagnosis

The tools are listed in Table 7-1. In this section, we describe the use of several of these tools. Firmware deployment tools are also described in 5.8, “Updating firmware” on page 211.

Table 7-1 Available tools to manage the X6 systems

Tool	Firmware deployment	OS installation assistance	System management	Problem diagnosis
Integrated Management Module II	Yes <sup>a</sup>	No	Yes	Yes
Dynamic System Analysis™	No	No	No	Yes
Bootable Media Creator	Yes	No	No	No
Scripting Toolkit	Yes	Yes	No	No
ServerGuide	No	Yes	No	No
Update Express System Pack Installer	Yes	No	No	No
IBM Systems Director	Yes	Yes <sup>b</sup>	Yes	Yes
Lenovo XClarity Administrator	Yes	Yes	Yes	Yes
Storage Configuration Manager™	Yes	No	Yes	Yes
Remote Control	No <sup>c</sup>	No <sup>c</sup>	Yes	No
Advanced Settings Utility	No <sup>c</sup>	No <sup>c</sup>	Yes	No
MegaRAID Storage Manager	No	No	Yes <sup>d</sup>	Yes
Serial over LAN	No	No	Yes	No

a. You can update only Integrated Management Module (IMM2) firmware by using the IMM2 web browser interface.

b. Only when the IBM Tivoli Provision Management for OS Deployment IBM Systems Director edition is installed.

c. By using this tool, the media that contains the firmware or operating system can be mounted.

d. This tool provides the management of RAID controllers and hard disks only.

## 7.2 Integrated Management Module II

The Integrated Management Module II (IMM2.1, referred to here simply as IMM2) offers the following overall features and functions:

- ▶ Provides diagnostics, virtual presence, and remote control to manage, monitor, troubleshoot, and repair from anywhere.

- ▶ Securely manages servers remotely and independently of the operating system state.
- ▶ Helps remotely configure and deploy a server from bare metal.
- ▶ Auto-discovers the scalable components, ports, and topology.
- ▶ Provides one IMM2 firmware for a new generation of servers.
- ▶ Helps system administrators easily manage large groups of diverse systems.
- ▶ Requires no special drivers.
- ▶ Works with Lenovo XClarity to provide secure alerts and status, which helps reduce unplanned outages.
- ▶ Uses standards-based alerting, which enables upward integration into various enterprise management systems.

The following methods are available to manage an IMM2:

- ▶ **Out-of-band management:** All management tasks are passed directly to the systems IMM2 via a network connection. No drivers are required for the IMM2 because it is configured with its own IP address and connected directly to the network.
- ▶ **In-band management:** All management tasks are passed to the system by using the operating system that is installed on it. The tasks can apply to the operating system or to the IMM2 that is installed on the system. If the tasks are to be passed to the IMM2, the relevant operating system driver must be installed for the IMM2.

In the following section, we describe the out-of-band and in-band initial configuration.

## 7.2.1 Configuring IMM2 for out-of-band-management

The IMM2 for the X6 server can be managed over one of the following ports:

- ▶ The dedicated system management port on the rear of each chassis.  
By using this dedicated port, the IMM2 can be connected to an isolated management network for improved security. This port is the default port.
- ▶ Port 1 on the installed ML2 Ethernet adapter. The IMM2 shares the ML2 adapter slot for access to one of the adapters Ethernet ports with the operating system.

The use of a port that is shared between IMM2 and the operating system eliminates the need for an extra port on your network switch. For more information, see 5.1.4, “IMM2 dedicated versus shared ML2 Ethernet port” on page 159.

Figure 7-1 shows the location of the IMM2 and ML2 ports on the rear of the server.




Figure 7-1 Ports over which the IMM2 can be managed out-of-band on X6.

By default, the ML2 slot is not shared for use with the IMM2. You must enable this feature in the Unified Extensible Firmware Interface (UEFI) of the server.

Complete the following steps to enable the IMM2 to use the ML2 adapter:

1. Ensure that you have an ML2 Ethernet adapter installed.
2. Start the server and press F1 when prompted.
3. Select **System Settings** from the System Configuration and Boot Management menu.
4. Select **Integrated Management Module** from the System Settings menu.
5. Select **Network Configuration** from the IMM. By using this menu, you can configure the network settings for the IMM2. Also, you can configure the IMM2 to share the use of ML2 or out-of-band management from this menu, as shown in Figure 7-2.

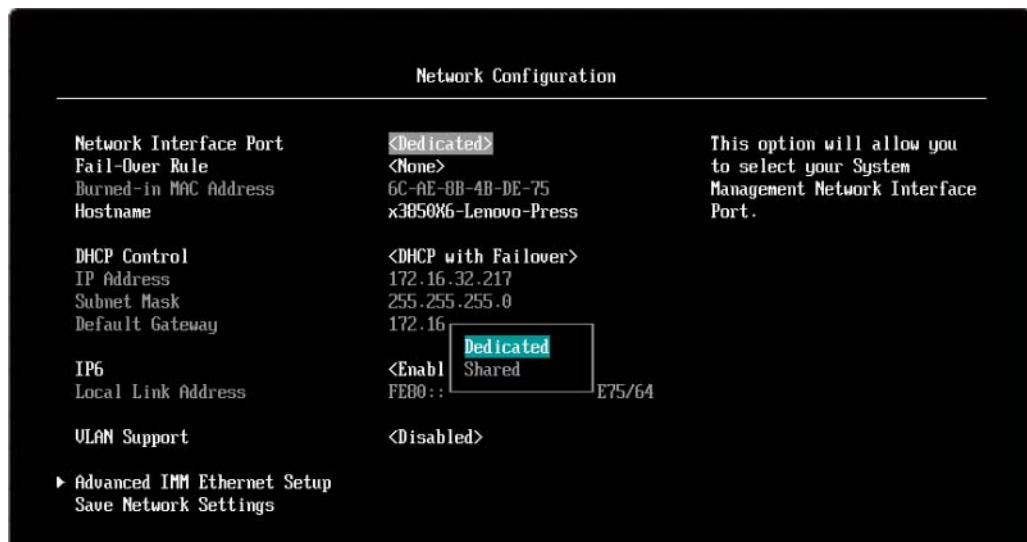


Figure 7-2 IMM2 Network Configuration panel

6. Set the Network Interface Port setting to Shared to allow the IMM2 to use the ML2 adapter.



7. For DHCP Control, choose **Static IP**.
8. For IP Address, enter the relevant IP address.
9. For Subnet Mask, enter the required subnet mask.
10. For Default Gateway, enter the required default gateway address.
11. When you complete the IP address configuration, press Esc three times to return to the System Configuration and Boot Management menu.
12. For Exit Setup, press the Y key when prompted to save and exit the Setup utility. The server restarts with the new settings.
13. Plug a network cable into the dedicated system management port or the ML2 adapter if you set the IMM2 to share its use according to the instructions. Ensure that you can ping the IP address of the IMM2 on the connected network port.

After the IMM2 is available in the network, you can log in to the IMM2 web interface by entering its IP address in a supported web browser, as shown in Figure 7-3.

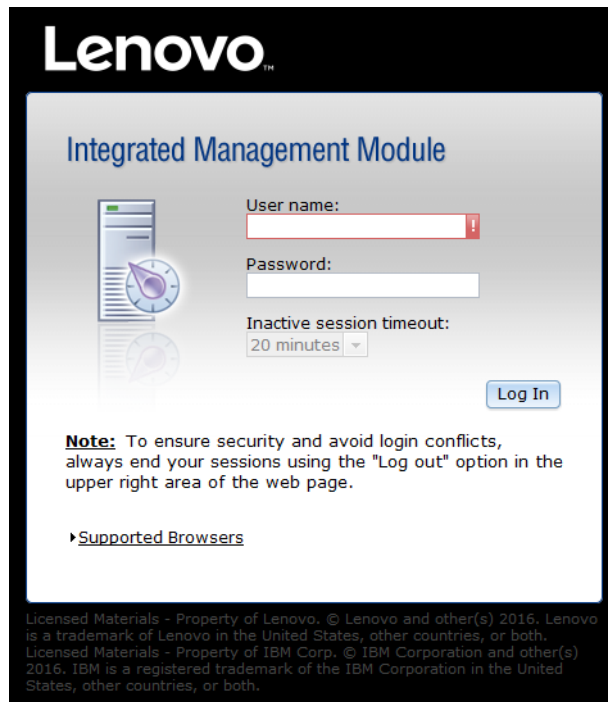


Figure 7-3 IMM2 Login web page

Enter the default user name USERID. This user name is case-sensitive. Enter the default password PASSWORD, in which 0 is the number zero.

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

<https://support.lenovo.com/docs/UM103336>

## 7.2.2 Configuring IMM2 in-band configuration

Managing an IMM2 in-band means managing the IMM2 through the operating system. For example, IBM Systems Director can update the firmware of an IMM2 via the operating system. The benefit of this approach is that you do not need to configure the IMM2 with its own dedicated IP address if there are insufficient IP addresses to allocate or you do not want to provide the IMM2 with its own IP address.

There is no actual configuration that is required within the IMM2 web interface so that the IMM2 can be managed in-band. However, you must ensure that the prerequisite drivers are installed so that the operating system can recognize the IMM2. All supported versions of Microsoft Windows Server 2008 and 2012, 2012 R2, VMware ESX, and Linux now include the prerequisite drivers for the X6 systems.

For more information about the supported operating systems, see the *Lenovo Operating System Interoperability Guide*, located here:

<http://lenovopress.com/osig>

## 7.3 Remote control

Controlling the X6 systems remotely can be done by using the IMM2.

You can perform the following common tasks with the remote control function:

- ▶ Control the power of the systems.
- ▶ Mount remote media, which includes CD/DVD-ROMs, supported ISO and firmware images, and USB devices.
- ▶ Create your own customized keyboard key sequences b using the soft key programmer.
- ▶ Customize your viewing experience.

### 7.3.1 Accessing the remote control feature in the IMM2

Complete the following steps to use the Remote Control feature on the IMM2 interface for the X6 system:

1. Log in to the IMM2 of the specific system that you want to control.
2. On the System Status page, select **Remote Control**, as shown in Figure 7-4.



Figure 7-4 Remote control from the IMM2

- To protect sensitive disk and KVM data during your session, click **Encrypt disk and KVM data during transmission** before starting Remote Control. For complete security, use Remote Control with SSL. You can configure SSL by selecting **IMM2 Management** → **Security** from the top menu.
- If you want exclusive remote access during your session, click **Start Remote Control in Single User Mode**. If you want to allow other users remote console (KVM) access during your session, click **Start Remote Control in Multi-user Mode**, as seen in Figure 7-5.

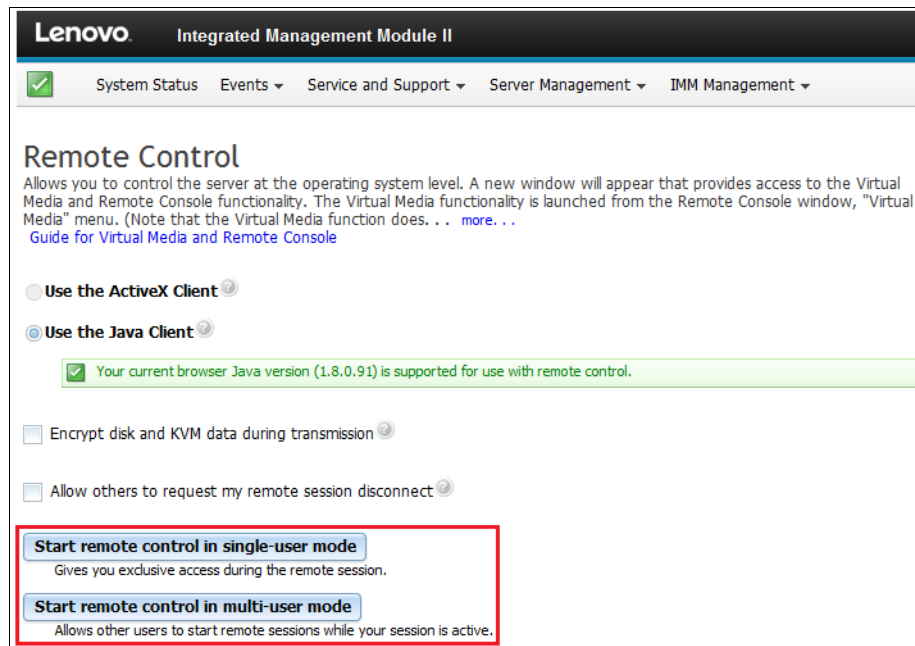


Figure 7-5 IMM2 Remote Control page

For more information about the various controls that are available to control the server, see *User's Guide for Integrated Management Module II*, which is available here:

<https://support.lenovo.com/docs/UM103336>

By remote controlling your server, you also can remotely mount an ISO image. For more information about how to remote mount an ISO image to your server, see 6.1, “Installing an OS without a local optical drive” on page 226.

## 7.4 Lenovo XClarity Administrator

Lenovo XClarity Administrator is a centralized resource management solution that reduces complexity, speeds response, and enhances the availability of Lenovo server systems and solutions.

XClarity Administrator provides agent-free hardware management for System x rack servers, including x3850 X6 and x3950 X6 servers and Flex System compute nodes and components, including the Chassis Management Module and Flex System I/O modules.

XClarity Administrator is a virtual appliance that is quickly imported into a Microsoft Hyper-V or VMware virtualized environment, which gives easy deployment and portability. The tool offers out-of-band agentless management to reduce complexity, which means that the

endpoints that are managed do not need special software agents or drivers to be installed or maintained.

XClarity Administrator offers the following functions:

- ▶ Discovery
- ▶ Inventory
- ▶ Monitoring
- ▶ Firmware updates
- ▶ Firmware compliance
- ▶ Configuration management
- ▶ Deployment of operating systems and hypervisors to bare metal servers
- ▶ Call home service

The administration dashboard is based on HTML 5. Such a foundation allows fast location of resources so tasks can be run quickly.

In this section, we describe key functions and tasks of XClarity Administrator that are relevant to the X6 servers. This section includes the following topics:

- ▶ 7.4.1, “X6 considerations” on page 250
- ▶ 7.4.2, “Discovering the IMM2 of an x3850 X6” on page 250
- ▶ 7.4.3, “Remote control” on page 253
- ▶ 7.4.4, “Hardware monitoring” on page 254
- ▶ 7.4.5, “Firmware updates” on page 256
- ▶ 7.4.6, “Operating system deployment” on page 257

## 7.4.1 X6 considerations

Consider the following points regarding the use of Lenovo XClarity Administrator with the x3850 X6 and x3950 X6:

- ▶ Only machine type 6241 of the x3850 X6 and x3950 X6 is supported. Machine type 3837 does not support XClarity Administrator.
- ▶ XClarity Administrator Fix Pack 1 is required to perform firmware updates of the x3850 X6 and x3950 X6.
- ▶ For the x3850 X6, the IMM2 must be connected to the network. For the x3950 X6, both IMM2 ports must be connected to the network.
- ▶ For the x3950 X6, you must set up automatic or custom partitioning by using the IMM before deploying a server pattern to the partition.

## 7.4.2 Discovering the IMM2 of an x3850 X6

After your IMM2 is configured for out-of-band management (see 7.2.1, “Configuring IMM2 for out-of-band-management” on page 245), you can discover the system from within XClarity Administrator. After the IMM2 is added to the XClarity Administrator console, you can perform management tasks against it.

Complete the following steps to discover and add the IMM2 to the XClarity Administrator console:

1. Log in to the XClarity Administrator web interface by browsing to the following website, where *servername* is the name or IP address of the virtual machine where XClarity Administrator is running:

`https://servername`

For example:

<https://xclarity-demo.lenovo.com>

<https://172.16.32.220>

After you log in, the Dashboard page is displayed, as shown in Figure 7-6.

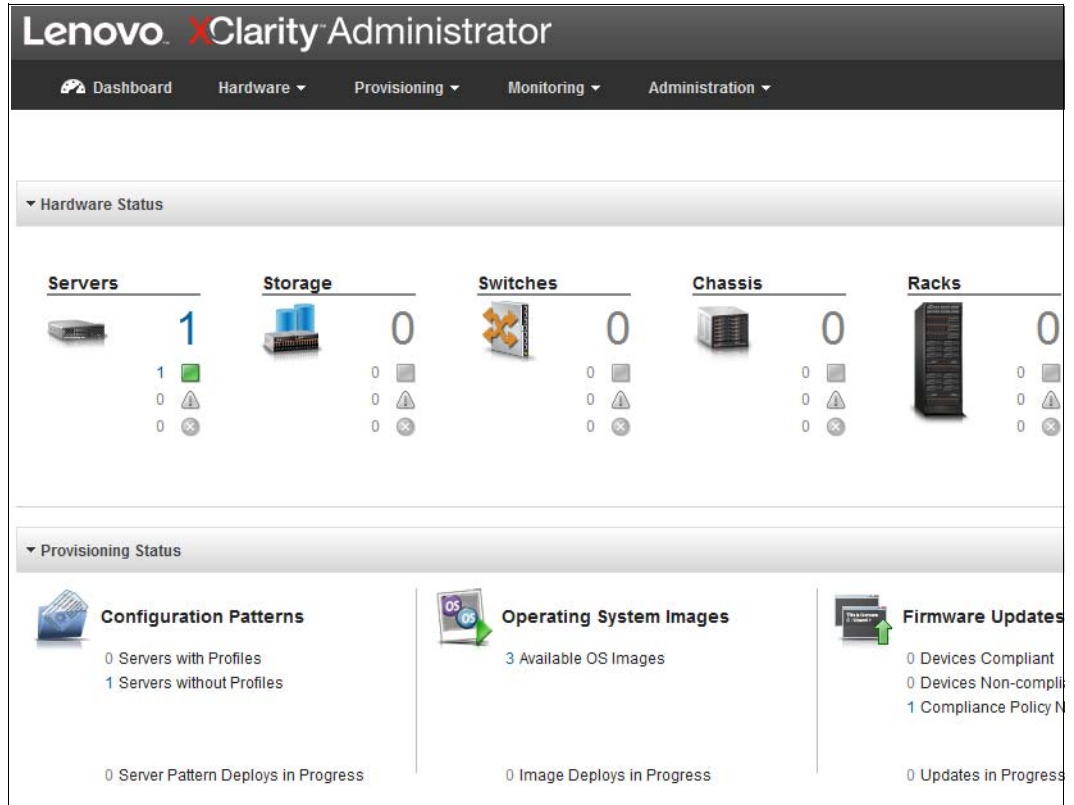


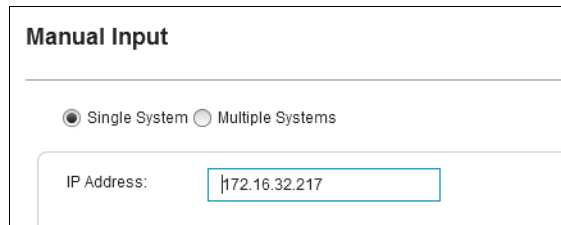
Figure 7-6 Lenovo XClarity Administrator Dashboard

2. Select **Hardware** → **Discover and Manage New Devices**, as shown in Figure 7-7.



Figure 7-7 Discover and Manage New Devices

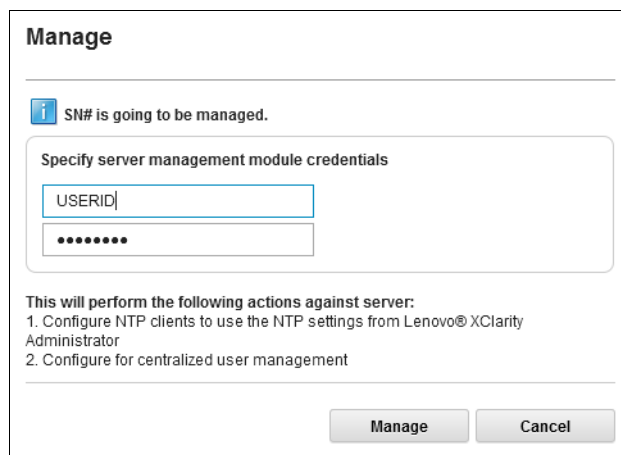
3. Enter the IP address of the IMM2 in the space that is provided under the IP address, as shown in Figure 7-8.



The 'Manual Input' dialog box features a title bar with the text 'Manual Input'. Below the title bar, there are two radio buttons: 'Single System' (which is selected) and 'Multiple Systems'. Underneath, there is a label 'IP Address:' followed by a text input field containing the IP address '172.16.32.217'.

Figure 7-8 IMM2 IP input

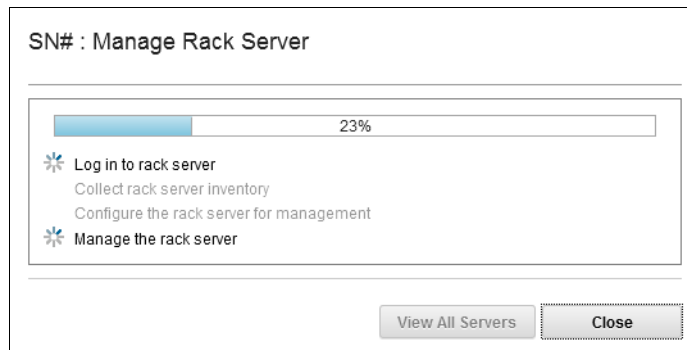
4. You also must specify the IMM2 credentials to gain access to the IMM2 and click **Manage**, as shown in Figure 7-9.



The 'Manage' dialog box has a title bar with the text 'Manage'. Below the title bar, there is an information icon and the text 'SN# is going to be managed.'. A section titled 'Specify server management module credentials' contains two input fields: the first is labeled 'USERID' and the second is a password field with masked characters. Below this section, there is a list of actions: 'This will perform the following actions against server:' followed by two numbered items: '1. Configure NTP clients to use the NTP settings from Lenovo® XClarity Administrator' and '2. Configure for centralized user management'. At the bottom right, there are two buttons: 'Manage' and 'Cancel'.

Figure 7-9 IMM2 credentials

A Discovery progress bar is shown, as shown in Figure 7-10.



The 'SN# : Manage Rack Server' dialog box features a title bar with the text 'SN# : Manage Rack Server'. Below the title bar, there is a progress bar showing 23% completion. Underneath, there is a list of tasks with asterisk icons: 'Log in to rack server' (with sub-items 'Collect rack server inventory' and 'Configure the rack server for management') and 'Manage the rack server'. At the bottom right, there are two buttons: 'View All Servers' and 'Close'.

Figure 7-10 Discovery progress bar

When the discovery process finishes and XClarity Administrator discovers the IMM2, you see a confirmation message, as shown in Figure 7-11.

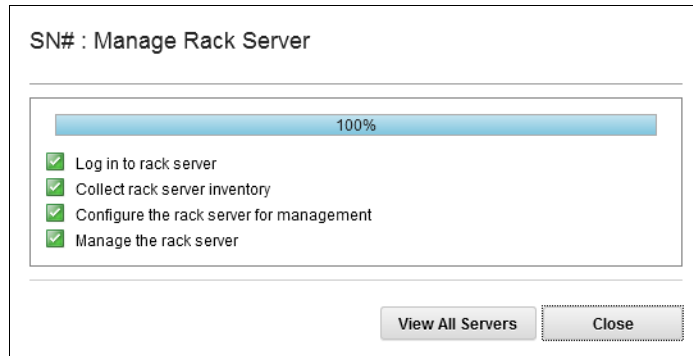


Figure 7-11 Discovery process ending

5. Click **Hardware** → **All Servers** to see the discovered IMM2 that is now listed, as shown in Figure 7-12.

Server	Status	Power	IP Addresses	Rack Name/Unit	Chassis/Bay	Product Name
imm2-x3850x6-brdwell.poc.bts.lab	Normal	On	172.30.30.152, 169.254.9...	Unassigned / ...	Unassigned /	

Figure 7-12 Discovered servers list

From this point, you can perform management tasks against the X6 server.

### 7.4.3 Remote control

XClarity Administrator provides remote control features for all managed hardware, including x3850 X6 and x3950 X6 servers.

Complete the following steps to open a remote control session to the server:

1. Select **Hardware** → **All Servers**, select the server from the list, and click **All Actions** → **Launch Remote Control**, as shown in Figure 7-13 on page 254.

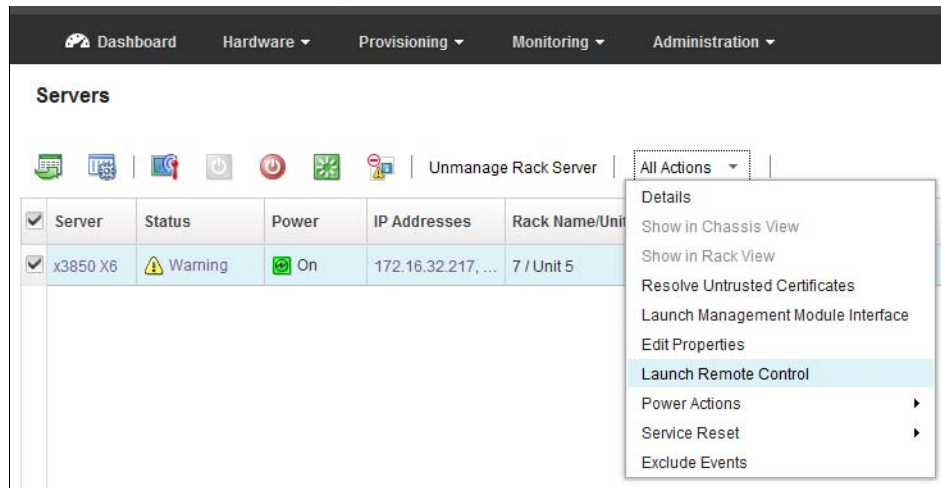


Figure 7-13 Launch Remote Control

2. The remote control Java applet appears, which provides virtual keyboard, video, and mouse to the server, as seen in Figure 7-14.

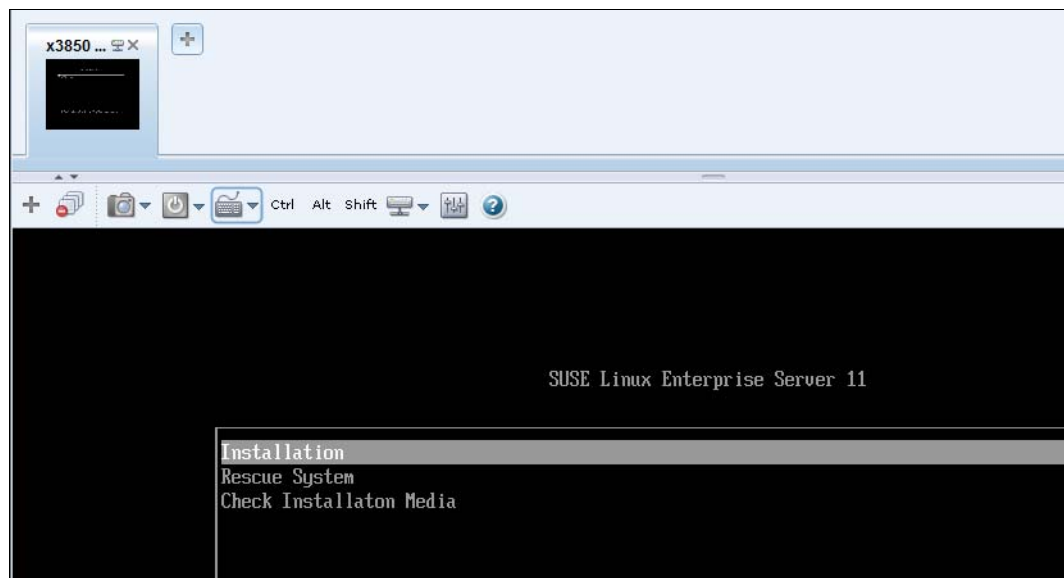


Figure 7-14 Remote control applet

With the remote control feature, you can control the power to the systems, mount remote media (including CD/DVD-ROMs, ISO-images, and USB devices), create your own customized keyboard key sequences by using the soft key programmer, and open new remote sessions to other servers that are managed by XClarity Administrator.

## 7.4.4 Hardware monitoring

XClarity Administrator provides a centralized view of events and alerts that are generated from managed endpoints, such as X6 servers. When an issue is detected by the IMM2, an event is passed to the XClarity Administrator. That event is then displayed in the alerts list that is available within the user interface.



To open the alerts page of the managed server, select **Hardware** → **All Servers**, select the server from the list, and click **Alerts** in the Status and Health section, as shown in Figure 7-15.

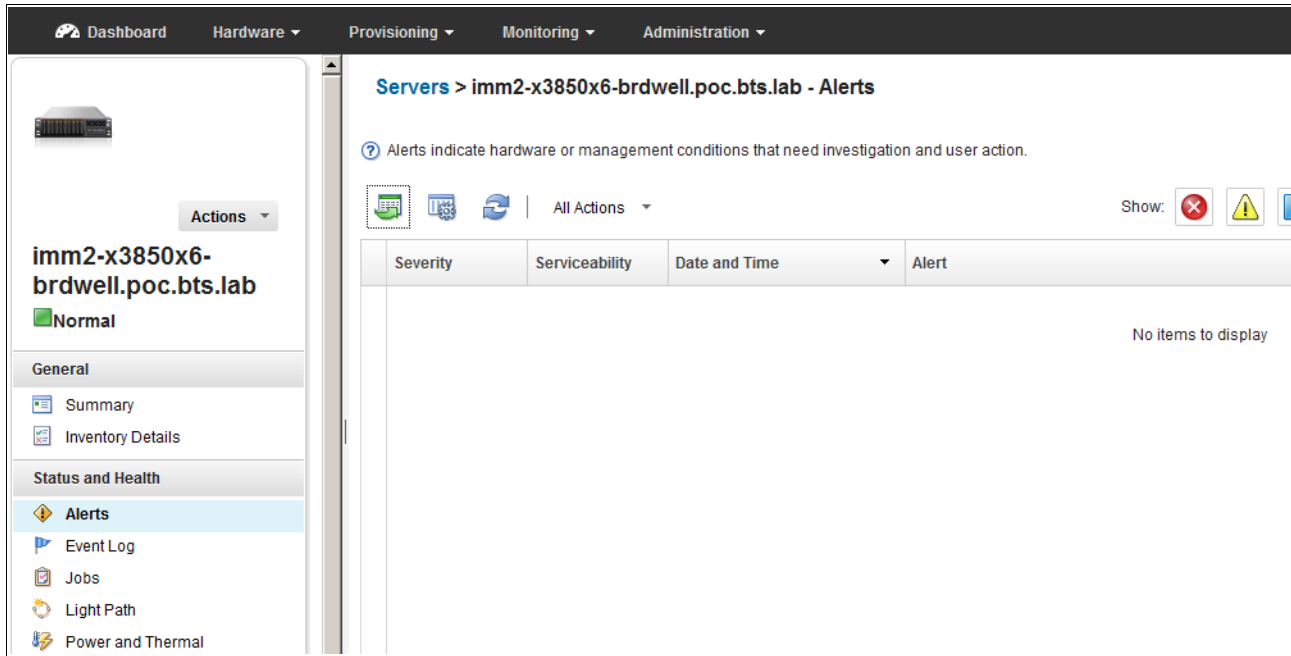


Figure 7-15 Alerts

You can also open the system event log of the server, check light path states, or power consumption. For example, to open Power Consumption History graph, browse to Power and Thermal of the Status and Health section, as shown in Figure 7-16.

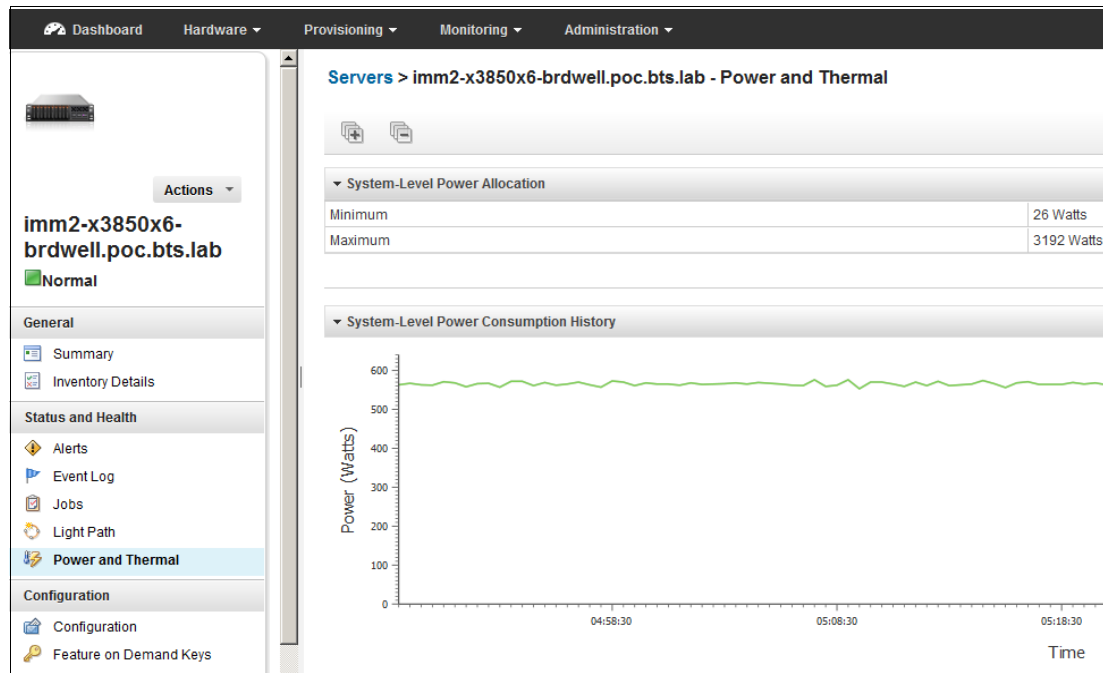


Figure 7-16 Power and Thermal

## 7.4.5 Firmware updates

Within XClarity Administrator, you can manage the firmware updates repository and apply and activate firmware updates for all managed endpoints. Compliance policies can be started to flag managed endpoints that do not comply with the defined firmware rules. Refreshing the repository and downloading updates requires an Internet connection. If XClarity Administrator has no Internet connection, you can manually import updates to the repository.

Complete the following steps to update the firmware of one of servers that is managed by Lenovo XClarity Administrator:

1. Click **Provisioning** → **Apply / Activate** under the Firmware Updates section, as shown in Figure 7-17.

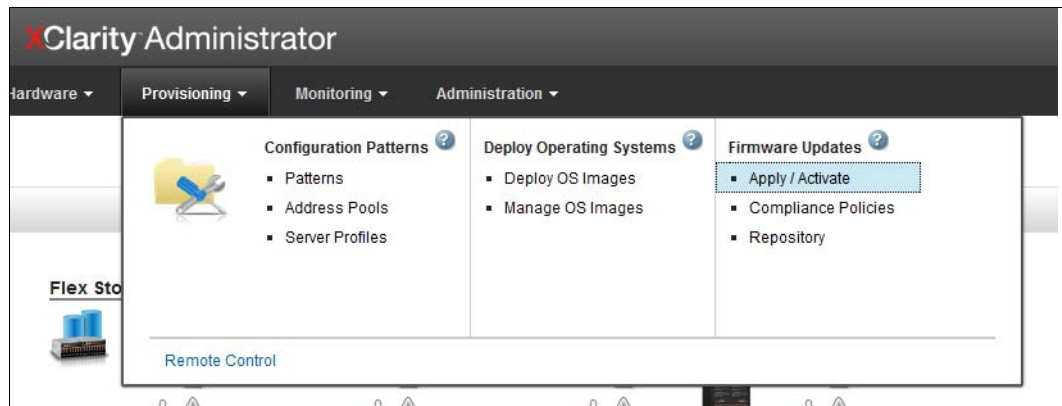


Figure 7-17 Firmware Update page selection

2. In the Firmware Updates: Apply / Activate page, select the required hardware components to update, and click the **Perform Updates** icon, as shown in Figure 7-18.

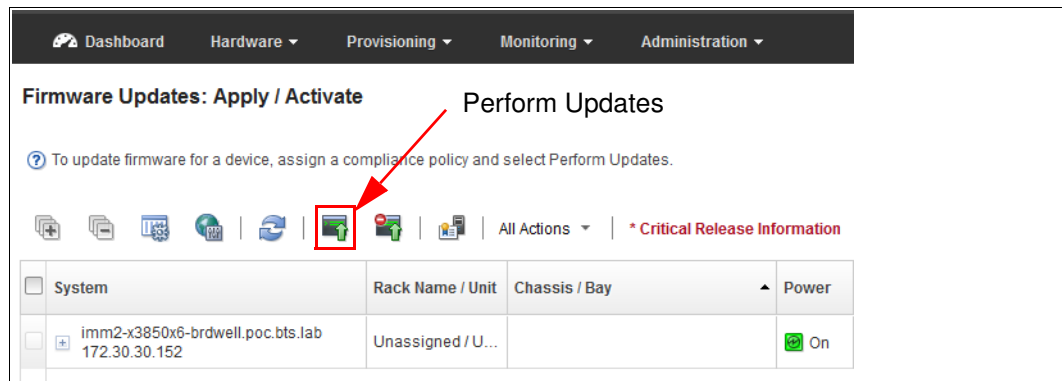


Figure 7-18 Firmware update selection

3. An Update Summary window opens, in which you can set Update Rules and Activation Rules. As shown in Figure 7-19 on page 257 the Update Rules is set to “Stop all updates on error” and Activation Rules is set to “Immediate activation”, which means an immediate server restart is needed for new firmware activation.

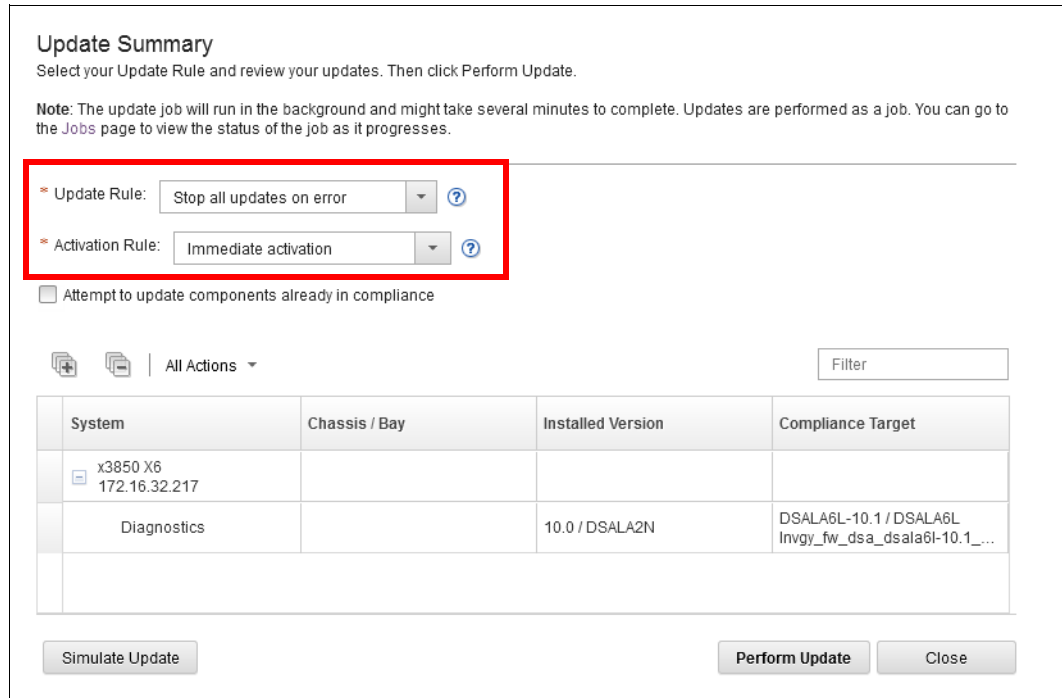


Figure 7-19 Firmware update parameters

- To start flashing, click **Perform Update**. When the Immediate activation option is chosen, you must confirm the operation by clicking **OK**, as shown in Figure 7-20.

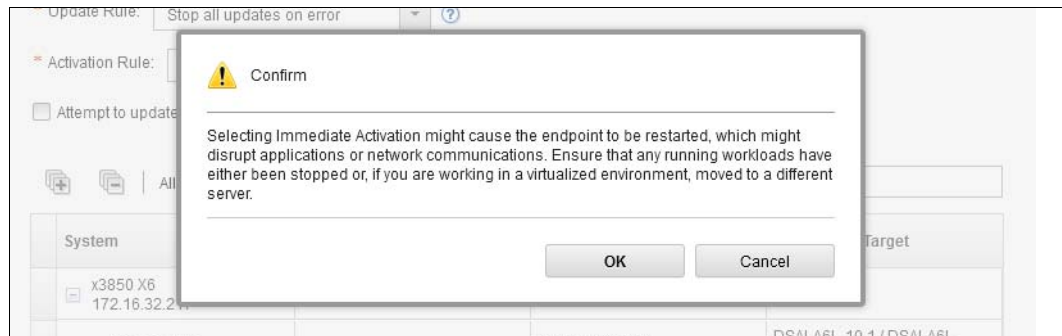


Figure 7-20 Start flashing confirmation

The firmware update process starts. You can check the firmware update status on the **Jobs** page. Once the firmware has been successfully updated, the server will restart automatically as requested.

## 7.4.6 Operating system deployment

XClarity Administrator can be used to manage the operating system images repository and deploy operating system images to managed servers.

For more information about how to mount and deploy an operating system image from XClarity Administrator, see 6.1, “Installing an OS without a local optical drive” on page 226.

Lenovo XClarity Administrator also integrates with managers, such as VMware vSphere and Microsoft System Center. This capability is described next.

## 7.4.7 Call home service

You can configure XClarity Administrator to automatically open a problem record and transfer diagnostic files to notify Lenovo Support when certain serviceable events are received from the managed endpoints, such as an unrecoverable memory error from a System x rack server. This automatic support notification is called *call home*.

Note that when call home is configured and enabled in Lenovo XClarity Administrator, call home is disabled on all managed chassis and servers to avoid duplicate problem records being created.

XClarity Administrator requires access to certain ports and Internet addresses for the call home function to work. Table 7-2 and Table 7-3 list the required ports and IP address.

Table 7-2 Ports that must be open for the call home feature in Lenovo XClarity Administrator

Port	Direction	Affected Devices	Purpose
80	Inbound/Outbound	Support website: Address: 129.42.0.0/18	Used for HTTP and DDP file downloads for call home
443	Inbound/Outbound	Client computers that access Lenovo XClarity Administrator	Used by HTTPS for web access and REST communications. Outbound direction is used for call home.

Table 7-3 Required connection for call home

DNS name	IP address range
esupport.ibm.com	129.42.0.0/18

Complete the following steps to enable the call home feature in Lenovo XClarity Administrator:

1. From the Lenovo XClarity Administrator menu bar, click **Administration** → **Service and Support**.
2. Click the **Call Home Configuration** tab.
3. Fill in the required fields (marked with \*) in the Configure call home section. Click **Apply**.
4. The **Enable Call Home** checkbox will be viewable.
5. Click **Enable Call Home**

For additional information on enabling call home and the features and functions, refer to the *Lenovo XClarity Administrator Planning and Implementation Guide*:

<http://lenovopress.com/sg248296>

## 7.5 Lenovo XClarity integrators

Lenovo XClarity Pro supports integration with VMware vSphere and Microsoft System Center by using XClarity Integrators. Upward integration provides automation and resiliency in virtual environments.

Each integrator integrates hardware predictive failure analysis (PFA) and microcode management and diagnostics into standard hypervisors, which provides the following capabilities:

- ▶ Manage resources from virtualization console
- ▶ Perform nondisruptive server updates
- ▶ Perform nondisruptive server starts
- ▶ Evacuate workloads on predicted hardware failure

The Lenovo XClarity Integrator plug-ins are available for the following virtualization platforms:

- ▶ Lenovo XClarity Integrator for VMware vCenter (requires a license)  
<https://support.lenovo.com/documents/LNVO-VMWARE>
- ▶ Lenovo XClarity Integrator for VMware vRealize Orchestrator (free download)  
<https://support.lenovo.com/documents/LNVO-VMRO>
- ▶ Lenovo XClarity Integrator for VMware vRealize Log Insight (free download)  
<https://solutionexchange.vmware.com/store/products/lenovo-xclarity-administrator-content-pack-for-vmware-vrealize-log-insight>
- ▶ Lenovo XClarity Integrator for Microsoft System Center (requires a license)  
<https://support.lenovo.com/documents/LNVO-MANAGE>
- ▶ Lenovo XClarity Integrator for Zenoss (requires a license)  
<https://support.lenovo.com/documents/LNVO-ZENOSS>

**Note:** The Lenovo XClarity Integrator requires a Lenovo customized ESXi version. The Lenovo customized version can be downloaded from this website:

[https://my.vmware.com/web/vmware/info/slug/datacenter\\_cloud\\_infrastructure/vmware\\_vsphere/6\\_0#custom\\_iso](https://my.vmware.com/web/vmware/info/slug/datacenter_cloud_infrastructure/vmware_vsphere/6_0#custom_iso)

If it is a generic installation, download and install the Lenovo Customization for ESXi offline bundle.

This bundle enables all management functions. Without the customized version or offline bundle installed, Lenovo XClarity Integrator for VMware vCenter provides limited management functionality.

If you purchased Lenovo XClarity Administrator licenses and want to add integration with VMware vCenter or Microsoft System Center, you can order the software license as listed in Table 7-4 or Table 7-5 on page 260, depending on your location.

**Note:** Lenovo XClarity Integrators for VMware vCenter and Microsoft System Center are included in the Lenovo XClarity Pro offerings.

*Table 7-4 Lenovo XClarity Integrator part numbers per managed server (NA, AP, Canada, and Japan)*

<b>Lenovo XClarity Integrator per managed server, for Microsoft System Center or VMware vCenter for United States, Asia Pacific, Canada, and Japan</b>	<b>Part number</b>
Lenovo XClarity Integrator for MSSC, Per Managed Srv w/1Yr S&S	00MT275
Lenovo XClarity Integrator for MSSC, Per Managed Srv w/3Yr S&S	00MT276
Lenovo XClarity Integrator for MSSC, Per Managed Srv w/5Yr S&S	00MT277
Lenovo XClarity Integrator for MSSC, w/IMMV2ADv, Per Managed Srv w/1Yr S&S	00MT278

<b>Lenovo XClarity Integrator per managed server, for Microsoft System Center or VMware vCenter for United States, Asia Pacific, Canada, and Japan</b>	<b>Part number</b>
Lenovo XClarity Integrator for MSSC, w/IMMV2ADv, Per Managed Srv w/3Yr S&S	00MT279
Lenovo XClarity Integrator for MSSC, w/IMMV2ADv, Per Managed Srv w/5Yr S&S	00MT280
Lenovo XClarity Integrator for VMware vCenter, Per Managed Srv w/1Yr S&S	00MT281
Lenovo XClarity Integrator for VMware vCenter, Per Managed Srv w/3Yr S&S	00MT282
Lenovo XClarity Integrator for VMware vCenter, Per Managed Srv w/5Yr S&S	00MT283
Lenovo XClarity Integrator f/VMw vCtr w/IMMV2Adv, Per Managed Svr w/1Yr S&S	00MT284
Lenovo XClarity Integrator f/VMw vCtr w/IMMV2Adv, Per Managed Svr w/3Yr S&S	00MT285
Lenovo XClarity Integrator f/VMw vCtr w/IMMV2Adv, Per Managed Svr w/5Yr S&S	00MT286

Table 7-5 *Lenovo XClarity Integrator part numbers per managed server (EMEA and Latin America)*

<b>Lenovo XClarity Integrator per managed server, for Microsoft System Center or VMware vCenter for Europe Middle East, Africa, and Latin America</b>	<b>Part number</b>
Lenovo XClarity Integrator for MSSC, Per Managed Srv w/1Yr S&S	00MT287
Lenovo XClarity Integrator for MSSC, Per Managed Srv w/3Yr S&S	00MT288
Lenovo XClarity Integrator for MSSC, Per Managed Srv w/5Yr S&S	00MT289
Lenovo XClarity Integrator for MSSC, w/IMMV2ADv, Per Managed Srv w/1Yr S&S	00MT290
Lenovo XClarity Integrator for MSSC, w/IMMV2ADv, Per Managed Srv w/3Yr S&S	00MT291
Lenovo XClarity Integrator for MSSC, w/IMMV2ADv, Per Managed Srv w/5Yr S&S	00MT292
Lenovo XClarity Integrator for VMware vCenter, Per Managed Srv w/1Yr S&S	00MT293
Lenovo XClarity Integrator for VMware vCenter, Per Managed Srv w/3Yr S&S	00MT294
Lenovo XClarity Integrator for VMware vCenter, Per Managed Srv w/5Yr S&S	00MT295
Lenovo XClarity Integrator f/VMw vCtr w/IMMV2Adv, Per Managed Svr w/1Yr S&S	00MT296
Lenovo XClarity Integrator f/VMw vCtr w/IMMV2Adv, Per Managed Svr w/3Yr S&S	00MT297
Lenovo XClarity Integrator f/VMw vCtr w/IMMV2Adv, Per Managed Svr w/5Yr S&S	00MT298

Table 7-6 *Zenoss part numbers for United States, Asia Pacific, Canada and Japan*

<b>Description</b>	<b>Part number</b>
Zenoss integrator per server (includes 1 year of service & support) 1 required for each physical server or host running virtualization software	
Lenovo XClarity Integrator for Zenoss	00KE784
Zenoss Resource Manager Virtual Server (includes 1 year of service & support) 1 required for each virtual machine	
Zenoss Resource Manager - 1 Virtual Server	00KE785
Zenoss Resource Manager - 10 Virtual Servers	00KE786
Zenoss Resource Manager - 100 Virtual Servers	00KE787

Description	Part number
Zenoss Resource Manager Physical Server (includes 1 year of service & support) 1 required for each network switch, storage array, and physical server without virtualization	
Zenoss Resource Manager - 1 Physical Server	00KE792
Zenoss Resource Manager - 10 Physical Servers	00KE793
Zenoss Resource Manager - 100 Physical Servers	00KE794
Zenoss Resource Manager - 1000 Physical Servers	00YC232
Zenoss Service Dynamics Virtual Server (includes 1 year of service & support) 1 required for each virtual machine	
Zenoss Service Dynamics - 1 Virtual Server	00YC233
Zenoss Service Dynamics - 10 Virtual Servers	00YC234
Zenoss Service Dynamics - 100 Virtual Servers	00YC235
Zenoss Service Dynamics Physical Server (includes 1 year of service & support) 1 required for each network switch, storage array, and physical server without virtualization	
Zenoss Service Dynamics - 1 Physical Server	00YC236
Zenoss Service Dynamics - 10 Physical Servers	00YC237
Zenoss Service Dynamics - 100 Physical Servers	00YC238
Zenoss Service Dynamics - 1000 Physical Servers	00YC239

Table 7-7 Zenoss part numbers for Europe, Middle East, Africa & Latin America

Description	Part number
Zenoss integrator per server (includes 1 year of service & support) 1 required for each physical server or host running virtualization software	
Lenovo XClarity Integrator for Zenoss	00KE795
Zenoss Resource Manager Virtual Server (includes 1 year of service & support) 1 required for each virtual machine	
Zenoss Resource Manager - 1 Virtual Server	00KE796
Zenoss Resource Manager - 10 Virtual Servers	00KE797
Zenoss Resource Manager - 100 Virtual Servers	00KE798
Zenoss Resource Manager Physical Server (includes 1 year of service & support) 1 required for each network switch, storage array, and physical server without virtualization	
Zenoss Resource Manager - 1 Physical Server	00KE799
Zenoss Resource Manager - 10 Physical Servers	00KE800
Zenoss Resource Manager - 100 Physical Servers	00KE801
Zenoss Resource Manager - 1000 Physical Servers	00YC240

Description	Part number
Zenoss Service Dynamics Virtual Server (includes 1 year of service & support) 1 required for each virtual machine	
Zenoss Service Dynamics - 1 Virtual Server	00YC241
Zenoss Service Dynamics - 10 Virtual Servers	00YC242
Zenoss Service Dynamics - 100 Virtual Servers	00YC243
Zenoss Service Dynamics Physical Server (includes 1 year of service & support) 1 required for each network switch, storage array, and physical server without virtualization	
Zenoss Service Dynamics - 1 Physical Server	00YC244
Zenoss Service Dynamics - 10 Physical Servers	00YC245
Zenoss Service Dynamics - 100 Physical Servers	00YC246
Zenoss Service Dynamics - 1000 Physical Servers	00YC247

Lenovo XClarity Integrator provides ease of management in the following ways:

- ▶ Non-disruptive server updates: Eliminate maintenance downtime through nondisruptive rolling firmware updates or server restarts, which automates the VM migration and update process of the hosts in a cluster environment without any workload interruption.
- ▶ Reliability, availability, serviceability (RAS): Lenovo XClarity Integrator can manage and set policies around all PFA in the system. It includes the ability to evacuate all VMs off one host, migrate individual VMs, or manage VMs before an outage affects them.

These features are described next.

### Concurrent firmware updates

Lenovo XClarity Integrator allows for concurrent firmware updates for ease of management and deployment. Lenovo XClarity Integrator uses the following procedure to roll out firmware updates concurrently:

1. The user starts an update, as shown in Figure 7-21.




Figure 7-21 Start a firmware update with upward integration module (UIM)



- Any virtual machines (VMs) that are running on the first server are moved from the first server to another server, as shown in Figure 7-22.

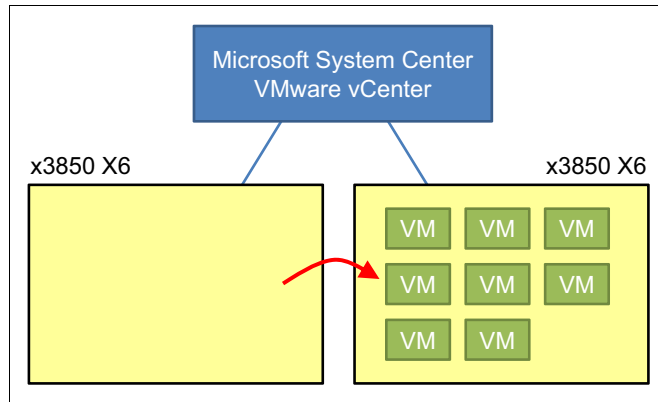


Figure 7-22 UIM evacuates first servers VMs for updating

- The server is taken offline and the updates are applied.
- After the updates are applied, the server is returned to the pool and workloads are restored to the server.
- Any running VMs are moved from the next server that is to be updated to another server, as shown in Figure 7-23.




Figure 7-23 UIM move VMs from the next server

- The server is taken offline and the updates are applied.
- After the updates are applied, the server is returned to the pool and workloads are restored to the server.
- The process repeats until all targeted machines are updated.

## Reliability, availability, serviceability

Lenovo XClarity Integrator now has upward integration with RAS events. Lenovo XClarity Integrator uses the following procedure when a RAS event is detected on a server:

1. A RAS event that is detected on the server (see Figure 7-24) is reported up the chain to Microsoft System Center, VMware vCenter or XClarity.




Figure 7-24 RAS event reports up the chain

2. The server is emptied and the workload resumes on a different server, as shown in Figure 7-25.




Figure 7-25 Workload resumes on a different server

3. The failed server is taken offline and is ready for maintenance.

For more information and a 90-day free trial of the Integrators, see the following resources:

- ▶ Lenovo XClarity Integrator for VMware vCenter, v4.0.2:

<https://support.lenovo.com/documents/LNVO-VMWARE>

- ▶ Lenovo XClarity Integrator Offerings for Microsoft System Center Management Solutions:

<https://support.lenovo.com/documents/LNVO-MANAGE>

## 7.6 Lenovo XClarity Energy Manager

Lenovo XClarity Energy Manager is a standalone piece of software that models data center physical hierarchy and monitors power and temperature at the server level and at the group level. By analyzing power and temperature data monitored, Energy Manager helps data center administrator improve business continuity and energy efficiency.

Features and functions of Energy Manager include the following:

- ▶ Monitoring energy consumption, estimating the power demand, and reallocating power to servers as needed.
- ▶ Intuitive, easy-to-use dashboard to access power and thermal information across multiple servers, racks and data centers.
- ▶ Monitoring the temperature and cooling capacity of servers.
- ▶ Sending notifications when certain events occur or when thresholds are exceeded.
- ▶ Limiting the amount of energy that an endpoint consumes using policies.
- ▶ Optimizing energy efficiency by monitoring real-time inlet temperatures, identifying low-usage servers based on out-of-band power data, measuring power ranges for different server models, and evaluating how servers accommodate new workloads based on the availability of resources.
- ▶ Reducing the power consumption to a minimum level to prolong service time during an emergency power event (such as a data center power failure).

The Lenovo XClarity Energy Manager home page is:

<http://shop.lenovo.com/us/en/systems/software/systems-management/xclarity-energy-manager/>

The following section discusses licensing, system requirements, and how to download, install, and set up Energy Manager.

### 7.6.1 Lenovo XClarity Energy Manager licensing

Energy Manager requires a license to be purchased per device for Lenovo and non-Lenovo servers. Table 7-8 shows the ordering part number information for a single-device, 5-node and 50-device licenses.

**Note:** When the x3950 X6 is partitioned, it counts as two separate nodes.

Table 7-8 Energy Manager part number ordering information

Part number	Description	Maximum supported
4L40E51621	XClarity Energy Manager Device License	1 device
4L40E51622	XClarity Energy Manager Device License Pack	5 device
4L40E51623	XClarity Energy Manager Device License Pack	50 device

You can find the number of available licenses and used licenses by clicking **About** in the top right hand corner of Energy Manager as seen in Figure 7-26.

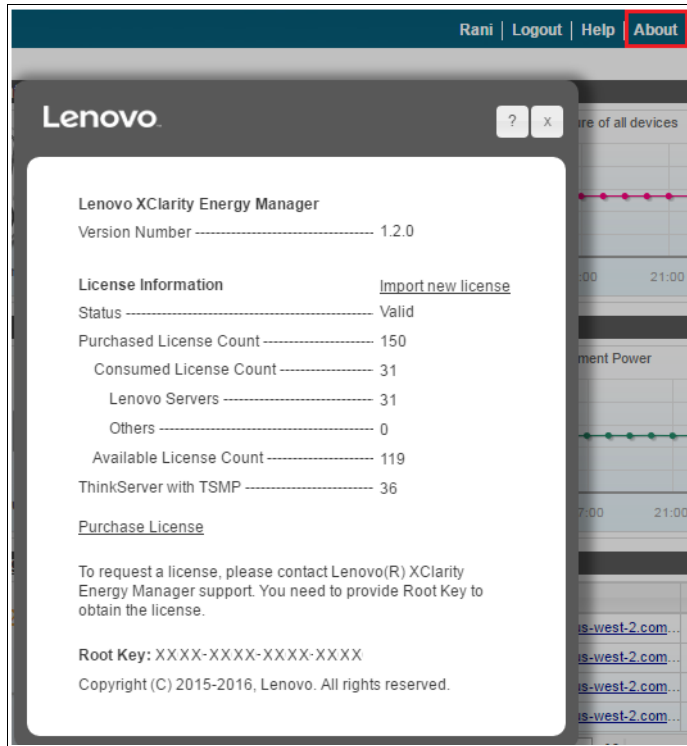


Figure 7-26 Energy Manager version and licensing

Note: A license file is bounded to a Root Key and vice versa. This means that a license file can only be imported to an Energy Manager instance with the same Root Key shown in the About dialog. The Root Key is generated based on the OS information. If an OS is reconfigured (for example, a network configuration change) or if the OS is reinstalled, the Root Key might change accordingly. This implies that you may need to request a new license file based on the new Root Key.

## 7.6.2 System requirements

The Energy Manager server where Energy Manager is installed is supported on the following operating systems:

- ▶ Windows Server 2008
- ▶ Windows Server 2012
- ▶ Windows Server 2012 R2
- ▶ Red Hat Enterprise Linux 6 (32-bit)
- ▶ Red Hat Enterprise Linux 6 (64-bit)
- ▶ Red Hat Enterprise Linux 7 (64-bit)
- ▶ SUSE Linux Enterprise Server 11 (32-bit)
- ▶ SUSE Linux Enterprise Server 11 (64-bit)
- ▶ SUSE Linux Enterprise Server 12 (64-bit)

It is recommended to install the software on a system with at least:

- ▶ A dual-core processor of 2.6 GHz or higher
- ▶ 8GB RAM
- ▶ 80GB of hard drive space

To access the Energy Manager interface, the following web browsers are supported:

- ▶ Mozilla Firefox 25
- ▶ Google Chrome 31
- ▶ Microsoft Internet Explorer 9 and above

Note: Energy Manager communicates with the managed nodes through multiple protocols including IPMI, SSH, SNMP, WS-MAN, HTTPS, and DCOM (WMI). These protocols should be allowed by the network firewall and the operating system firewall between Energy Manager and the managed nodes.

### 7.6.3 Download and install Lenovo XClarity Energy Manager

You can download the Lenovo XClarity Energy Manager from the following web page:

<https://support.lenovo.com/downloads/ds101160>

Once downloaded, for Windows, run the installer and follow the wizard for the installation. For Linux, unzip the package and then launch the executable for the installation. The software will be installed on the server where you are launching the installation package from.

The following are the instructions for installing Energy Manager on Windows:

1. Run the EXE file you downloaded from the above web page.
2. At the standard InstallShield welcome screen, click **Next**.
3. Accept the license agreement and click **Next**.
4. Enter a User Name and Organization for where the software is being installed, as shown in Figure 7-27. Click **Next**.

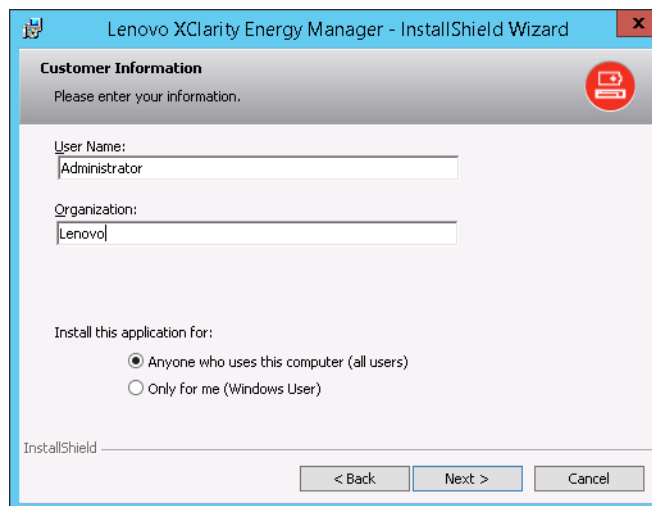


Figure 7-27 Customer information

5. Specify the directory where the software is to be installed and click **Next**.

- Specify the SNMP and Java RMI ports to be used in Figure 7-28. By default Energy Manager uses port 162 to listen for SNMP traps. It is important to check that this port is available and able to be used by Energy Manager without conflict. If the port is already in use, then set a different port for SNMP traps.

Energy Manager also uses port 1099 by default to listen to Java Remote Method Invocation (RMI) clients. If this port will cause a conflict, then set a different port. Figure 7-28 displays the SNMP and RMI settings page. Click **Next** once the ports have been set.

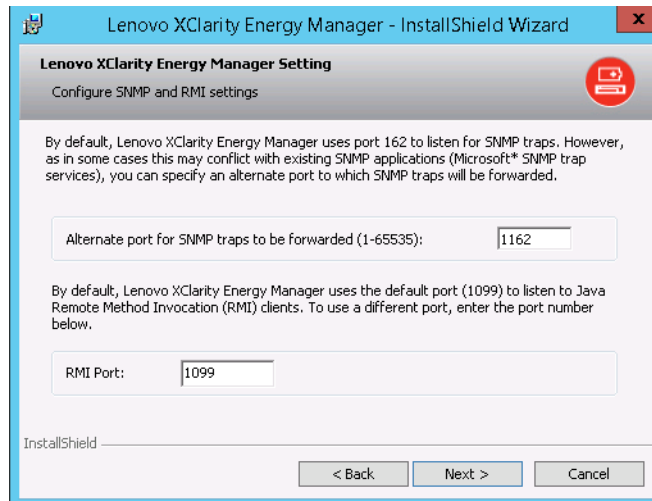


Figure 7-28 SNMP and RMI port settings

- Specify the port to be used for the web service that will be used to access the Energy Manager interface in Figure 7-29. By default, TLS (Transport Layer Security) is set to enabled. When TLS is enabled, Energy Manager communicates via port 8643 by default.

If TLS is disabled, the communication from the browser is not secure. The port used by default when TLS is disabled is 8688. If this port is already in use you can set a different port. Click **Next** to continue the installation.

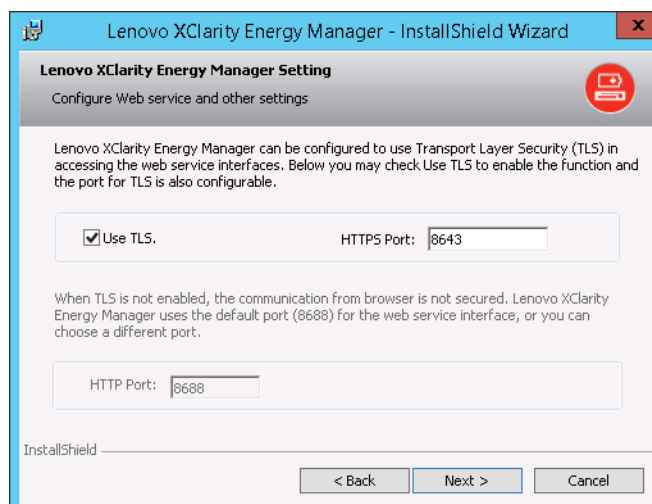


Figure 7-29 Configure web service

8. Set the sampling frequency and granularity of data as shown in Figure 7-30.

The frequency setting refers to the interval at which data will be collected. Data can be set to be collected every 30, 60, 180, 300, 360 or 600 seconds via the pulldown menu. In this example frequency is set to 60, which means collection of data will occur every minute.

The granularity refers to the interval at which the reporting graphs will be updated with new data. Granularity can be set to either 30, 60, 180, 300, 360, 600, 1800 or 3600 seconds. In this example, granularity is set to 180, so the tools graphs will be updated every 3 minutes with the new data.

Click **Next**.




Figure 7-30 Set frequency and granularity of data

In Figure 7-31, enter a username and password. These are the credentials that you will need to access the web interface of Energy Manager. From the web interface you can create and manage the data center hierarchy. Refer to 7.6.5, “Setting up Lenovo XClarity Energy Manager” on page 274 for information on the data center hierarchy.

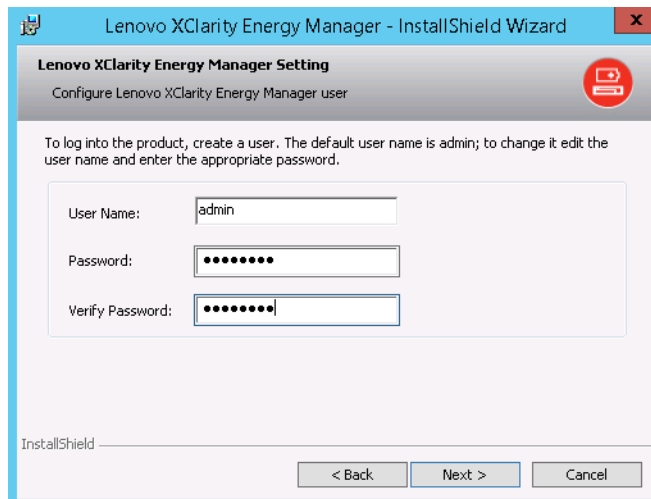


Figure 7-31 Set user name and password

9. Energy Manager uses a Java keystore (JKS) to manage security certificates. Set the JKS username and password as shown in Figure 7-32. Click **Next**.

Note: A Java Keystore (which is included with Java) is used when accessing secure websites (for instance, HTTP over TLS / HTTPS). These secure websites generate an SSL certificate. When accessing Energy Manager via your browser, you are accessing a secure site. By creating this keystore file and keystore password, it will be used when the browser interface connects to Energy Manager. The purpose of this is to encrypt and decrypt user page requests as well as the pages that are returned by the Energy Manager server. Your keystore password can be unique to the other required passwords for Energy Manager.

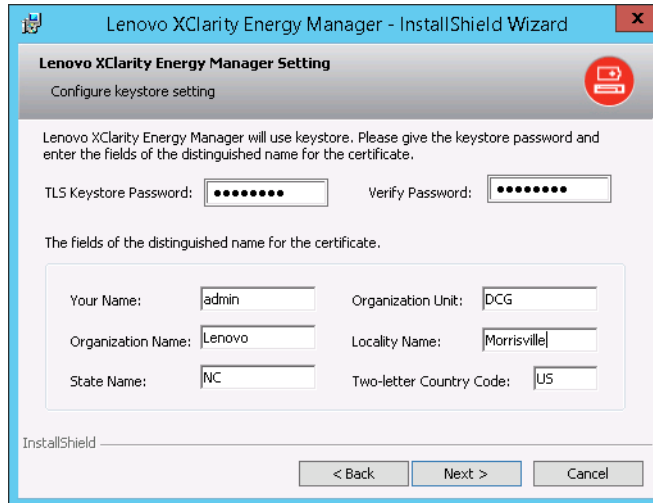


Figure 7-32 Keystore access

10. Energy Manager has an embedded database server. As seen in Figure 7-33, you can set the database attributes including username, an open/unused port, password and installation path. Once set, click **Next**.

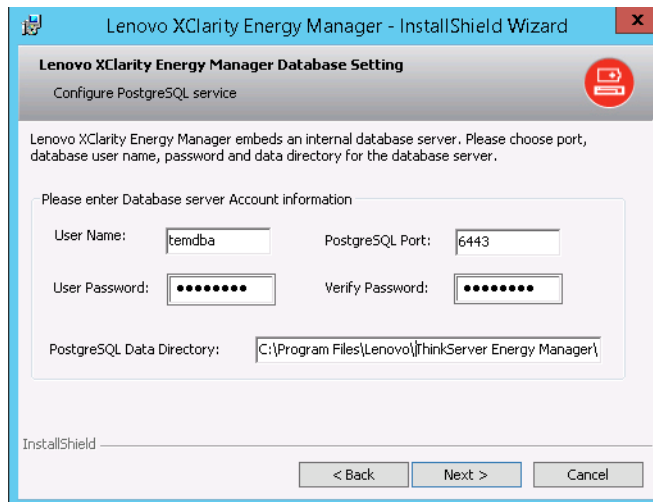


Figure 7-33 Embedded database installation



11. If changes to any settings are needed, click the **Back** button to make the changes now. Otherwise, to begin the installation, click **Install** as shown in Figure 7-34.

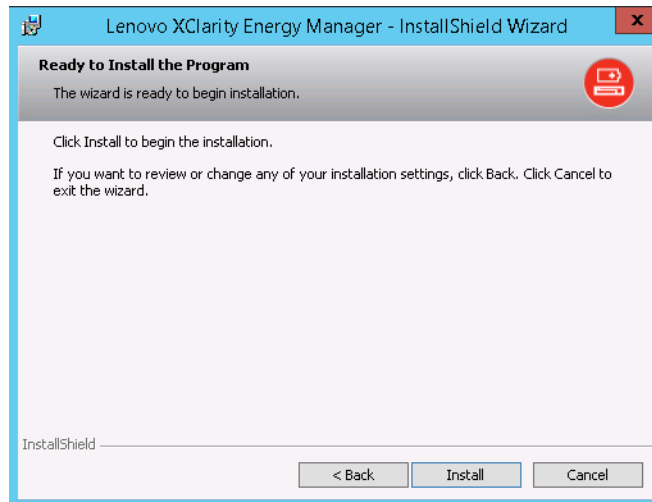


Figure 7-34 Begin installation

12. When the installation is complete, the installation wizard will display a successful installation message. Click the **Finish** button to exit the wizard. You can now launch the web interface of Energy Manager, as described in the next section.

## 7.6.4 Launching Lenovo XClarity Energy Manager

When Energy Manager has been successfully installed there are two ways to launch it.

- ▶ By typing the URL directly in the address bar of a browser.  
Enter the following default URL in a supported web browser to launch the Energy Manager if you enabled Transport Layer Security (TLS) during the installation procedure as shown in Figure 7-29 on page 268. Enter the port you specified during installation if you didn't accept the default.  
`https://localhost:8643/energymanager/`  
If TLS is not enabled enter the following URL:  
`http://localhost:8688/energymanager/`
- ▶ By clicking the Lenovo XClarity Energy Manager shortcut. A shortcut is created if Energy Manager is installed on a Windows server.

Once launched the login page will appear as seen in Figure 7-35. Enter the user ID and password that was set during installation (Figure 7-31 on page 269).

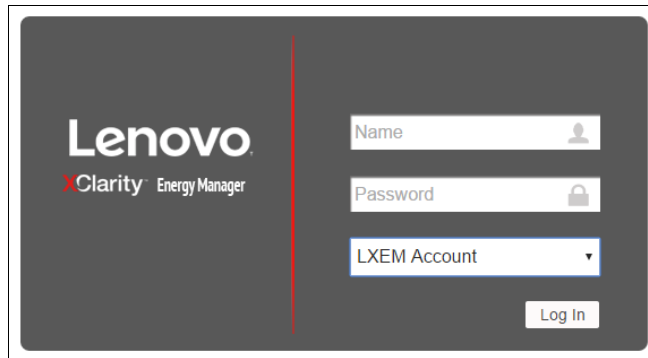


Figure 7-35 Energy Manager login screen

Once logged in to Energy Manager, the dashboard is displayed, Figure 7-36. The Dashboard provides the overall health status of the data center. It shows the current power and cooling status, the historical power and temperature trend, the hot rooms and hotspots, the power and space capacity information, and the critical events. These are displayed in individual information boxes which are called *gadgets*. In Figure 7-36, the gadgets do not currently display any data because we have not yet discovered any devices.

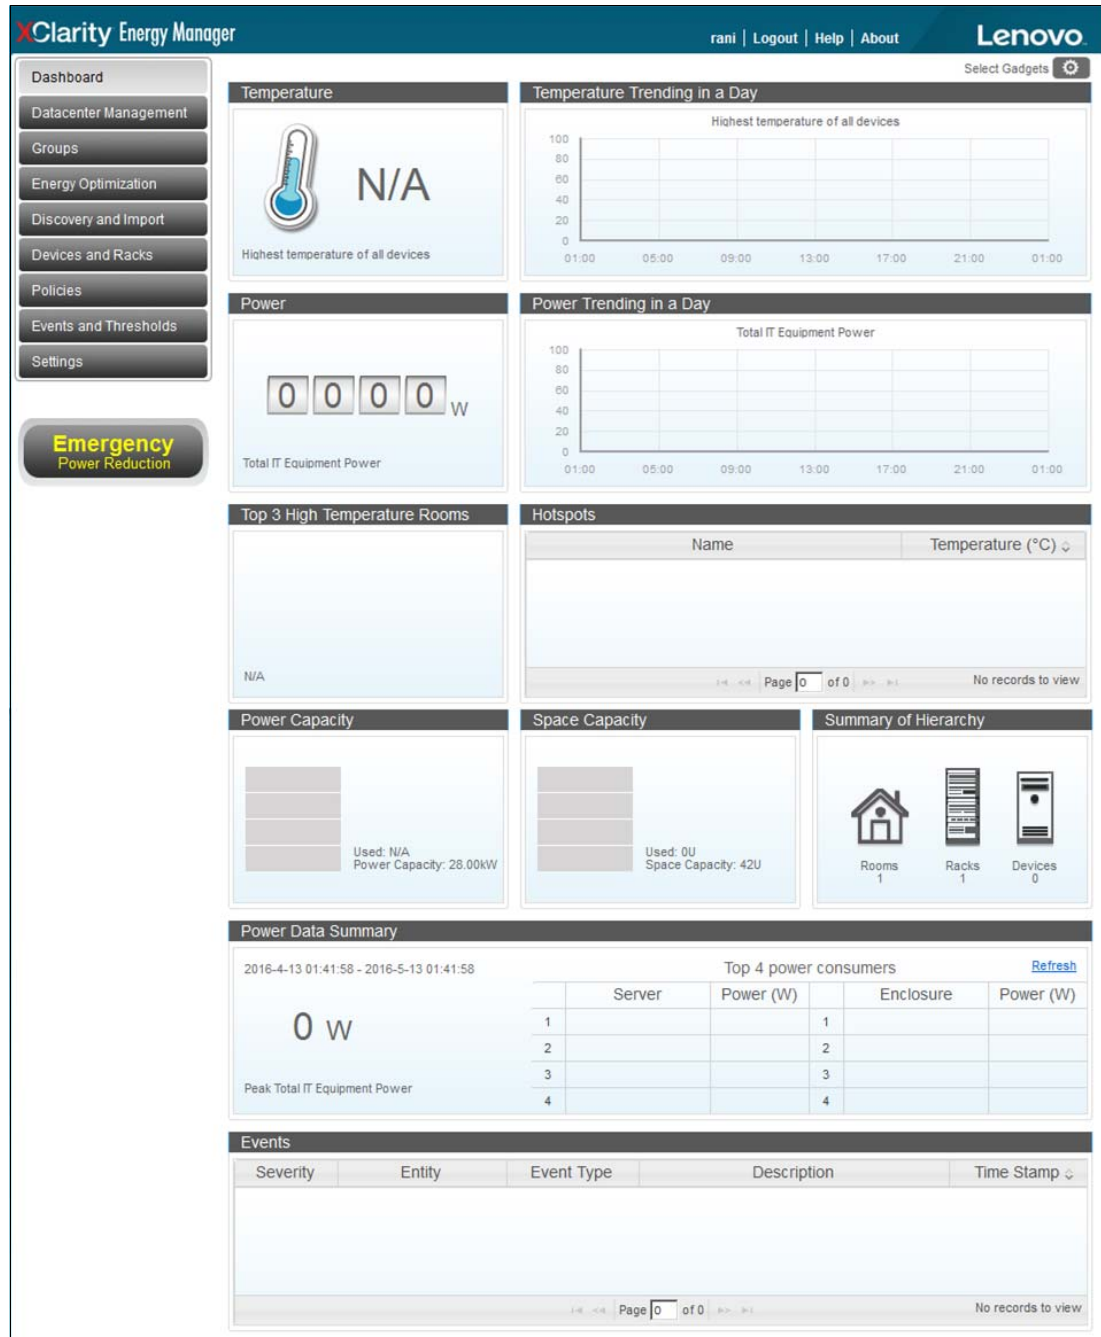


Figure 7-36 Lenovo XClarity Energy Manager Dashboard

You can customize the Dashboard by adding and deleting gadgets that are of interest to you. To add and delete gadgets, click the **Select Gadgets** button on the top right hand corner of the Dashboard. Check or uncheck the checkbox next to each Gadget description to delete or add that Gadget to the Dashboard, as shown in Figure 7-37 on page 274.

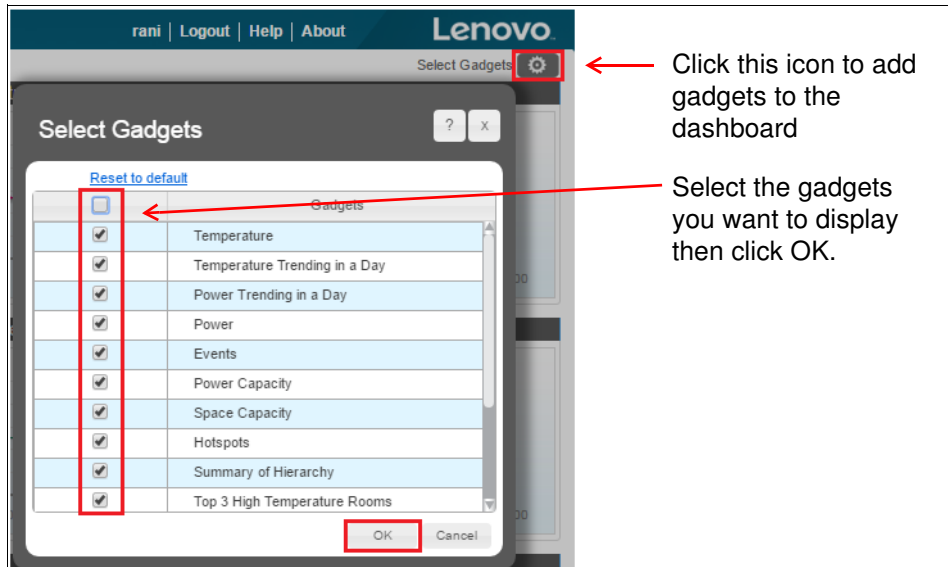


Figure 7-37 Customizing the Dashboard with gadgets

## 7.6.5 Setting up Lenovo XClarity Energy Manager

To start monitoring and managing power and temperature of devices with Lenovo XClarity Energy Manager, first create a hierarchy so when a device is discovered, its location can be specified in the hierarchy.

Energy Manager provides several ways to set up the data center hierarchy. The hierarchy is as follows:

- ▶ Data Centers: Where you can add rooms
- ▶ Rooms: Where you can add rows
- ▶ Rows: Where you can add racks
- ▶ Racks: Where you can add devices (chassis, server, PDU, UPS, etc)
- ▶ Devices: That are monitored by Energy Manager for power and temperature trends

To set up a hierarchy, click the Datacenter Management button from the left hand menu and create entries for a Data Center, Room, Row and Rack as seen in Figure 7-38.

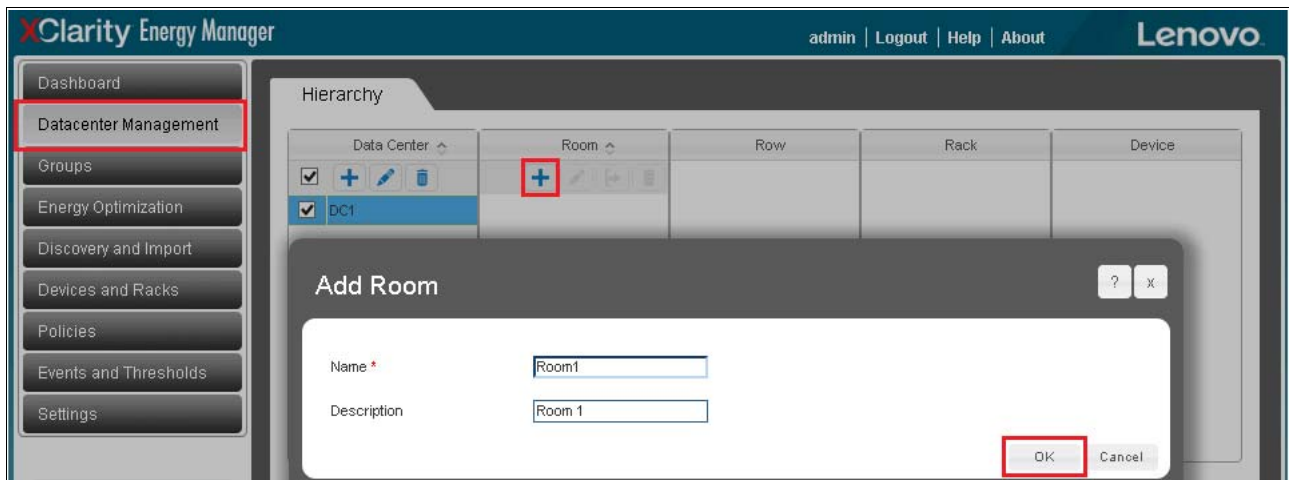


Figure 7-38 Creating a hierarchy

When adding a rack to the hierarchy, enter the total power available (in Watts) within that rack as seen in Figure 7-39. To determine the total power available in your rack, refer to the PDU Technical Reference Guides for information on the PDUs power capacities in your rack. The PDU Technical Reference Guides are located at the following web page:

<https://support.lenovo.com/documents/LNVO-POWINF>

Check the box for **PDU Power as Rack Power** if you want to use the power reading of the PDU(s) in your rack as the IT equipment power of the rack. Click **OK** to add the rack.

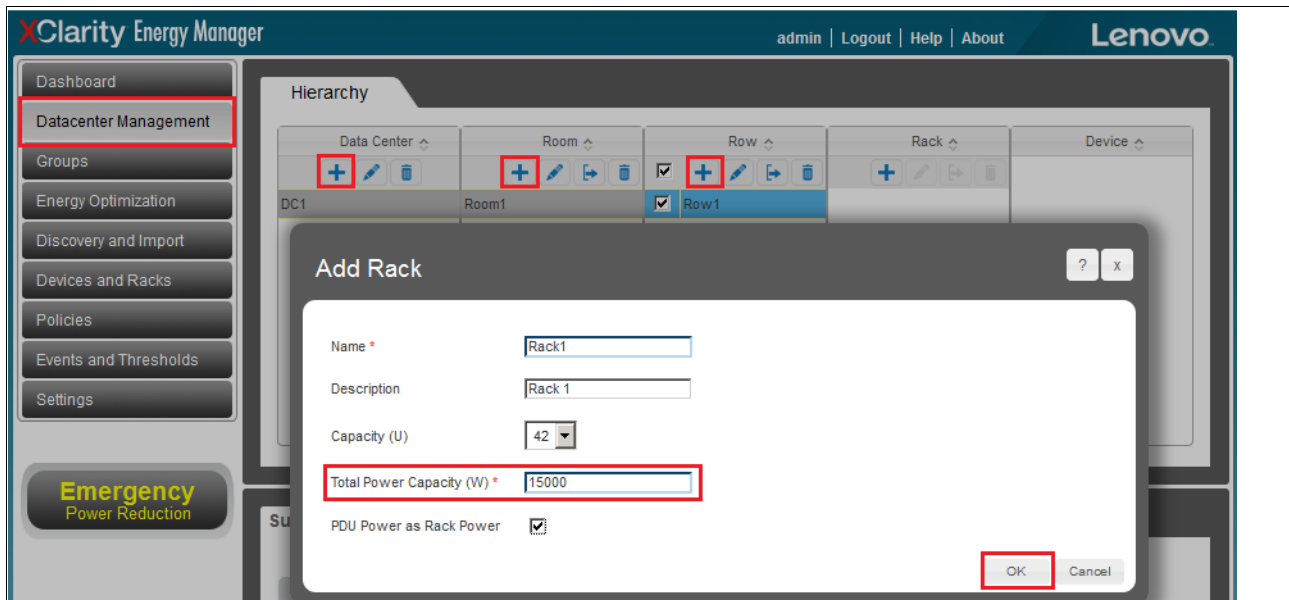


Figure 7-39 Adding a rack to the hierarchy

## Discovering devices with Lenovo XClarity Energy Manager

There are two ways to discover devices and add them to the hierarchy.

1. Discovery: To discover devices in the network, start from **Discovery and Import** on the left side navigation menu as seen in Figure 7-40 on page 276.

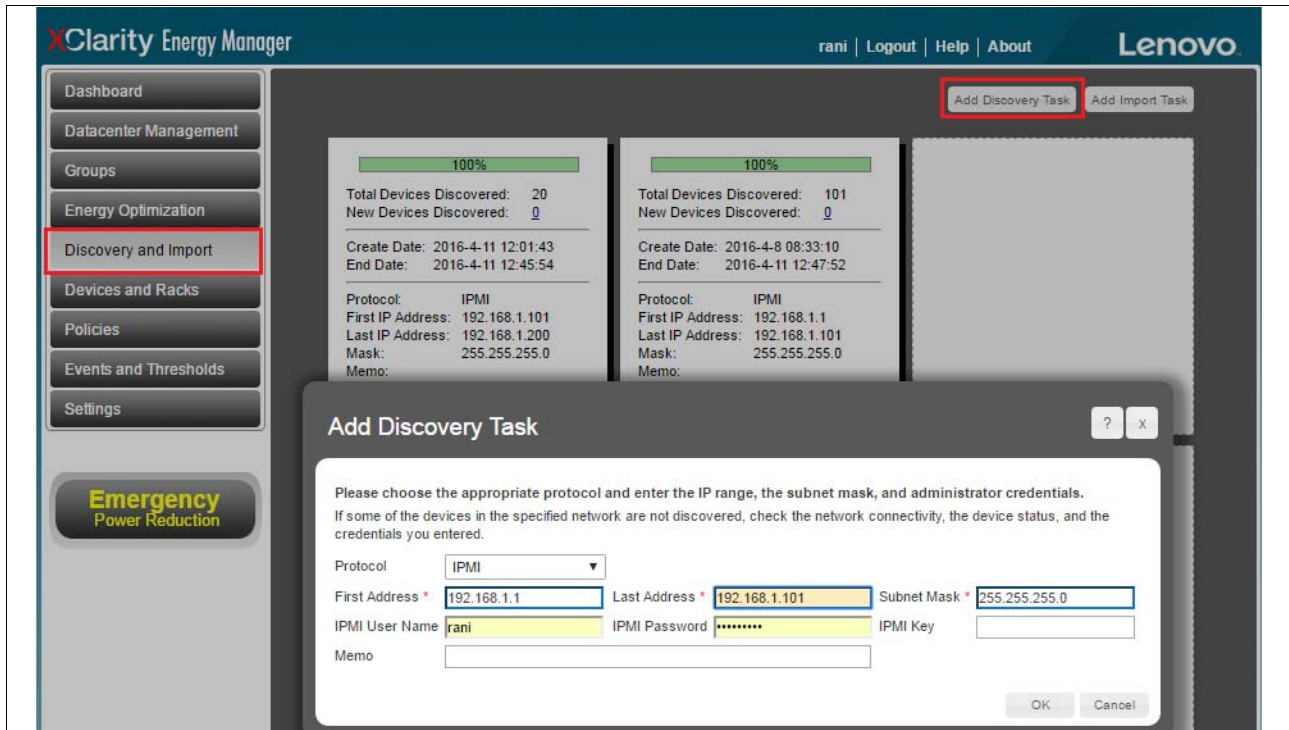


Figure 7-40 Discovery and import

- Import devices (or an entire hierarchy) from an xls: To import devices or the hierarchy from an Excel file, start from the **Import** button in **Devices and Racks** as seen in Figure 7-41. For the Excel file requirements, refer to the *Lenovo XClarity Energy Manager User Guide* located at this URL:

<http://support.lenovo.com/us/en/downloads/ds101160>

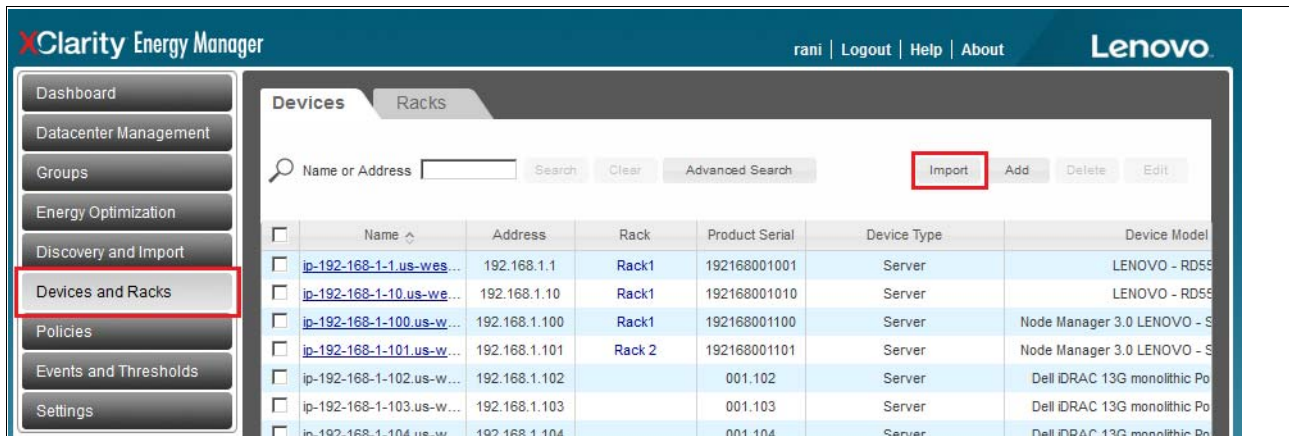


Figure 7-41 Discovery and import from Excel file

Once you have devices discovered in Energy Manager, you can add them to the hierarchy previously created.

## Adding devices to the hierarchy

To add your discovered devices to a rack within the hierarchy, select the Datacenter, room, row, and rack that you previously created in the **Datacenter Management** page. Click the **+** icon to add the discovered device(s) to that rack, as seen in Figure 7-42.

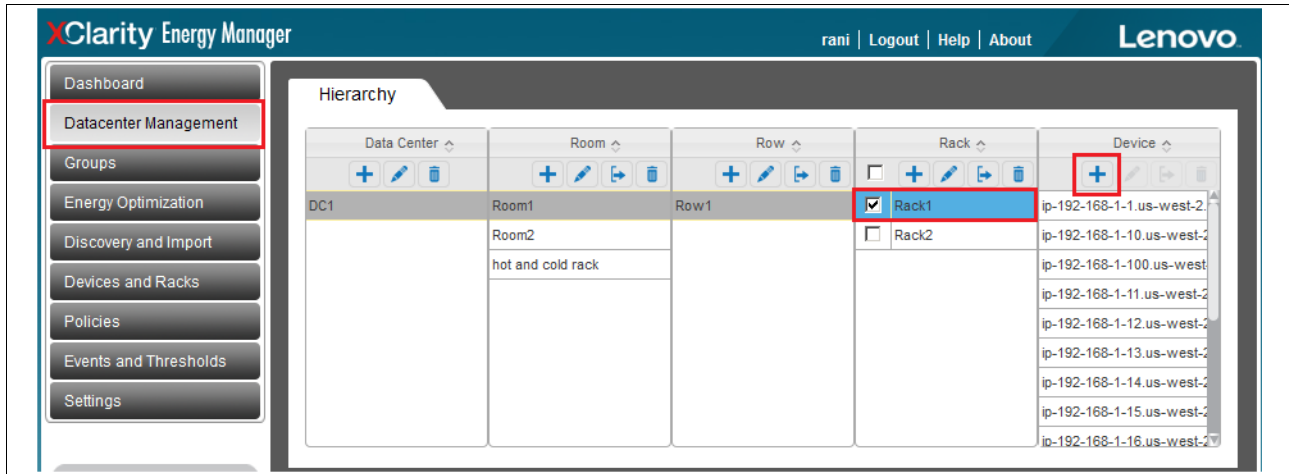


Figure 7-42 Adding discovered devices to a rack

## Managing and viewing devices in the hierarchy

Once the hierarchy is set up and discovery is made, the dashboard will now display information based on the gadgets you have selected to view. An example of the dashboard is shown in Figure 7-43.



Figure 7-43 Energy Manager Dashboard

The Datacenter Management page gives you complete control over the hierarchy allowing you to add, edit, move and delete data centers, rooms, rows, racks and devices as needed.



The hierarchy is interactive and will update information displayed in the GUI based on the selection made in the hierarchy. For instance, when a room is selected in the hierarchy, information for the selected room is displayed in the Summary tab, as seen in Figure 7-44. You can view stats as well as how many racks are in the room, and how many devices exist.

The screenshot shows a GUI with a hierarchy view and a summary view. The hierarchy view has columns for Data Center, Room, Row, Rack, and Device. Under Data Center 'DC1', 'Room1' is selected. The summary view for 'Room1' shows:

- Temperature:** 28°C (Highest Inlet Temperature)
- Power:** 9232W used, 40000W capacity
- Space:** 33U used, 84U capacity
- Room Statistics:** Name: Room1, Description: , Racks: 2, Devices: 33

**Events Table:**

Severity	Entity	Event Type	Description	Time Stamp
!	ip-192-168-1-100.us...	COMMUNICATION_WITH_NODE...	IPMI session has not been creat...	2016-4-11 12:49:00
!	ip-192-168-1-100.us...	COMMUNICATION_WITH_NODE...	Receive timeout, state =TIMEOUT	2016-4-11 12:48:54
!	ip-192-168-1-100.us...	COMMUNICATION_WITH_NODE...	Receive timeout, state =TIMEOUT	2016-4-11 12:48:48

Figure 7-44 Hierarchy room selection displays room statistics

Likewise, when a rack is selected, information for that particular rack is displayed, as seen in Figure 7-45. The number of devices and their location in that rack is displayed, as well as information on the highest and lowest inlet temperatures, the overall and remaining power capacity for that rack, and U space left in the rack. Temperature trending, policies and threshold information is also accessible from here. These are discussed further in “Setting policies” on page 282 and “Setting thresholds” on page 283.

The screenshot shows a management interface with a 'Hierarchy' section at the top. It has five columns: Data Center, Room, Row, Rack, and Device. Under 'Rack', 'Rack1' is selected and highlighted with a red box. Below this, a 'Summary' panel is visible with three main metrics: Temperature (28°C), Power (4561W), and Space (17U). To the right of these metrics, there is a 'Name Rack1' section with 'Description' and 'Devices 17'. Below the summary, there is an 'Events' table with columns for Severity, Entity, Event Type, Description, and Time Stamp. The table contains multiple rows of communication timeout events. On the far right, a list of device IP addresses is visible, including ip-192-168-1-1.us-west-2 through ip-192-168-1-16.us-west-2.

Figure 7-45 Hierarchy rack selection displays rack information

The rack itself in the Summary tab is interactive and will update information based on the device selected in the rack. Information such as IP address, serial and model numbers and power draw is displayed. If enabled, you can also turn the device on and off, as seen in Figure 7-46.

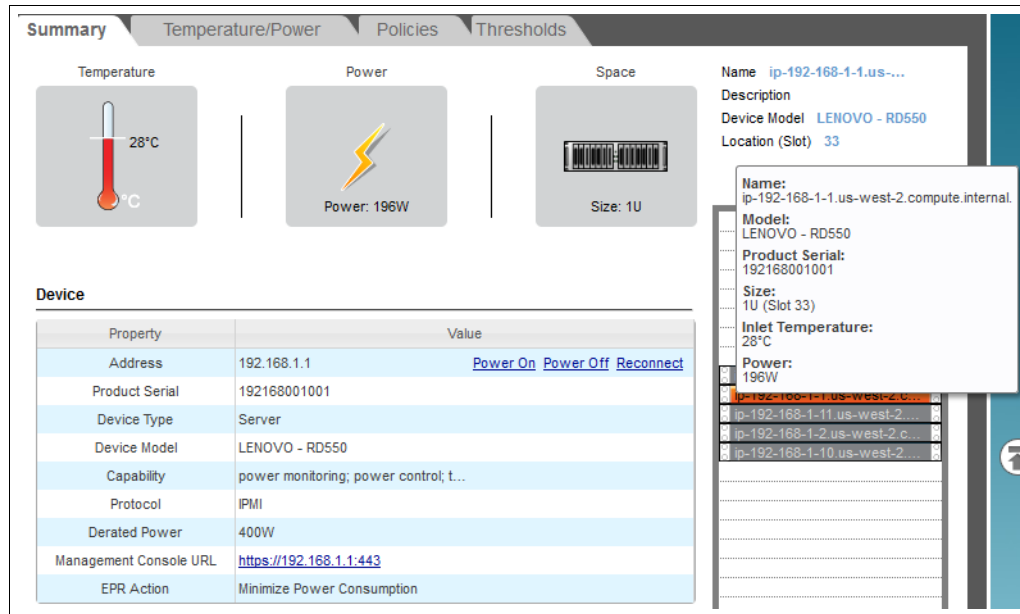


Figure 7-46 Device selected in rack

From the tabs, temperature/power trending, policies and thresholds can also be viewed for each individual server.

## Setting and monitoring groups in the hierarchy

Energy Manager allows creation of groups of devices to manage. The group can be managed, monitored, and configured in the same way as in the hierarchy.

To create a group: On the **Groups** page, click the **+** button under the **Group List** to add a group. Specify the name and an optional description in the popup dialog, and then click OK. You will see your group added to the Group List.

Select the group just created and click the **+** button under **Device Name** to add a device to it. Select **data center** → **room** → **row** → **rack** → **device** as seen in Figure 7-47.

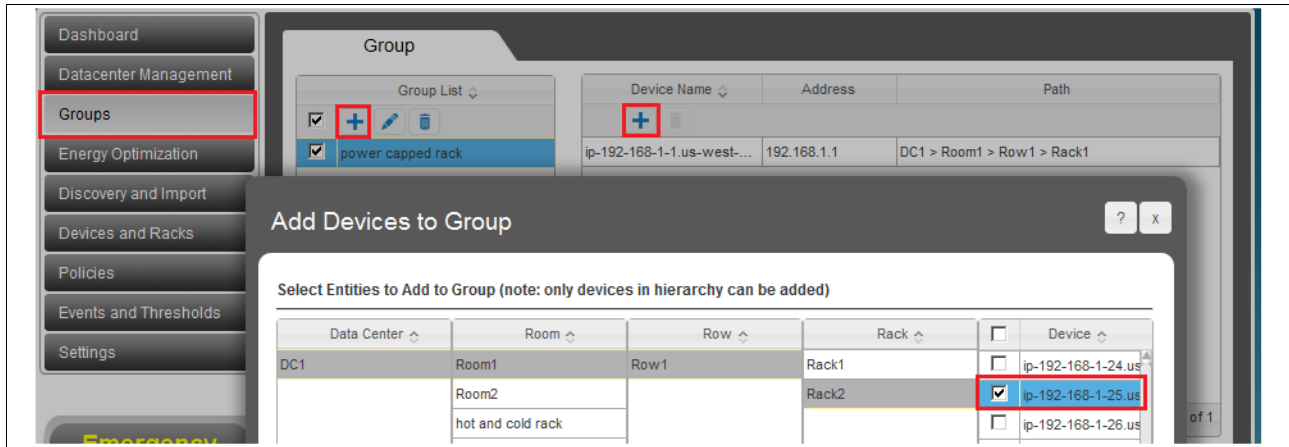


Figure 7-47 Creating a group

You can also select a rack for all devices in the rack to be added to a group.

## SNMP traps

SNMP traps can be used to assign a recipient to receive events triggered. This makes it easier to manage the events in 3rd-party event management systems. Energy Manager events are defined in the Management Information Base (MIB) file. And the MIB file is installed at “<installation path>\conf\DCMConsole-MIB-V1.mib”.

To add a trap receiver, go to the Settings page and click **Add Receiver**, fill in the Destination IP Address or Hostname, Port, and Community String fields, then click **OK**.

## Email alerting

Energy Manager allows you to subscribe to alerts. This is done in the **Settings** page under the “**Email Subscription**” tab. To subscribe to alerts and events: Go to the Settings page → click Add Subscriber → fill in the email server configuration → and check the “Subscribe threshold-based events only” (if you are only after threshold based events) → Click ok.

## Setting policies

You can use policies to limit the amount of power that an entity consumes. There are two ways in which a policy can be applied; either to a group entity or to an individual device.

To set a policy: Click the **Datacenter Management** page. Select a device from the hierarchy by selecting **data center** → **room** → **row** → **rack** → **device**. Click the **Policies** tab. Click the **Add** button.

In the popup dialog, specify the policy name and select the policy type from the drop-down list. There are two types of policies available:

- ▶ Custom Power Limit: If this is selected, Energy Manager will generate an alert when the actual power consumption is higher than the threshold you configured.
- ▶ Minimum Power: If this is selected, Energy Manager throttles the device power to the minimum (so you do not need to specify a threshold).

Set the schedule for the policy as either a permanent policy, for a specific time or as recurrent. Figure 7-48 displays the policy page.

The screenshot shows the 'Add Policy' configuration page. On the left is a navigation menu with items like Dashboard, Datacenter Management, Groups, Energy Optimization, Discovery and Imaging, Devices and Racks, and Policies. The main content area is titled 'Add Policy' and contains the following fields:

- Entity:** ip-192-168-1-1.us-west-2....
- Name:** Power Capped Servers
- Reserve Budget (W):** 0
- Policy Enabled:**
- Policy Type:** Custom Power Limit
- Power Cap (W):** 230

Figure 7-48 Setting a policy

## Setting thresholds

When a threshold is set it monitors a device or group for that condition. When a condition is met or exceeded an event is triggered and listed in the Summary tab. There are two types of thresholds that can be set:

- ▶ **Power thresholds:** Collected data is compared with the device or group power consumption (in units of Watts)
- ▶ **Temperature thresholds:** Collected data is compared with the device or groups temperature (in units of Celsius degrees)

To set a threshold: Click the **Datacenter Management** page. Select a device from the hierarchy by selecting **data center** → **room** → **row** → **rack** → **device**. Click the **Thresholds** tab. Click the **Edit** option to set the threshold. Figure 7-49 displays the Thresholds page.

The screenshot shows the 'Thresholds' configuration page for the device 'ip-192-168-1-1.us-west-2.compute.internal'. The page has tabs for Summary, Temperature/Power, Policies, and Thresholds. It is divided into two sections:

**Power Thresholds**

Event Type	Condition	Threshold (W)
IT Equipment Power	Greater	340

Buttons: OK, Cancel

**Temperature Thresholds**

Event Type	Condition	Threshold (°C)
Highest Inlet Temperature	Greater	<a href="#">Edit</a> <a href="#">Clear</a>
Lowest Inlet Temperature	Less	<a href="#">Edit</a> <a href="#">Clear</a>
Average Inlet Temperature	Greater	<a href="#">Edit</a> <a href="#">Clear</a>
Average Inlet Temperature	Less	<a href="#">Edit</a> <a href="#">Clear</a>

Figure 7-49 Setting a threshold

## 7.6.6 Energy optimization with Energy Manager

Energy Manager analyzes different data to help optimize the energy efficiency within the datacenter. There are four types of optimization strategies available, Cooling Analysis, Low-Utilization Servers, Server Power Characteristics and Workload Placement. These are discussed in further detail below. Figure 7-50 shows the Energy Optimization page.



Figure 7-50 Energy Optimization page

## Cooling Analysis

The cooling analysis page provides real-time monitoring data of the inlet temperatures of each device. The results are published in a bar graph where X-axis represents temperature values and Y-axis represents the percentage of servers at the corresponding temperature, as seen in Figure 7-51.

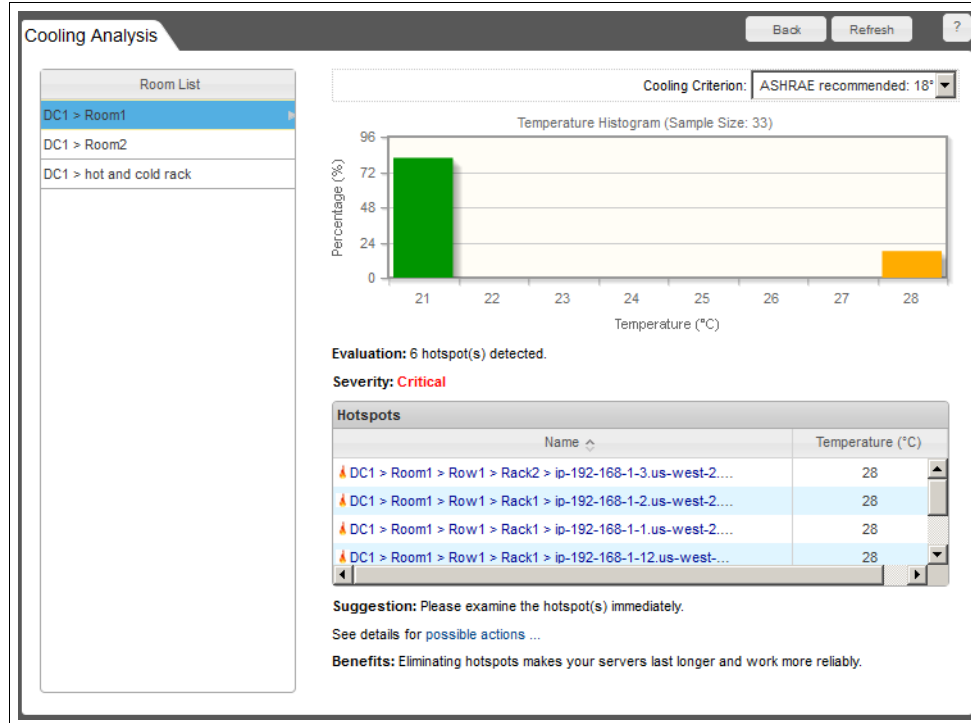


Figure 7-51 Cooling analysis page

Energy Manager will identify the servers that are causing hotspots in the data center and suggest an action to eliminate the hotspot, as seen in Figure 7-52. In this example, there are 6 servers with inlet temperatures that are higher than 27 degrees Celcius.

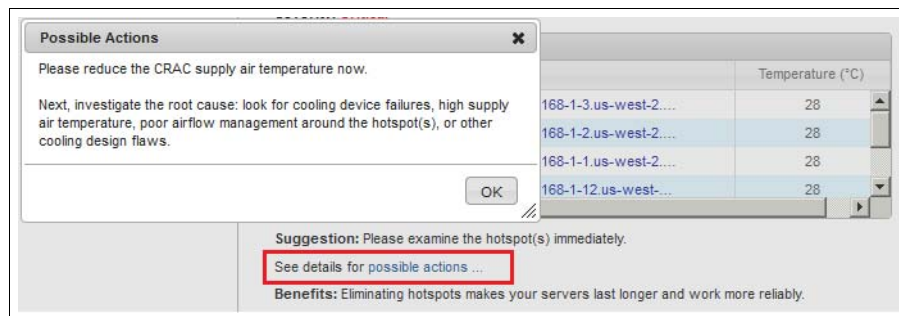


Figure 7-52 Cooling analysis page - possible action

## Low-utilization servers

Energy Manager evaluates and tries to identify low-utilization servers based on out-of-band power data. Using a set of heuristics, server utilization is estimated according to the power history with all the raw data that is collected. Given the statistics of utilization data estimated, low-utilization servers are identified.

By clicking **Analyze**, low-utilization servers will be identified and listed as potential targets for consolidation to optimize energy efficiency. The tdate and ime that the analysis was done is shown in the bottom left of the screen as seen in Figure 7-53.

Low-Utilization Servers			
Name	Average Utilization Estimated	99%ile Utilization Estimated*	Estimated Energy Saving (kWh) per Week after Consolidation
<a href="#">ip-192-168-1-19.us-west-2.compute.internal</a>	0.01%	0.53%	32
<a href="#">ip-192-168-1-101.us-west-2.compute.inter...</a>	0.01%	0.53%	32
<a href="#">ip-192-168-1-86.us-west-2.compute.internal</a>	0.01%	0.53%	32
<a href="#">ip-192-168-1-12.us-west-2.compute.internal</a>	0.01%	0.53%	32
<a href="#">ip-192-168-1-111.us-west-2.compute.inter...</a>	0.01%	0.55%	31
<a href="#">ip-192-168-1-112.us-west-2.compute.inter...</a>	0.01%	0.55%	31
<a href="#">ip-192-168-1-108.us-west-2.compute.inter...</a>	0.01%	0.55%	31
<a href="#">ip-192-168-1-90.us-west-2.compute.internal</a>	0.01%	0.57%	30
<a href="#">ip-192-168-1-21.us-west-2.compute.internal</a>	0.01%	0.57%	29
<a href="#">ip-192-168-1-119.us-west-2.compute.inter...</a>	0.01%	0.58%	29
<a href="#">ip-192-168-1-114.us-west-2.compute.inter...</a>	0.01%	0.58%	29
<a href="#">ip-192-168-1-15.us-west-2.compute.internal</a>	1.52%	1.98%	34
<a href="#">ip-192-168-1-39.us-west-2.compute.internal</a>	1.59%	2.08%	32
<a href="#">ip-192-168-1-1.us-west-2.compute.internal</a>	1.60%	2.08%	32

Analyzed at 2016-4-30 01:23:56 Page 1 of 4 View 1 - 14 of 49

Figure 7-53 Low utilization servers detected

### Server Power Characteristics

The Server Power Characteristics page represents the power usage of a server, as seen in Figure 7-54.

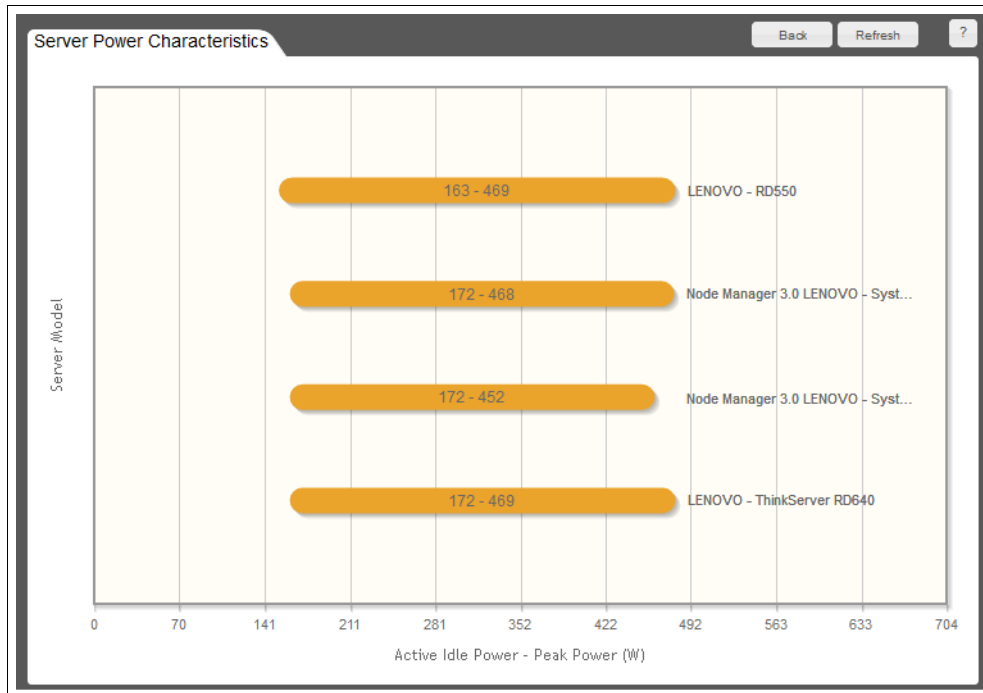


Figure 7-54 Server power characteristics

The X-axis shows the power values and the Y-axis shows the server model. The descriptions next to the bars represent the power ranges measured for those server models.



In this example, the top bar reads 163-469. This means that, for all the servers that are this certain model, the lowest power observed was 163 Watts and the highest power observed was 469 Watts.

Each bar can be clicked on to show detailed power characteristics for the server type as seen in Figure 7-55 which shows a break down of how peak power and idle power are distributed.



Figure 7-55 Power distribution characteristics

## Workload placement

The workload placement page evaluates how likely a server would be able to accommodate a new workload being added to it based on the current resource utilization and availability of the server and the resources needed for the new workload.

To analyze a server for additional workload, the new workload in terms of CPU, memory, and I/O needs to be known. An example search is shown in Figure 7-56.

Name	Evaluation Score
ip-192-168-1-92.us-west-2.compute.internal.	100.00
ip-192-168-1-93.us-west-2.compute.internal.	100.00
ip-192-168-1-94.us-west-2.compute.internal.	100.00
ip-192-168-1-95.us-west-2.compute.internal.	100.00
ip-192-168-1-96.us-west-2.compute.internal.	100.00
ip-192-168-1-97.us-west-2.compute.internal.	100.00
ip-192-168-1-98.us-west-2.compute.internal.	100.00
ip-192-168-1-99.us-west-2.compute.internal.	100.00
ip-192-168-1-100.us-west-2.compute.internal.	100.00
ip-192-168-1-101.us-west-2.compute.internal.	100.00

Figure 7-56 Workload analysis

This option is only available on servers with CUPS monitoring capabilities.

## 7.7 Upward Integration Module for Zenoss

Lenovo System x Upward Integration Module (UIM) for Zenoss extends the Zenoss Resource Manager heterogeneous monitoring tool with System x and Flex System infrastructure inventory information and events.

The solution helps IT Operations transform from an endpoint management focus to a service-oriented focus by automatically discovering and modeling heterogeneous hardware and software resources in the environment and immediately creating an asset-service relationship context. This end-to-end visualization helps IT Operations isolate and address problems faster and reduce mean time to resolution (MTTR).

The following key features are included:

- ▶ Agentless software plug-ins are installed on the Zenoss Resource Manager server
- ▶ Automatically discovers System x and Flex System servers in the environment, and makes them available for visualization and monitoring from the Zenoss Resource Monitor graphical user interface (GUI)
- ▶ Visibility into System x and Flex System hardware and firmware inventory data from a unified monitoring console
- ▶ Simple Network Management Protocol v3 (SNMPv3) based inventory collection, availability and health monitoring, and trap handling
- ▶ Faster incident resolution through holistic and precise root cause analysis, with automated error and exceptions handling, such as restarting system components

For more information about the UIM for Zenoss, see this web page:

<http://www.zenoss.com/solution/lenovo>

## 7.8 Advanced Settings Utility

By using the Advanced Settings Utility (ASU), you can modify your server firmware settings from a command line. It supports multiple operating systems, such as Linux, Solaris, and Windows, including Windows Preinstallation Environment (PE). UEFI and IMM2 firmware settings can be modified on the X6 platform.

You can perform the following tasks by using the ASU:

- ▶ Modify the UEFI CMOS settings without the need to restart the system and access the F1 menu.
- ▶ Modify selected settings in an IMM-based server for the IMM firmware and Lenovo x86 server firmware.
- ▶ Modify CMM VPD settings on IMM-based servers.
- ▶ Modify the iSCSI start settings. To modify the iSCSI settings through ASU, you must first manually configure the iSCSI settings through the server setup utility.
- ▶ Remotely modify all of the settings through an Ethernet connection.
- ▶ Mount an ISO file or USB/CD/DVD to a remote IMM-based system (the IMM must be available to the network and be accessible).
- ▶ Feature on Demand (FoD) activation key management.
- ▶ Uses Secure Boot Configuration to manage secureboot-related keys and signatures in UEFI.

ASU supports scripting environments through batch-processing mode. For more information about downloading ASU and *Advanced Settings Utility User's Guide*, see the following Advanced Settings Utility website:

<https://support.lenovo.com/documents/LNV0-ASU>

### 7.8.1 Using ASU to configure settings in IMM2-based servers

ASU 9.x supports configuring settings on servers with IMM2s, such as the X6 servers. The ASU uses the same set of commands and syntax that is used by previous versions of the ASU tool.

In IMM2-based servers, you configure all firmware settings through the IMM2. The ASU can connect to the IMM2 locally (in-band) through the keyboard console style (KCS) interface or through the LAN over USB interface. The ASU can also connect remotely over the LAN (out-of-band).

When the ASU runs any command on an IMM2-based server, it attempts to connect and automatically configure the LAN over a USB interface if it detects that this interface is not configured. The ASU also provides a level of automatic and default settings. You have the option of specifying that the automatic configuration process is skipped if you manually configured the IMM2 LAN over a USB interface. We advise that the ASU configure the LAN over a USB interface.

For more information, see this ASU web page:

<https://support.lenovo.com/documents/LNV0-ASU>

**Tip:** After you use the ASU to change settings, you must reset the IMM2 before you flash new firmware; otherwise, the changes to the settings might be lost. To reset the IMM2, use the `asu rebootimm` command.

Complete the following steps to download, install, and connect to the IMM2 by using a Windows operating system:

1. Create a directory that is named ASU.
2. Download the ASU Tool for your operating system (32-bit or 64-bit) from the following web page and save it in the ASU directory:

<https://support.lenovo.com/documents/LNVO-ASU>

3. Unpack the utility:

- For Windows, double-click the `filename.exe`, where `filename` is the name for the Advanced Settings Utility file for Windows that you downloaded.
- For Linux, from a shell command prompt, enter one of the following commands and press Enter:

If the `.tgz` file for ASU was downloaded, use the `tar -zxvf filename.tgz` command.

Where `filename` is the name of the Advanced Settings Utility file for Linux that you downloaded. The files are extracted to the same directory.

If the `.rpm` file for ASU was downloaded, use the `rpm -Uvh filename.rpm` command.

Where `filename` is the name of the Advanced Settings Utility file for Linux that you downloaded. The files are extracted to the `/opt/IBM/toolscenter/asu` directory.

4. Run a command, such as the `asu show` command, in-band or out-of-band, by using the commands that are listed in Table . This command confirms that the connection and utility work.

For in-band, run the `asu show--kcs` command.

For out-of-band, complete the following steps:

- a. Ping the IMM2 to ensure that you have a network connection to the IMM2. The default IP address is `192.168.70.125`.
- b. Run the following command:

```
asu show --host target_IMM_external_IP_address --user target_IMM_User_ID
--password target_IMM_password
```

**Note:** If the ASU is connecting remotely to the IMM2 over the LAN, there is no requirement for the remote operating system of the targeted IMM2 to be online. The ASU can connect to the IMM2 remotely when the server is connected to power or is using standby power.

## 7.8.2 Command examples

In this section, we provide a brief overview for the most commonly used ASU commands. For more information about for all of the ASU commands, see *Advanced Settings Utility User's Guide*, which is available at this web page:

<https://support.lenovo.com/documents/LNVO-ASU>

## Show all settings

At the command line, enter the **asu show** command. The command output is shown in Example 7-1.

*Example 7-1 Output of the asu show command*

---

```
IMM2.SSH_SERVER_KEY=Installed
IMM2.SSL_HTTPS_SERVER_CERT=Private Key and Cert/CSR not available.
IMM2.SSL_HTTPS_SERVER_CSR=Private Key and Cert/CSR not available.
IMM2.SSL_LDAP_CLIENT_CERT=Private Key and Cert/CSR not available.
IMM2.SSL_LDAP_CLIENT_CSR=Private Key and Cert/CSR not available.
IMM2.SSL_SERVER_DIRECTOR_CERT=Private Key and Cert/CSR not available.
IMM2.SSL_SERVER_DIRECTOR_CSR=Private Key and Cert/CSR not available.
IMM2.SSL_CLIENT_TRUSTED_CERT1=Not-Installed
IMM2.SSL_CLIENT_TRUSTED_CERT2=Not-Installed
IMM2.SSL_CLIENT_TRUSTED_CERT3=Not-Installed
IMM2.PowerRestorePolicy=Restore
IMM2.ThermalModePolicy=Normal
```

---

## Show all UEFI settings

At the command line, enter the **asu show uefi** command. The command output is shown in Example 7-2.

*Example 7-2 Output of the asu show uefi command*

---

```
uEFI.OperatingMode=Custom Mode
uEFI.QuietBoot=Enable
uEFI.TurboModeEnable=Enable
uEFI.TurboBoost=Power Optimized
uEFI.ProcessorEistEnable=Enable
uEFI.ProcessorCcxEnable=Disable
uEFI.ProcessorC1eEnable=Enable
uEFI.HyperThreading=Enable
uEFI.EnableCoresInSbsp=All
uEFI.ExecuteDisableBit=Enable
uEFI.ProcessorVmxEnable=Enable
uEFI.ProcessorDataPrefetch=Enable
```

---

## Show all IMM2 settings

At the command line, enter the **asu show imm** command. The command output is shown in Example 7-3.

*Example 7-3 Output of the asu show imm command*

---

```
IMM2.SSH_SERVER_KEY=Installed
IMM2.SSL_HTTPS_SERVER_CERT=Private Key and Cert/CSR not available.
IMM2.SSL_HTTPS_SERVER_CSR=Private Key and Cert/CSR not available.
IMM2.SSL_LDAP_CLIENT_CERT=Private Key and Cert/CSR not available.
IMM2.SSL_LDAP_CLIENT_CSR=Private Key and Cert/CSR not available.
```

---

## Show power on/off state

At the command line, enter the **asu immapp powerstate --host [IP address]--user [USERID] --password [PASSWORD]** command. The command output is shown in Example 7-3.

*Example 7-4 Output of the power on/off state*

---

```
C:\asu>asu.exe immapp powerstate --host 10.10.10.1 --user USERID --password PASSWORD
Lenovo Advanced Settings Utility version 10.0.87F
Licensed Materials - Property of Lenovo
(C) Copyright Lenovo Corp. 2007-2012 All Rights Reserved
Connected to IMM at IP address 10.10.10.1
Server Power is currently On!
```

---

To power on and off, enter the following commands:

- ▶ `asu immapp poweronos --host [IP address]--user [USERID] --password [PASSWORD]`
- ▶ `asu immapp poweroffos --host [IP address]--user [USERID] --password [PASSWORD]`

## 7.9 MegaRAID Storage Manager

In this section, we describe the MegaRAID Storage Manager (MSM) software. For more information, see *Installation and User's Guide*, which is available at this website:

<https://support.lenovo.com/documents/MIGR-5073015>

With MSM, you can configure, monitor, and maintain storage configurations on ServeRAID-M controllers. The MegaRAID Storage Manager GUI makes it easy to create and manage storage configurations.

You can use MSM to manage local or remote RAID controllers and configure MSM for remote alert notifications. A command-line interface also is available.

To download the latest MegaRAID Storage Manager software and the *Installation and User's Guide*, see this website:

<https://support.lenovo.com/documents/MIGR-5073015>

### 7.9.1 MegaRAID Storage Manager installation

You must have administrator or root equivalent operating system privileges to install and fully access the MegaRAID Storage software.

The following setup options are available:

- ▶ **Complete:** This option installs all program features.
- ▶ **Client:** This option installs the required components to remotely view and configure servers.
- ▶ **Server:** This option installs only the required components for remote server management.
- ▶ **Stand Alone:** This option installs only the required components for local server management.

### 7.9.2 Drive states

A *drive group* is one or more drives that are controlled by the RAID controller.

The following drive states are possible:

- ▶ **Online**

This drive can be accessed by the RAID controller and is part of the virtual drive.

- ▶ **Unconfigured Good**  
This drive functions normally but is not configured.
- ▶ **Hot Spare**  
This drive is powered up and ready for use as a spare if an online drive fails. This drive can be Dedicated or Global Hot Spare.
- ▶ **Failed**  
A failed drive was originally configured as Online or Hot Spare, but the firmware detected an unrecoverable error.
- ▶ **Rebuild**  
Data is written to this drive to restore full redundancy for a virtual drive.
- ▶ **Unconfigured Bad**  
The firmware detects an unrecoverable error on this drive.
- ▶ **Missing**  
This drive was online but was removed from its location.
- ▶ **Offline**  
This drive is part of a virtual drive but has invalid data as far as the RAID configuration is concerned.

### 7.9.3 Virtual drive states

A *virtual drive* is a partition in a drive group that is made up of contiguous data segments on the drives.

The following virtual drive states are possible:

- ▶ **Optimal**  
The virtual drive operating condition is good. All configured drives are online.
- ▶ **Degraded**  
The virtual drive operating condition is not optimal. One of the configured drives failed or is offline.
- ▶ **Partially Degraded**  
The operating condition in an RAID-6 virtual drive is not optimal. One of the configured drives failed or is offline.
- ▶ **Failed**  
The virtual drive failed.
- ▶ **Offline**  
The virtual drive is not available to the RAID controller.

### 7.9.4 StorCLI utility for storage management

In this section, we describe the StorCLI utility. The StorCLI utility is a command-line interface (CLI) application. You can use this utility to configure, monitor, and maintain ServeRAID SAS RAID controllers and the devices that connect to them.

For more information about how to use StorCLI, see the *ServeRAID-M Software User's Guide*, available from:

<https://support.lenovo.com/docs/UM103976>

You can also use the **StorCLI -h** command for help.

**Note:** StorCLI is the successor to MegaCLI however MegaCLI commands can be run on the Storage Command Line (StorCLI) tool. A single binary is output for the StorCLI commands and its equivalent MegaCLI commands.

## Creating a virtual drive by using the CLI

In this example, we have two hard disk drives in slots one and two. Both hard disk drives must be in *Unconfigured Good* state, as shown in Figure 5-26 on page 179.

Complete the following steps to create a virtual drive by using the CLI:

1. Run the following command to locate the Enclosure Device ID and the Slot Number of both hard disk drives:

```
MegaCli -PDList -aAll
```

2. Run the following command to create an RAID-1:

```
MegaCli -CfgLdAdd -R1[252:1,252:2] -a0
```

Example 7-5 shows the resulting output.

*Example 7-5 Output from command MegaCli -CfgLdAdd -R1[252:1,252:2] -a0*

---

```
Adapter 0: Created VD 1
```

```
Adapter 0: Configured the Adapter!!
```

```
Exit Code: 0x00
```

---

3. The virtual drive is successfully created.

## More command examples

Table 7-9 lists command examples that use the MegaCli command and the equivalent StorCLI command.

Table 7-9 MegaCLI and StorCLI common commands

MegaCLI	StorCLI	Function
MegaCli -h -Help ?	storcli -help -h ?	Display help menus
MegaCli -AdpAllInfo -aALL	storcli /cx show all	Display controller properties for all installed adapters
MegaCLI -v	storcli -v	Display the software version on the controller
MegaCLI -adpCount	storcli show ctrlcount	Display the number of connected controllers
MegaCLI -CfgFreeSpaceInfo -aN -a0,1,2 -aALL	storcli /cx show freespac	Display available free space that is on the controller
MegaCLI-GetPreservedCacheList -aALL	storcli /cx show all	Display preserved cache status
MegaCLI -AdpGetTime -aN	storcli /cx show time	Display the controller time



MegaCLI	StorCLI	Function
MegaCli -AdpBIOS -Dsply -aALL	storcli /cx show bios	Schedule a consistency check
MegaCLI -AdpCcSched -Info	storcli /cx show cc/ConsistencyCheck	Display consistency check and parameters in progress, if any
MegaCli -AdpBbuCmd -GetBbuStatus -aN -a0,1,2 -aALL	storcli /cx/bbu show status	Display battery information, firmware status, and the gas gauge status

## 7.10 Managing NVMe drives and software RAID recovery

In this section we describe the NVMe drive replacement procedure and software RAID recovery for Linux and Windows operating systems. We show you how to locate a failed NVMe drive, how to gracefully hot-remove it from the server while the system is running and how to recover the software RAID after drive replacement.

### 7.10.1 Software RAID initialization in Linux

Linux natively supports software RAID technology and the **mdadm** utility is a standard RAID management tool available in most Linux distributions. mdadm supports the most common RAID levels like RAID-0, RAID-1, RAID-5, RAID-10. In this section we show how to initialize software RAID5 consisting of four NVMe drives on RHEL 7.2.

To create a new software RAID array, follow these steps:

1. Check that all installed NVMe drives are recognized by OS

Use `lspci` command to get the list of recognized NVMe drives, as shown in Figure 7-57:

```
[root@rhel7-n6h1ne8 ~]# lspci |egrep -i "ssd|953"
41:00.0 Non-Volatile memory controller: Intel Corporation PCIe Data Center SSD (rev 01)
49:00.0 Non-Volatile memory controller: Intel Corporation PCIe Data Center SSD (rev 01)
4e:00.0 Non-Volatile memory controller: Intel Corporation PCIe Data Center SSD (rev 01)
51:00.0 Non-Volatile memory controller: Intel Corporation PCIe Data Center SSD (rev 01)
```

Figure 7-57 `lspci` command output

As you can see in previous Figure 7-57 on page 295, Linux has found four NVMe drives installed in the server. At the beginning of each line you will see the unique PCIe address of each NVMe drive. You can get more information about any drive using its PCIe address, as shown in Figure 7-58:

```
[root@rhel7-n6h1ne8 ~]# lspci -s 49:00.0 -v
49:00.0 Non-Volatile memory controller: Intel Corporation PCIe Data Center SSD (rev 01) (prog-if 02 [NVM Express])
    Subsystem: Intel Corporation DC P3700 SSD [2.5" SFF]
    Physical Slot: 19
    Flags: bus master, fast devsel, latency 0, IRQ 38
    Memory at e7cfc000 (64-bit, non-prefetchable) [size=16K]
    Expansion ROM at e7c00000 [disabled] [size=64K]
    Capabilities: [40] Power Management version 3
    Capabilities: [50] MSI-X: Enable+ Count=32 Masked-
    Capabilities: [60] Express Endpoint, MSI 00
    Capabilities: [100] Advanced Error Reporting
    Capabilities: [150] Virtual Channel
    Capabilities: [180] Power Budgeting <?>
    Capabilities: [190] Alternative Routing-ID Interpretation (ARI)
    Capabilities: [270] Device Serial Number 55-cd-2e-41-4c-9d-08-a1
    Capabilities: [2a0] #19
    Kernel driver in use: nvme
```

Figure 7-58 Detailed lspci output for specific NVMe drive

As shown in previous Figure 7-58, using verbose (-v) mode of the lspci command you can get PCIe slot number and serial number of the drive, which we will use later during the drive replacement procedure.

2. Check that every NVMe drive has an associated block device. Use the ls command to locate NVMe block devices in /dev directory, as shown in Figure 7-59:

```
[root@rhel7-n6h1ne8 ~]# ls /dev/nvme*
/dev/nvme0 /dev/nvme0n1 /dev/nvme1 /dev/nvme1n1 /dev/nvme2 /dev/nvme2n1
/dev/nvme3 /dev/nvme3n1
```

Figure 7-59 List of NVMe block devices in /dev directory

Every NVMe drive is represented in the OS as /dev/nvmeXn1 device. As shown in previous Figure 7-59, four block devices were created.

You can also check the drive capacity using the parted utility, as shown in Figure 7-60:

```
[root@rhel7-n6h1ne8 ~]# parted /dev/nvme0n1 print
Error: /dev/nvme0n1: unrecognised disk label
Model: Unknown (unknown)
Disk /dev/nvme0n1: 1600GB
Sector size (logical/physical): 512B/512B
Partition Table: unknown
Disk Flags:
[root@rhel7-n6h1ne8 ~]#
```

Figure 7-60 parted utility output

3. Create a new software RAID using NVMe drives

Using the mdadm utility you can initialize a new array /dev/md0. In this example we create a RAID-5 array consisting of four drives /dev/nvme0n1 /dev/nvme1n1 /dev/nvme2n1 /dev/nvme3n1, as shown in Figure 7-61:

```
[root@rhel7-n6h1ne8 ~]# mdadm -C /dev/md0 --force --level=raid5 --bitmap=internal
--raid-devices=4 --assume-clean /dev/nvme0n1 /dev/nvme1n1 /dev/nvme2n1 /dev/nvme3n1
mdadm: Defaulting to version 1.2 metadata
mdadm: array /dev/md0 started.
```

Figure 7-61 Software RAID5 initialization

To check the status of the array, run the following commands, as shown in Figure 7-62 and Figure 7-63 on page 298:

```
[root@rhel7-n6h1ne8 ~]# cat /proc/mdstat
Personalities : [raid6] [raid5] [raid4]
md0 : active raid5 nvme3n1[3] nvme2n1[2] nvme1n1[1] nvme0n1[0]
      4688047104 blocks super 1.2 level 5, 512k chunk, algorithm 2 [4/4] [UUUU]
      bitmap: 0/12 pages [0KB], 65536KB chunk

unused devices: <none>
```

Figure 7-62 Array status

As you can see in previous Figure 7-62 on page 297, the array is in active state and all four drives are available.

```
[root@rhel7-n6h1ne8 ~]# mdadm --detail /dev/md0
/dev/md0:
    Version : 1.2
  Creation Time : Fri May 20 19:21:42 2016
    Raid Level : raid5
    Array Size : 4688047104 (4470.87 GiB 4800.56 GB)
  Used Dev Size : 1562682368 (1490.29 GiB 1600.19 GB)
  Raid Devices : 4
  Total Devices : 4
    Persistence : Superblock is persistent

  Intent Bitmap : Internal

    Update Time : Fri May 20 19:21:42 2016
      State : clean
Active Devices : 4
Working Devices : 4
Failed Devices : 0
  Spare Devices : 0

    Layout : left-symmetric
    Chunk Size : 512K

    Name : rhel7-n6h1ne8.poc.bts.lab:0 (local to host
rhel7-n6h1ne8.poc.bts.lab)
    UUID : 6acdc9c0:a56492f6:d13cfd69:fd81253
    Events : 0

    Number Major Minor RaidDevice State
       0     259     2        0    active sync  /dev/nvme0n1
       1     259     1        1    active sync  /dev/nvme1n1
       2     259     3        2    active sync  /dev/nvme2n1
       3     259     0        3    active sync  /dev/nvme3n1
```

Figure 7-63 Detailed information about the array

As you can see in the previous example (Figure 7-63), the total array size is 4480.56 GB, all drives are active and in sync state, and the array has no failed drives.

4. You can also use the mdadm command to generate config file, as shown in Figure 7-64:

```
[root@rhel7-n6h1ne8 ~]# mdadm -E -s
ARRAY /dev/md/0 metadata=1.2 UUID=6acdc9c0:a56492f6:d13cfd69:fd81253
name=rhel7-n6h1ne8.poc.bts.lab:0
[root@rhel7-n6h1ne8 ~]# mdadm -E -s >> /etc/mdadm.conf
```

Figure 7-64 mdadm configuration file creation

When you complete this procedure, you will have a working software RAID array. You can use it as a regular block device, in other words, you can create partitions and file systems and mount it to the file system tree.

## 7.10.2 NVMe drive hot-replacement in Linux

In this section we cover the hot-replacement procedure of failed NVMe drives in RHEL7.2. Hot-replacement means that we perform a graceful hot-swap procedure on the running system with no interruption in service or downtime.

**Note:** Not every Linux distribution supports NVMe drive hot-replacement; it depends on the Linux kernel version. Here is the list of distributions and Linux kernels that were validated at the time of writing:

- ▶ RHEL 7.0 and higher, kernel 3.10.0-123.el7.x86\_64 and higher
- ▶ RHEL 6.6 and higher, kernel 2.6.32-500.el6.x86\_64 and higher
- ▶ SLES 12, kernel 3.12.28-2 rc 3 and higher

According to Intel white paper, *Hot-Plug Capability of NVMe SSDs in Server Platforms*, you must set the following kernel parameter: `pci=pcie_bus_perf`. To do that, add that line as the kernel boot argument to the bootloader configuration file (`grub.cfg` or `elilo.conf`).

For more information please refer to the Intel white paper:

<http://www.intel.com/content/dam/www/public/us/en/documents/white-papers/333596-hot-plug-capability-nvme-ssds-paper.pdf>

In this section we simulate the outage of one of the NVMe drive to demonstrate the hot-replacement concept. Follow the described procedure below, to perform a graceful NVMe drive hot-replacement procedure:

1. Make sure that required Linux kernel is running and `pci` kernel parameter has required value.
2. Run the following command to check the running kernel version and its boot parameters, as shown in Figure 7-65:

```
[root@rhel7-n6h1ne8 ~]# cat /proc/cmdline
BOOT_IMAGE=/vmlinuz-3.10.0-327.el7.x86_64 root=/dev/mapper/rhel-root ro
rd.lvm.lv=rhel/root rd.lvm.lv=rhel/swap rhgb quiet pci=pcie_bus_perf
```

Figure 7-65 kernel boot parameters

3. Mark the failed NVMe drive as a “failed” drive in `mdadm` configuration

Let’s assume that one of the installed NVMe drive has failed, `/dev/nvme1n1`, for example. First, mark this drive as a “failed” drive in the `mdadm` configuration, as shown in Figure 7-66:

```
[root@rhel7-n6h1ne8 ~]# mdadm --manage /dev/md0 --fail /dev/nvme1n1
mdadm: set /dev/nvme1n1 faulty in /dev/md0
```

Figure 7-66 Failed drive designation

To verify that the array status has changed, run the following command, as shown in Figure 7-67:

```
[root@rhel7-n6h1ne8 ~]# cat /proc/mdstat
Personalities : [raid6] [raid5] [raid4]
md0 : active raid5 nvme1n1[1](F) nvme2n1[2] nvme3n1[3] nvme0n1[0]
      4688047104 blocks super 1.2 level 5, 512k chunk, algorithm 2 [4/3] [U_UU]
      bitmap: 0/12 pages [0KB], 65536KB chunk

unused devices: <none>
```

Figure 7-67 Array status

As you can see in Figure 7-67, nvme1n1 drive is in the failed state and the array now has only 3 active drives.

4. Determine the PCIe address and PCIe slot number used by the failed drive

Run two commands to locate the failed NVMe drive in the server. First, find out the PCIe address of the nvme1n1 drive. To do that, run the command shown in Figure 7-68:

```
[root@rhel7-n6h1ne8 ~]# find /sys/devices |egrep 'nvme1[0-9]?$'
/sys/devices/pci0000:40/0000:40:02.2/0000:47:00.0/0000:48:02.0/0000:49:00.0/nvme/nvme1
```

Figure 7-68 PCIe address location of the failed drive

As you can see, the failed nvme1n1 drive has PCIe address 0000:49:00.0. To determine the PCIe slot number of the failed drive, use the `lspci` command, as shown in Figure 7-58 on page 296, or un the following command, as shown in Figure 7-69:

```
[root@rhel7-n6h1ne8 ~]# grep '49:00' /sys/bus/pci/slots/*/address
/sys/bus/pci/slots/19/address:0000:49:00
```

Figure 7-69 PCIe slot number determination

As you can see, both commands show the same result – the nvme1n1 drive is located in PCIe slot 19, the upper drive bay in the Storage book (see Figure 5-38 on page 186).

5. Power off the failed NVMe drive

To gracefully power off the failed NVMe drive located in PCIe slot 19, run the following command, as shown in Figure 7-70:

```
[root@rhel7-n6h1ne8 ~]# echo 0 > /sys/bus/pci/slots/19/power
```

Figure 7-70 Power off the failed NVMe drive

Verify that the drive is shut down and is no longer represented in the OS using the `lspci` and `lsblk` commands, as shown in following Figure 7-71:

```
[root@rhel7-n6h1ne8 ~]# lspci |egrep -i "ssd|953"
41:00.0 Non-Volatile memory controller: Intel Corporation PCIe Data Center SSD (rev 01)
4e:00.0 Non-Volatile memory controller: Intel Corporation PCIe Data Center SSD (rev 01)
51:00.0 Non-Volatile memory controller: Intel Corporation PCIe Data Center SSD (rev 01)
[root@rhel7-n6h1ne8 ~]# lsblk
NAME            MAJ:MIN RM  SIZE RO TYPE  MOUNTPOINT
sda              8:0    0 744.1G  0 disk
├─sda1           8:1    0   200M  0 part  /boot/efi
├─sda2           8:2    0   128M  0 part
├─sda3           8:3    0 353.2G  0 part
├─sda4           8:4    0   500M  0 part  /boot
├─sda5           8:5    0   104G  0 part
├─┌─rhel-root   253:0    0   100G  0 lvm   /
│ └─rhel-swap   253:1    0     4G  0 lvm   [SWAP]
├─sda6           8:6    0     2G  0 part
├─sda7           8:7    0    80G  0 part
└─nvme0n1       259:2    0   1.5T  0 disk
    ├─md0        9:0    0   4.4T  0 raid5
    └─nvme2n1    259:3    0   1.5T  0 disk
        ├─md0        9:0    0   4.4T  0 raid5
        └─nvme3n1    259:0    0   1.5T  0 disk
            └─md0        9:0    0   4.4T  0 raid5
```

Figure 7-71 `lspci` and `lsblk` output

As you can see, `lspci` shows that now only three NVMe drives are available, `lsblk` also shows three drives: `nvme0n1`, `nvme2n1` and `nvme3n1`, which are combined in RAID-5.

#### 6. Replace the failed NVMe drive

As shown previously, the failed `nvme1n1` drive is located in PCIe slot 19, in the Storage Book bay 7. Now it is safe to remove the NVMe drive from the Storage Book. Then insert the new drive (same model and capacity) into the Storage Book.

### 7.10.3 NVMe drive hot-plug and software RAID recovery in Linux

As described in the previous section 7.10.2, “NVMe drive hot-replacement in Linux” on page 299, you can perform a hot-removal operation for the failed NVMe drive by following the procedure given. When the drive is replaced, power it on and then recover the related software RAID array. To do that, follow the next procedure:

#### 1. Power on the new drive

Similar to the power off procedure, as shown in Figure 7-70 on page 300, run the following command:

```
[root@rhel7-n6h1ne8 ~]# echo 1 > /sys/bus/pci/slots/19/power
```

Figure 7-72 Power on the new NVMe drive

If you insert the new NVMe drive in the same Storage Book bay where the failed drive was located, the PCIe slot number remains the same; in this example it is PCIe slot 19.

#### 2. Ensure that the new drive has been successfully started and recognized by the OS by using the `lspci` and `lsblk` commands, as shown in Figure 7-73:

```
[root@rhel7-n6h1ne8 ~]# lspci |egrep -i "ssd|953"
41:00.0 Non-Volatile memory controller: Intel Corporation PCIe Data Center SSD (rev 01)
49:00.0 Non-Volatile memory controller: Intel Corporation PCIe Data Center SSD (rev 01)
4e:00.0 Non-Volatile memory controller: Intel Corporation PCIe Data Center SSD (rev 01)
51:00.0 Non-Volatile memory controller: Intel Corporation PCIe Data Center SSD (rev 01)
[root@rhel7-n6h1ne8 ~]# lsblk
NAME                MAJ:MIN RM   SIZE RO TYPE MOUNTPOINT
sda                  8:0    0 744.1G 0 disk
  ..sda1             8:1    0  200M 0 part /boot/efi
  ..sda2             8:2    0   128M 0 part
  ..sda3             8:3    0 353.2G 0 part
  ..sda4             8:4    0   500M 0 part /boot
  ..sda5             8:5    0   104G 0 part
  · · ·rhel-root     253:0   0   100G 0 lvm  /
  · · ·rhel-swap     253:1   0    4G  0 lvm  [SWAP]
  ..sda6             8:6    0    2G  0 part
  ..sda7             8:7    0   80G  0 part
nvme0n1             259:2   0   1.5T 0 disk
  ..md0              9:0    0   4.4T 0 raid5
nvme1n1            259:1   0   1.5T 0 disk
nvme2n1             259:3   0   1.5T 0 disk
  ..md0              9:0    0   4.4T 0 raid5
nvme3n1             259:0   0   1.5T 0 disk
  ..md0              9:0    0   4.4T 0 raid5
```

Figure 7-73 *lspci and lsblk output*

As you can see, `lspci` has shown that the new NVMe drive in PCIe slot 10 has been recognized by the Linux kernel. The `lsblk` command also has shown that the appropriate block device – `nvme1n1` has been created. However, notice that `nvme1n1` is not associated with an array yet.

### 3. Recover the software RAID array

The existing RAID-5 now is in a degraded state: only 3 of 4 drive are available and active. You must add a new drive to the array to recover it. To do that, run the following command:

```
[root@rhel7-n6h1ne8 ~]# mdadm --manage /dev/md0 --add /dev/nvme1n1
mdadm: added /dev/nvme1n1
```

Figure 7-74 *Adding a new NVMe drive to the existing array*

You can check the RAID status using the commands in Figure 7-75.

```
[root@rhel7-n6h1ne8 ~]# cat /proc/mdstat
Personalities : [raid6] [raid5] [raid4]
md0 : active raid5 nvme1n1[4] nvme2n1[2] nvme3n1[3] nvme0n1[0]
      4688047104 blocks super 1.2 level 5, 512k chunk, algorithm 2 [4/3] [U_UU]
      [=====>.....] recovery = 31.8% (497008068/1562682368) finish=92.0min
      speed=192951K/sec
      bitmap: 0/12 pages [0KB], 65536KB chunk

unused devices: <none>
```

Figure 7-75 *Recovery process*



As you can see in previous Figure 7-75, the nvme1n1 drive has been added successfully and the array has started a recovery process. That may take some time and affect the array performance. When it's done, you will have a redundant and fully operating RAID.

### 7.10.4 NVMe drive hot-replacement and software RAID recovery in Windows

In this section we describe how to perform the hot-replacement procedure for a failed NVMe drive in Windows 2012 R2. By hot-replacement we mean a replacement procedure of the failed drive on the running system without any interruption in service or downtime. Then we show how to recover software RAID after drive replacement.

The NVMe drive and software RAID initialization in Windows 2012 R2 is very simple and straightforward. We don't cover this procedure in this section. For more information about NVMe drive initialization in Windows, refer to 5.4.4, "Using NVMe drives with Microsoft Windows Server 2012 R2" on page 192.

To demonstrate the hot-replacement procedure, assume there are four NVMe drives installed in the server and combined in one array using software RAID-5. The Windows Disk Management tool shows that we have disk X: (volume label "NVMe"), and a dynamic RAID-5 volume with capacity of 4470.87 GB. The status of the volume is Healthy, and all related NVMe drives are online, as shown in Figure 7-76:

Volume	Layout	Type	File S...	Status	Capacity
	Simple	Basic		Healthy (EFI System Partition)	200 MB
	Simple	Basic		Healthy (Primary Partition)	104.00 GB
(C:)	Simple	Basic	NTFS	Healthy (Boot, Page File, Cras...	353.17 GB
(D:)	Simple	Basic	RAW	Healthy (Primary Partition)	500 MB
NVMe (X:)	RAID-5	Dynamic	NTFS	Healthy	4470.87 GB

Disk	Layout	Type	File S...	Status	Capacity
Disk 0	Basic	744.00 GB	Online	200 I Heal	286.14 GB
		(C:)	353.17 GB NTFS	Healthy (Boot,	
		(D:)	500 I Heal	Healthy (Prim	Unallocated
Disk 1	Dynamic	1490.29 GB	Online	NVMe (X:) 1490.29 GB NTFS Healthy	
Disk 2	Dynamic	1490.29 GB	Online	NVMe (X:) 1490.29 GB NTFS Healthy	
Disk 3	Dynamic	1490.29 GB	Online	NVMe (X:) 1490.29 GB NTFS Healthy	
Disk 4	Dynamic	1490.29 GB	Online	NVMe (X:) 1490.29 GB NTFS Healthy	

Figure 7-76 Initial state of the array

Assume one of the NVMe drives has failed—**Disk 1**, for example. The array in this case is still operating, but it has failed redundancy. In **Disk Manager** you may see the following picture, as shown in Figure 7-77:






 <b>Disk 1</b> Unknown 1490.42 GB Not Initialized	1490.42 GB Unallocated
 <b>Disk 2</b> Dynamic 1490.29 GB Online	<b>NVMe (X:)</b> 1490.29 GB NTFS Failed Redundancy
 <b>Disk 3</b> Dynamic 1490.29 GB Online	<b>NVMe (X:)</b> 1490.29 GB NTFS Failed Redundancy
 <b>Disk 4</b> Dynamic 1490.29 GB Online	<b>NVMe (X:)</b> 1490.29 GB NTFS Failed Redundancy
 <b>Missing</b> Dynamic 1490.29 GB Missing	<b>NVMe (X:)</b> 1490.29 GB NTFS Failed Redundancy

Figure 7-77 Disk 1 is failed

To perform the hot-replacement and RAID recovery operations, follow the procedure described below:

1. Put the failed drive in Offline mode

To do that, right-click on the **Disk 1** and choose **Offline** from the pop-up menu, as shown in Figure 7-78.

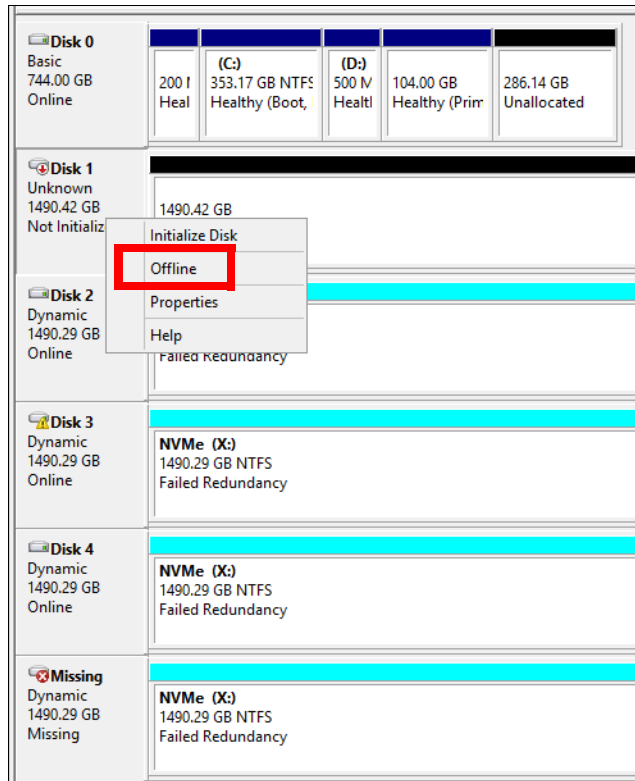


Figure 7-78 Offline mode

There are several steps to find out the physical location of a failed NVMe drive in the Storage book.

- Select the failed disk (**Disk 1**) in **Disk Manager** and open the **Properties** window, as shown in Figure 7-79:

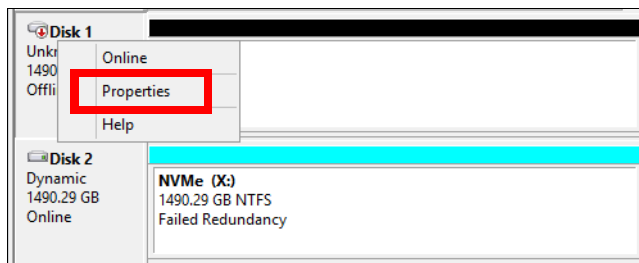


Figure 7-79 Disk Properties

- In the **Properties** window, on the **General** tab you can see the PCIe slot number of the associated NVMe drive, as shown in Figure 7-80:

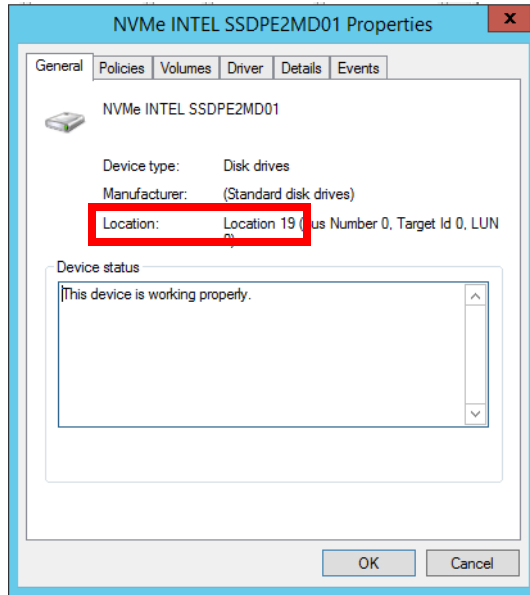


Figure 7-80 Properties window, General tab, disk location

As you can see on the **General** tab, the drive location is 19. That means the NVMe drive is located in PCIe slot 19 (bay 7 is the Storage book). For more information about NVMe drives location in the Storage book, refer to 5.4.1, “NVMe drive placement” on page 184.

2. Power off the failed NVMe drive

Next, shut down the device from OS. Open the **Devices and Printers** window. All NVMe drives installed in the server should appear, as shown in Figure 7-81:

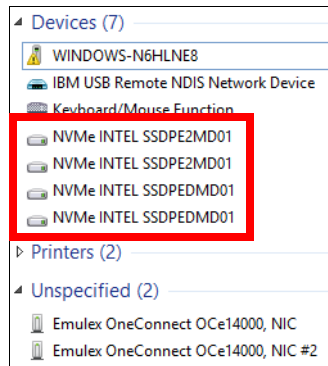


Figure 7-81 Devices and Printers

Right-click on one of the NVMe drives and select **Properties** from the pop-up menu, as shown in Figure 7-82:

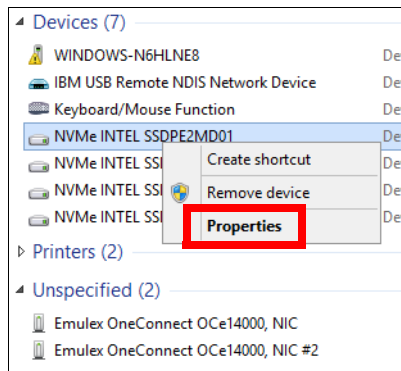


Figure 7-82 Device properties

Check the **Hardware** tab of the **Properties** window and locate the NVMe drive in Location 19, as shown in Figure 7-83:

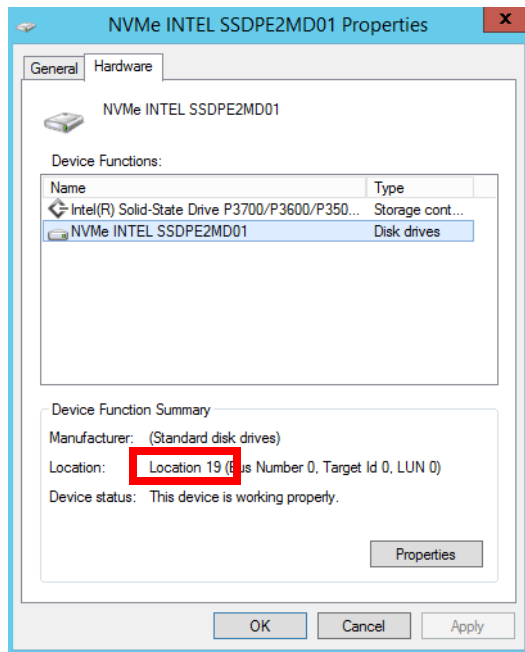


Figure 7-83 Drive location

Remove the failed drive from the OS. Right-click on the appropriate drive in the **Devices and Printers** window and select the **Remove device** action, as shown in Figure 7-84:

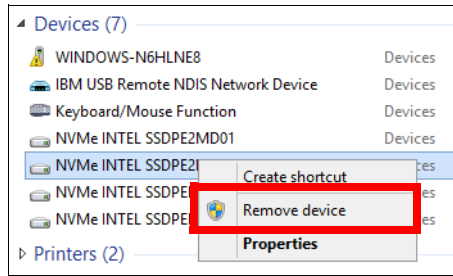


Figure 7-84 Remove device

3. Replace the failed NVMe drive

Next, replace the failed NVMe drive located in PCIe slot 19 (Storage book bay 7) using the same drive model with equal drive capacity.

4. Rescan available devices

When an NVMe drive replacement is performed, you must rescan available devices from the OS. To do that, open **Device Manager**, choose **Action** from menu and select **Scan for hardware changes**, as shown in Figure 7-85:

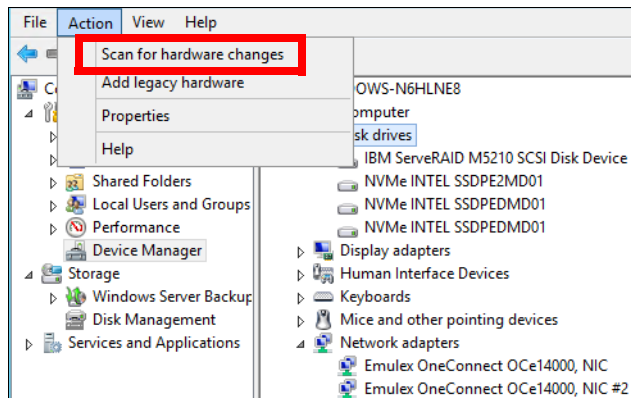


Figure 7-85 Scan for new devices

As you can see, only three NVMe drives are visible to the OS. When the scanning process is finished, you will see a new NVMe drive in the list of available devices:

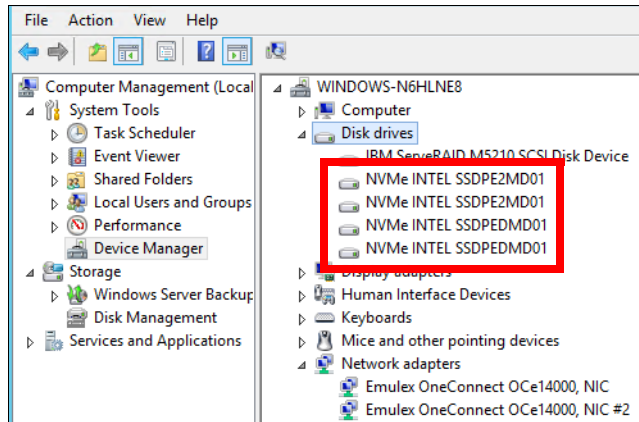


Figure 7-86 New NVMe drive is discovered

### 5. Initialize a new NVMe drive

Next, initialize the new NVMe drive installed in the server. Open **Disk Manager**, where you should see a new “not initialized” drive (**Disk 1**), as shown in Figure 7-87:

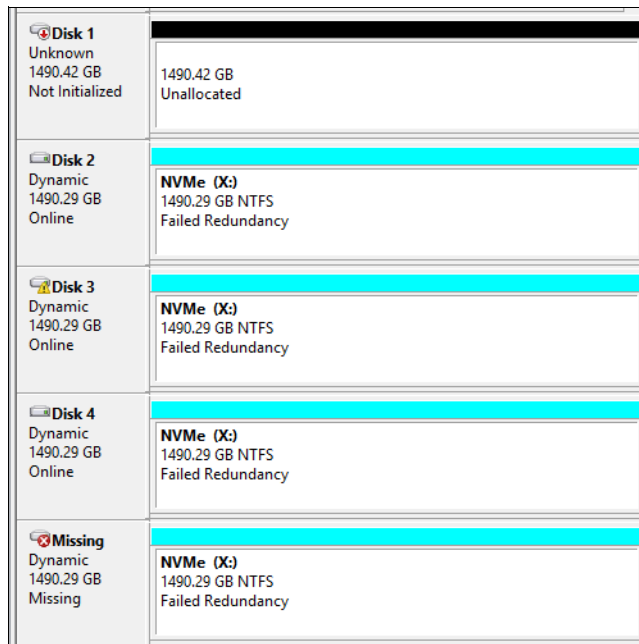


Figure 7-87 New not initialized drive

Right-click on the new drive (Disk 1) and select **Initialize disk**, as shown in Figure 7-88:

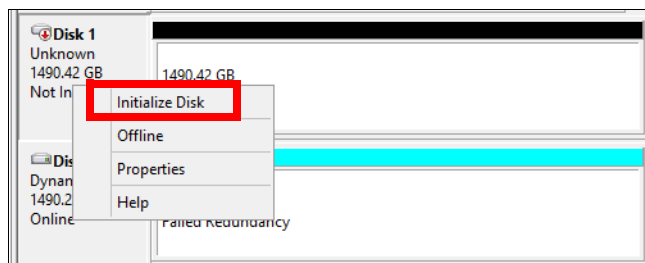


Figure 7-88 Initialize a new drive

## 6. Repair volume

When disk initialization is complete you can start the array recovery procedure. To do that, right-click on the **NVMe** volume (disk X:) and select the **Repair Volume** option from the pop-up menu:

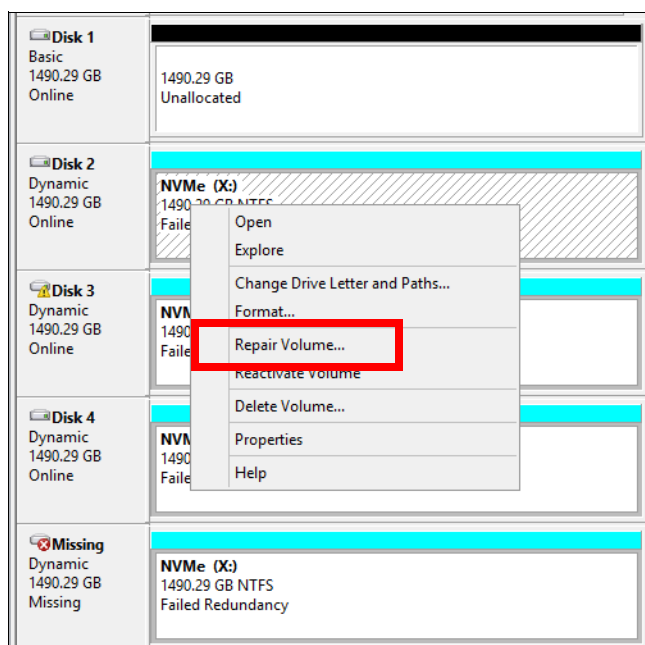


Figure 7-89 Repair Volume



Choose the new drive (Disk 1) to replace the failed drive in the array, as shown in Figure 7-90:

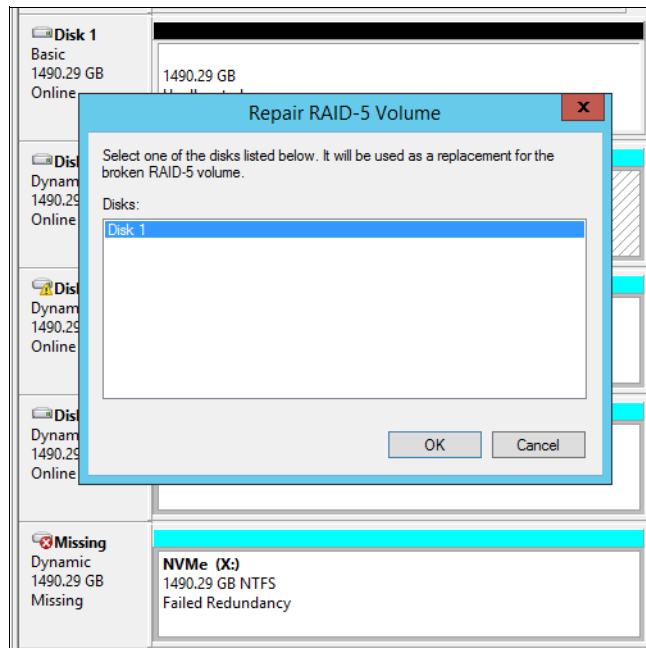


Figure 7-90 Select a new drive for the array

Confirm the action:

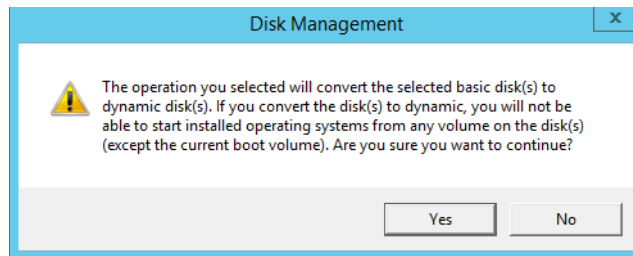


Figure 7-91 Confirm the action

## 7. Check the synchronization status

When the new drive is added to the array, the synchronization process will run automatically. Check the synchronization status of the array in **Disk Manager**, as shown in Figure 7-92:

Volume	Layout	Type	File System	Status
	Simple	Basic		Healthy (EFI System Partition)
	Simple	Basic		Healthy (Primary Partition)
(C:)	Simple	Basic	NTFS	Healthy (Boot, Page File, Crash Dump, Prim
(D:)	Simple	Basic	RAW	Healthy (Primary Partition)
NVMe (X:)	RAID-5	Dynamic	NTFS	Resynchronizing : (27%)

Disk	Layout	Type	File System	Status
Disk 0	Basic	744.00 GB	Online	
Disk 1	Dynamic	1490.29 GB	Online	
Disk 2	Dynamic	1490.29 GB	Online	
Disk 3	Dynamic	1490.29 GB	Online	
Disk 4	Dynamic	1490.29 GB	Online	

Figure 7-92 Resyncing process

When the resyncing process is complete, the array will become redundant and fully operational again.

## 7.11 Serial over LAN

*Serial over LAN* (SoL) is a mechanism that enables the input and output of the serial port of a managed system to be redirected in the network over TCP/IP. SoL provides a means to manage servers remotely by using a CLI over a Telnet or Secure Shell (SSH) connection.

SoL can give you remote access to your X6 servers UEFI and power-on self-test (POST) messages. By using SoL, you can log in to the machine remotely. It also can give you access to special operating system functions during start.

In the x3850 X6, the serial port is shared with the IMM2. The IMM2 can take control of the shared serial port to perform text console redirection and to redirect serial traffic by using SoL.

This section includes the following topics:

- ▶ 7.11.1, “Enabling SoL in UEFI” on page 313
- ▶ 7.11.2, “Enabling SoL in the operating system” on page 313

- ▶ 7.11.3, “Starting an SoL connection” on page 317

### 7.11.1 Enabling SoL in UEFI

To enable SoL from the UEFI interface at start, press F1 and select **System Settings** → **Devices and I/O Ports** → **Console Redirection Settings**. Table 7-10 lists the UEFI settings that must be set.

Table 7-10 Settings in UEFI for SoL

Value	x3850 X6 and x3950 X6
<b>General settings</b>	
COM Port 1	Enable
COM Port 2	Enable
Remote Console	Enable
Serial Port Sharing	Enable
Serial Port Access Mode	Shared
Legacy Option ROM Display	COM Port 2
<b>COM settings</b>	
Baud Rate	115200
Data Bits	8
Parity	None
Stop Bits	1
Terminal Emulation	VT100
Active After Boot	Enable
Flow Control	Hardware

**Tip:** Terminal Emulation can be set to VT100 or ANSI; however, when Linux operating systems are configured, ensure that the OS settings match the terminal emulation that is selected in the hardware.

### 7.11.2 Enabling SoL in the operating system

In this section, we describe the settings to enable SoL in the following operating systems:

- ▶ Windows Server 2008
- ▶ Windows Server 2012
- ▶ Linux

#### Windows Server 2008

Complete the following steps to enable the Microsoft Emergency Messaging Service (EMS) and the Special Administration Console (SAC). You must have administrator privileges to perform these steps:

1. Start a command prompt **Start** → **Run** → **cmd**.

2. Enter the **bootcfg** command. The command output is shown in Example 7-6.

*Example 7-6 Output of the bootcfg command*

---

```
C:\>bootcfg

Boot Loader Settings
-----
timeout:30
default:multi(0)disk(0)rdisk(0)partition(1)\WINDOWS

Boot Entries
-----
Boot entry ID:    1
OS Friendly Name: Windows Server 2003, Enterprise
Path:            multi(0)disk(0)rdisk(0)partition(1)\WINDOWS
OS Load Options: /noexecute=optout /fastdetect

C:\>
```

---

3. Examine the output. If there is more than one boot entry, determine the default entry.
4. Enable EMS by using the **bootcfg /ems on /port com2 /baud 115200 /id 1** command. As shown in Example 7-7, the default boot entry has the ID 1.

*Example 7-7 Output of the bootcfg /ems on /port com2 /baud 115200 /id 1 command*

---

```
C:\>bootcfg /ems on /port com2 /baud 115200 /id 1
SUCCESS: Changed the redirection port in boot loader section.
SUCCESS: Changed the redirection baudrate in boot loader section.
SUCCESS: Changed the OS entry switches for line "1" in the BOOT.INI file.
```

---

5. Run **bootcfg** again to verify that the EMS is activated, as shown in Example 7-8.

*Example 7-8 Output of the bootcfg command*

---

```
C:\>bootcfg

Boot Loader Settings
-----
timeout:          30
default:          multi(0)disk(0)rdisk(0)partition(1)\WINDOWS
redirect:         COM2
redirectbaudrate:115200

Boot Entries
-----
Boot entry ID:    1
OS Friendly Name: Windows Server 2003, Enterprise
Path:            multi(0)disk(0)rdisk(0)partition(1)\WINDOWS
OS Load Options: /noexecute=optout /fastdetect /redirect

C:\>
```

---

6. Restart the server to make the changes effective.

## Windows Server 2012

The Microsoft EMS is enabled on servers. Use the following syntax commands for more functions on Windows Server 2012:

```
bootcfg /ems {ON | OFF | EDIT} [/s <Computer> [/u <Domain>\<User> /p <Password>]]  
[/port {COM1 | COM2 | COM3 | COM4 | BIOSSET}] [/baud {9600 | 19200 | 38400 | 57600  
| 115200}] [/id <OSEntryLineNum>]
```

For more information about parameter descriptions, see this website:

<https://technet.microsoft.com/en-us/library/cc753847.aspx>

### ***IMM setting***

Complete the following steps to change the CLI mode for the COM port for use with EMS:

1. Log in to the web interface of the IMM2.
2. Browse to **IMM Management** → **IMM Properties**, as seen in Figure 7-93.

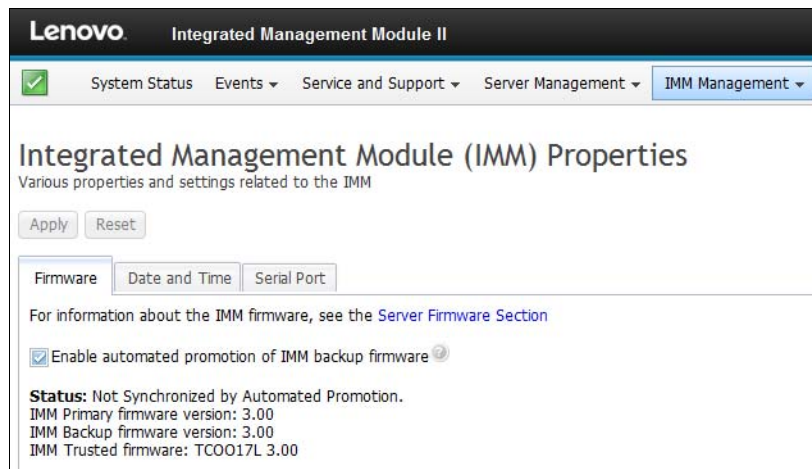


Figure 7-93 IMM2 properties pane

3. From the IMM2 properties window, select the **Serial Port** tab, as seen in Figure 7-94.

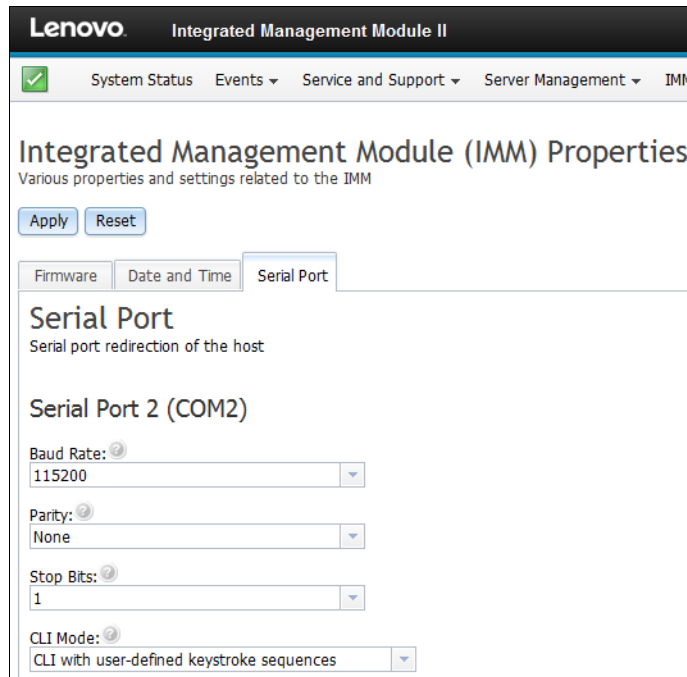


Figure 7-94 IMM2 Serial Redirect/CLI settings

4. Change the CLI mode to **CLI with EMS compatible keystroke sequences**.
5. Click **Apply** to save the changes.

For more information about Microsoft EMS and the SAC, see the following documents:

- ▶ *Boot Parameters to Enable EMS Redirection:*  
<http://msdn.microsoft.com/en-us/library/ff542282.aspx>
- ▶ *Special Administration Console (SAC) and SAC commands:*  
<http://msdn.microsoft.com/en-us/library/cc785873>

## Editing Linux files

You must edit two files in Linux to ensure that the console redirection still works after the operating system loads. The same files are changed for Red Hat Linux (RHEL) and SUSE Linux.

Edit the following files:

- ▶ /boot/grub/menu.lst
- ▶ /etc/inittab

**RHEL 6:** If you installed RHEL 6 in UEFI mode, you must edit the /boot/efi/EFI/redhat/grub.conf file instead of the /boot/grub/menu.lst file.

### ***Menu.lst or grub.conf***

Add the parameter that is highlighted in bold in the file as shown in Example 7-9.

*Example 7-9 Example of the grub.conf*

---

```
# grub.conf generated by anaconda
```

```

#
# Note that you do not have to rerun grub after making changes to this file
# NOTICE: You have a /boot partition. This means that
#           all kernel and initrd paths are relative to /boot/, eg.
#           root (hd0,1)
#           kernel /vmlinuz-version ro root=/dev/mapper/VolGroup-lv_root
#           initrd /initrd-[generic-]version.img
#boot=/dev/sda1
device (hd0) HD(1,800,64000,699900f5-c584-4061-a99f-d84c796d5c72)
default=0
timeout=5
splashimage=(hd0,1)/grub/splash.xpm.gz
hiddenmenu
title Red Hat Enterprise Linux (2.6.32-71.el6.x86_64)
    root (hd0,1)
    kernel /vmlinuz-2.6.32-71.el6.x86_64 ro root=/dev/mapper/VolGroup-lv_roo
t rd_LVM_LV=VolGroup/lv_root rd_LVM_LV=VolGroup/lv_swap rd_NO_LUKS rd_NO_MD rd_N
O_DM LANG=en_US.UTF-8 SYSFONT=latarcyrheb-sun16 KEYBOARDTYPE=pc KEYTABLE=de cras
hkernel=auto console=ttyS1,115200n8 rhgb quiet
    initrd /initramfs-2.6.32-71.el6.x86_64.img
[root@localhost redhat]#

```

---

### ***/etc/inittab***

Add the parameter that is highlighted in bold at the end of the `/etc/inittab` file, as shown in Example 7-10.

*Example 7-10 The `/etc/inittab` file*

```

id:3:initdefault:
co:2345:respawn:/sbin/agetty ttyS1 115200 vt100-nav
[root@localhost etc]#

```

---

## 7.11.3 Starting an SoL connection

In this section, we describe how to connect to X6 through SoL.

### **Connecting to the X6 server by using SoL**

Complete the following steps to use Telnet or SSH to connect to the X6 server:

1. Start a Telnet/SSH session to the IMM2 IP address.
2. Log in to the IMM2. The default user id is USERID and the default password is PASSWORD, where the 0 is a zero.
3. The IMM CLI main page opens, as seen Figure 7-95.



*Figure 7-95 IMM CLI main page*

4. Start SoL by running the **console 1** command.

The SoL console starts and you see whatever is transmitted over the SoL connection (for example, the UEFI setup windows).

For more information about the use of SoL commands, see the Integrated Management Module II User's Guide:

<https://support.lenovo.com/docs/UM103336>



# Abbreviations and acronyms

<b>AC</b>	alternating current	<b>ECC</b>	error checking and correcting
<b>ACID</b>	atomicity, consistency, isolation, and durability	<b>EIA</b>	Electronic Industries Alliance
<b>ACPI</b>	advanced control and power interface	<b>EMEA</b>	Europe, Middle East, Africa
<b>AES</b>	Advanced Encryption Standard	<b>EMS</b>	Emergency Messaging Service
<b>AES-NI</b>	Advanced Encryption Standard New Instructions	<b>ERP</b>	enterprise resource planning
<b>AMM</b>	Advanced Management Module	<b>ESA</b>	Electronic Service Agent
<b>ANSI</b>	American National Standards Institute	<b>ETS</b>	Enhanced Technical Support
<b>APIC</b>	Advanced Programmable Interrupt Controller	<b>FC</b>	Fibre Channel
<b>ASU</b>	Advanced Settings Utility	<b>FDR</b>	fourteen data rate
<b>BIOS</b>	basic input output system	<b>FSM</b>	Flex System Manager
<b>BM</b>	bridge module	<b>GB</b>	gigabyte
<b>BMC</b>	Baseboard Management Controller	<b>GPU</b>	Graphics Processing Unit
<b>BTU</b>	British Thermal Unit	<b>GT</b>	Gigatransfers
<b>CD</b>	compact disk	<b>GUI</b>	graphical user interface
<b>CIM</b>	Common Information Model	<b>HBA</b>	host bus adapter
<b>CLI</b>	command-line interface	<b>HD</b>	high definition
<b>CMOS</b>	complementary metal oxide semiconductor	<b>HDD</b>	hard disk drive
<b>CNA</b>	Converged Network Adapter	<b>HPC</b>	high performance computing
<b>COM</b>	Component Object Model	<b>HS</b>	hot-swap
<b>CPU</b>	central processing unit	<b>HTTP</b>	Hypertext Transfer Protocol
<b>CRC</b>	cyclic redundancy check	<b>HV</b>	high voltage
<b>CRM</b>	Customer Relationship Management	<b>I/O</b>	input/output
<b>CRU</b>	customer replaceable units	<b>IB/E</b>	InfiniBand/Ethernet
<b>CTO</b>	configure-to-order	<b>IBM</b>	International Business Machines
<b>DC</b>	domain controller	<b>ID</b>	identifier
<b>DCS</b>	Data Center Services	<b>IEC</b>	International Electrotechnical Commission
<b>DCU</b>	data cache unit	<b>IMM</b>	integrated management module
<b>DDR</b>	Double Data Rate	<b>IOPS</b>	I/O operations per second
<b>DHCP</b>	Dynamic Host Configuration Protocol	<b>IP</b>	Internet Protocol
<b>DIMM</b>	dual inline memory module	<b>IPMI</b>	Intelligent Platform Management Interface
<b>DNS</b>	Domain Name System	<b>ISD</b>	IBM Systems Director
<b>DSA</b>	Dynamic System Analysis	<b>ISO</b>	International Organization for Standards
<b>DVD</b>	Digital Video Disc	<b>IT</b>	information technology
<b>DW</b>	data warehousing	<b>ITSO</b>	International Technical Support Organization
		<b>JBOD</b>	just a bunch of disks
		<b>KB</b>	kilobyte
		<b>KCS</b>	keyboard console style

<b>KVM</b>	keyboard video mouse	<b>RAS</b>	remote access services; row address strobe
<b>KW</b>	kilowatts	<b>RDHX</b>	rear door heat exchanger
<b>LAN</b>	local area network	<b>RDIMM</b>	registered DIMM
<b>LC</b>	line cord	<b>RHEL</b>	Red Hat Enterprise Linux
<b>LCD</b>	liquid crystal display	<b>RISC</b>	reduced instruction set computer
<b>LDAP</b>	Lightweight Directory Access Protocol	<b>ROC</b>	RAID-on-card
<b>LED</b>	light emitting diode	<b>ROM</b>	read-only memory
<b>LER</b>	Live Error Recovery	<b>RPM</b>	revolutions per minute
<b>LP</b>	low profile	<b>RSA</b>	Remote Supervisor Adapter
<b>LRDIMM</b>	load-reduced DIMM	<b>RSS</b>	Receive-side scaling
<b>LUN</b>	logical unit number	<b>SAC</b>	Special Administration Console
<b>LV</b>	low voltage	<b>SAN</b>	storage area network
<b>MAC</b>	media access control	<b>SAS</b>	Serial Attached SCSI
<b>MB</b>	megabyte	<b>SATA</b>	Serial ATA
<b>MCA</b>	Machine Check Architecture	<b>SBB</b>	system building block
<b>MESI</b>	modified exclusive shared invalid	<b>SCM</b>	Supply Chain Management
<b>MLC</b>	multi-level cell	<b>SCO</b>	Santa Cruz Operation, Inc
<b>MR</b>	MegaRAID	<b>SED</b>	self-encrypting drive
<b>MSM</b>	MegaRAID Storage Manager	<b>SFP</b>	small form-factor pluggable
<b>MTM</b>	machine type model	<b>SLC</b>	Single Level Cell
<b>NBD</b>	next business day	<b>SLES</b>	SUSE Linux Enterprise Server
<b>NEMA</b>	National Electrical Manufacturers Association	<b>SMI</b>	scalable memory interconnect
<b>NI</b>	new instructions	<b>SMX</b>	Safer Mode Extensions
<b>NL</b>	nearline	<b>SNMP</b>	Simple Network Management Protocol
<b>NMI</b>	non-maskable interrupt	<b>SQL</b>	Structured Query Language
<b>NOS</b>	network operating system	<b>SSD</b>	solid state drive
<b>NTLM</b>	NT LAN Manager	<b>SSH</b>	Secure Shell
<b>NUMA</b>	Non-Uniform Memory Access	<b>SSL</b>	Secure Sockets Layer
<b>OLAP</b>	online analytical processing	<b>SSM</b>	Service and Support Manager
<b>OLTP</b>	online transaction processing	<b>TB</b>	terabyte
<b>OS</b>	operating system	<b>TCP/IP</b>	Transmission Control Protocol/Internet Protocol
<b>PCI</b>	Peripheral Component Interconnect	<b>TDP</b>	thermal design power
<b>PCI-SIG</b>	PCI Special Interest Group	<b>TPM</b>	Trusted Platform Module
<b>PDU</b>	power distribution unit	<b>TTY</b>	teletypewriter
<b>PE</b>	Preinstallation Environment	<b>TXT</b>	text
<b>PFA</b>	Predictive Failure Analysis	<b>UCO</b>	USB Conversion Option
<b>POST</b>	power-on self test	<b>UEFI</b>	Unified Extensible Firmware Interface
<b>PXE</b>	Preboot eXecution Environment	<b>UIM</b>	Upward Integration Module
<b>QPI</b>	QuickPath Interconnect	<b>UPS</b>	uninterruptible power supply
<b>RAID</b>	redundant array of independent disks	<b>USB</b>	universal serial bus
<b>RAM</b>	random access memory	<b>UXSP</b>	UpdateXpress System Packs

<b>UXSPI</b>	UpdateXpress System Packs Installer
<b>VAC</b>	Volts AC
<b>VD</b>	virtual drive
<b>VFA</b>	Virtual Fabric Adapter
<b>VGA</b>	video graphics array
<b>VLAN</b>	virtual LAN
<b>VM</b>	virtual machine
<b>VPD</b>	vital product data
<b>VPI</b>	Virtual Protocol Interconnect
<b>VT</b>	Virtualization Technology
<b>WW</b>	world wide



# Related publications

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

## Lenovo Press publications

The following Lenovo Press publications provide more information about the topics in this document. You can view these and other books, papers, Product Guides, and videos at the Lenovo Press website:

<http://lenovopress.com>

### Product Guides

The following Product Guides are available:

- ▶ Lenovo System x3850 X6:  
<http://lenovopress.com/tips1250>
- ▶ Lenovo System x3950 X6:  
<http://lenovopress.com/tips1251>

### Books and papers

The following books and papers are available:

- ▶ In-memory Computing with SAP HANA on Lenovo X6 Systems:  
<https://lenovopress.com/sg248086>
- ▶ SAP Business Suite on Lenovo X6 Systems:  
<http://lenovopress.com/redp5073>
- ▶ Oracle and Lenovo X6 Validated Configuration Guide  
<https://lenovopress.com/redp5218>
- ▶ Optimizing Memory Performance of Lenovo Servers Based on Intel Xeon E7 v3 Processors  
<https://lenovopress.com/lp0048>
- ▶ Quantifying the Power Savings by Upgrading to DDR4 Memory on Lenovo Servers  
<https://lenovopress.com/lp0083>
- ▶ Lenovo XClarity Administrator Planning and Implementation Guide:  
<http://lenovopress.com/sg248296>

## Other publications

For more information, see the Information Center, including the *Installation and Service Guide* for the x3850 X6 and x3950 X6, which is available at this website:

[http://publib.boulder.ibm.com/infocenter/systemx/documentation/topic/com.lenovo.sysx.6241.doc/product\\_page.html](http://publib.boulder.ibm.com/infocenter/systemx/documentation/topic/com.lenovo.sysx.6241.doc/product_page.html)

## Online resources

For more information, see the following resources:

- ▶ Lenovo System x3850 X6 and x3950 X6 product pages:
  - <http://shop.lenovo.com/us/en/systems/servers/mission-critical/x3850-x6/>
  - <http://shop.lenovo.com/us/en/systems/servers/mission-critical/x3950-x6/>
- ▶ Lenovo Information Center
  - Installation and Service Guide
  - Rack Installation Instructions

[http://publib.boulder.ibm.com/infocenter/systemx/documentation/topic/com.lenovo.sysx.6241.doc/product\\_page.html](http://publib.boulder.ibm.com/infocenter/systemx/documentation/topic/com.lenovo.sysx.6241.doc/product_page.html)
- ▶ ServerProven hardware compatibility page for the x3850 X6 and x3950 X6
  - E7 v2: <http://www.lenovo.com/us/en/serverproven/xseries/6241.shtml>
  - E7 v3: <http://www.lenovo.com/us/en/serverproven/xseries/6241E7xxxxV3.shtml>
  - E7 v4: <http://www.lenovo.com/us/en/serverproven/xseries/6241E7xxxxV4.shtml>
- ▶ Power Guides:
  - <https://support.lenovo.com/documents/LNVO-POWINF>
- ▶ Power Configurator:
  - <https://support.lenovo.com/documents/LNVO-PWRCONF>
- ▶ Configuration and Option Guide:
  - <https://support.lenovo.com/documents/SCOD-3ZVQ5W>
- ▶ xREF - System x Reference:
  - <http://lenovopress.com/xref>
- ▶ Lenovo Support Portal:
  - x3850 X6:
    - <http://support.lenovo.com/us/en/products/Servers/Lenovo-x86-servers/Lenovo-System-x3850-X6/6241>
  - x3950 X6:
    - <http://support.lenovo.com/us/en/products/Servers/Lenovo-x86-servers/Lenovo-System-x3950-X6/6241>
- ▶ IBM System Storage Interoperation Center:
  - <http://www.ibm.com/systems/support/storage/ssic>

**Lenovo**

**Lenovo System x3850 X6 and x3950 X6 Planning and Implementation Guide**



(0.5" spine)  
0.475" x 0.873"  
250 <-> 459 pages









# Lenovo System x3850 X6 and x3950 X6 Planning and Implementation Guide

---

**Covers the sixth generation Enterprise X-Architecture servers**

**Provides technical information about all server features**

**Explains what planning activities you need to do**

**Describes how to manage the server with Lenovo XClarity Administrator**

The increasing demand for cloud computing and business analytical workloads by enterprises to meet business needs drives innovation to find new ways to build informational systems. Clients are looking for cost-optimized fit-for-purpose IT solutions that manage large amounts of data, easily scale performance, and provide reliable real-time access to actionable information.

Built on decades of innovation, the Lenovo X6 servers are fast, agile, and resilient.

Fast application performance means immediate access to actionable information.

Agile system design helps to reduce acquisition costs and host multiple generations of technology in a single server.

Resilient platforms maximize application uptime and promote easy integration in virtual environments.

Lenovo X6 servers continue to lead the way as the shift toward mission-critical scalable databases, business analytics, virtualization, enterprise applications, and cloud applications accelerates.

This Lenovo Press book provides product, planning, and implementation information. In the first few chapters, we provide detailed technical information about the four-socket x3850 X6 and eight-socket x3950 X6. This information is most useful in designing, configuring, and planning to order a server solution. In the later chapters of the book, we provide detailed configuration and setup information to get your server operational.

This book is for clients, Lenovo Business Partners, and Lenovo employees that want to understand the features and capabilities of the X6 portfolio of servers and want to learn how to install and configure the servers for use in production.



---

**BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE**

**At Lenovo Press, we bring together experts to produce technical publications around topics of importance to you, providing information and best practices for using Lenovo products and solutions to solve IT challenges.**