

The Lenovo logo, consisting of the word "Lenovo" in white sans-serif font on a dark grey rectangular background.

Lenovo ThinkSystem SR950 Planning and Implementation Guide

**Covers the new flagship
high-performance ThinkSystem
server**

**Provides technical information about
all server features**

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

**Describes how to manage the server
with various Lenovo management
tools**

David Watts
Rani Doughty





Lenovo ThinkSystem SR950 Planning and Implementation Guide

31 January 2019

Note: Before using this information and the product it supports, read the information in “Notices” on page 275.

Last update on 31 January 2019

This edition applies to Lenovo ThinkSystem SR950, machine types 7X11, 7X12, and 7X13, with Intel Xeon Scalable Family processors.

© Copyright Lenovo 2019. All rights reserved.

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

Contents

Preface	vii
The team who wrote this book	vii
Comments welcome	viii
Do you have the latest version?	viii
 Summary of changes	 ix
31 January 2019	ix
23 October 2018	x
3 October 2018	x
11 September 2018	x
 Chapter 1. Introduction	 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
1.2.1 Scalability and performance	5
1.2.2 Availability and serviceability	6
1.2.3 Manageability and security	7
1.2.4 Energy efficiency	7
1.3 Positioning	8
1.4 Storage versus in-memory data	10
1.5 Flash storage	11
1.6 Energy efficiency	12
1.7 Services offerings	12
1.8 About this book	13
 Chapter 2. SR950 Technology	 15
2.1 System design	16
2.1.1 Compute trays	16
2.1.2 Storage tray	19
2.1.3 I/O tray	20
2.1.4 Power supplies	22
2.1.5 Fans	23
2.1.6 Mid-Chassis interconnects	24
2.2 System architecture	26
2.3 Processors	28
2.3.1 Gold and Platinum product family	29
2.3.2 Gold and Platinum features	30
2.4 Memory	36
2.4.1 Memory RAS modes	37
2.4.2 Memory protection features	38
2.5 Internal storage	40
2.5.1 SAS/SATA SSD technology	40
2.5.2 PCIe NVMe drives	41
2.5.3 M.2 drives	42
2.5.4 Flash Storage Adapters	43

2.6 Unified Extensible Firmware Interface (UEFI)	43
2.7 xClarity Controller (XCC)	44
2.8 Scalability	45
Chapter 3. Product information	47
3.1 Components and connectors	48
3.2 Server configurations	52
3.3 Server configuration upgrades	56
3.4 Standard specifications	58
3.5 Models	60
3.6 Processor options	61
3.7 Memory options	64
3.8 Internal storage	66
3.8.1 Backplanes and drive bays	66
3.8.2 Adapters and cabling	69
3.8.3 M.2 drives	70
3.8.4 SED encryption key management with ISKLM	71
3.9 Controllers for internal storage	72
3.10 Internal drive options	74
3.11 Optical drives	78
3.12 I/O expansion options	78
3.12.1 Slot availability by server configuration	81
3.13 Network adapters	83
3.14 Fibre Channel host bus adapters	86
3.15 SAS adapters for external storage	86
3.16 Flash Storage adapters	88
3.17 GPU adapters and coprocessors	88
3.18 Cooling	88
3.19 Power supplies	89
3.20 Integrated virtualization	90
3.21 Systems management	90
3.21.1 Local management	90
3.21.2 System status with XClarity Mobile	91
3.21.3 Remote management	92
3.21.4 Lenovo XClarity Administrator	93
3.21.5 Lenovo XClarity Essentials	93
3.21.6 Lenovo XClarity Energy Manager	94
3.22 Security	94
3.23 Operating system support	94
Chapter 4. Infrastructure planning	97
4.1 Physical and electrical specifications	98
4.2 Operating environment	99
4.3 Rack selection and rack options	101
4.4 Rack installation	103
4.4.1 Installing in a non-Lenovo rack	104
4.5 Floor clearance	104
4.6 Using Rear Door Heat eXchanger (RDHX)	105
4.7 Power guidelines	107
4.7.1 Considerations	107
4.7.2 Power supply redundancy	108
4.7.3 Power supply installation order	109
4.8 Power monitoring	109

4.8.1	XClarity Controller power and cooling	109
4.8.2	Lenovo XClarity Administrator power and cooling	112
4.9	Cooling considerations	112
4.10	Uninterruptible power supply units	113
4.11	PDU and line cord selection	114
4.11.1	Server-to-PDU power cord options	114
4.11.2	PDU selection	116
Chapter 5.	Preparing the hardware	119
5.1	Preparing the hardware with XClarity Provisioning Manager (UEFI)	120
5.1.1	Accessing Lenovo XClarity Provisioning Manager	120
5.1.2	Viewing system inventory	124
5.1.3	UEFI settings	125
5.1.4	Operating modes	127
5.1.5	Configuring common UEFI settings	128
5.1.6	Updating firmware with LXPM	134
5.1.7	Configuring RAID with LXPM	135
5.1.8	Configure Networking with LXPM	139
5.1.9	Installing an Operating System with LXPM	141
5.2	Preparing the hardware with XClarity Controller	142
5.2.1	Accessing XCC	142
5.2.2	Trouble shooting the XCC communications	147
5.2.3	Updating the firmware with XCC	147
5.2.4	Configure power policies with XCC	149
5.2.5	Configuring RAID with XCC	153
5.2.6	Configure networking with XCC	157
5.2.7	Installing an operating system with XCC	158
5.3	Other Lenovo tools for updating firmware	159
5.3.1	XClarity Essentials UpdateXpress	159
5.3.2	XClarity Essentials Bootable Media Creator (BoMC)	165
5.3.3	XClarity Essentials OneCLI	171
5.3.4	Lenovo XClarity Administrator	173
5.4	Overview of firmware deployment tools: Using the right tool	176
5.5	Backing up the server configuration	177
5.6	Tethering with XClarity Mobile	178
5.6.1	XClarity mobile app pre-requisites	179
5.6.2	Installing and setting up XClarity Mobile app	180
5.6.3	Connecting XClarity Mobile directly to the SR950 server	181
Chapter 6.	Operating system installation	183
6.1	Options for OS installation	184
6.2	Installing an OS with Lenovo tools	185
6.2.1	Installing an OS with XClarity Controller (XCC)	185
6.2.2	Installing an OS with Lenovo XClarity Provisioning Manager (LXPM)	190
6.2.3	Installing an OS with XClarity Administrator (LXCA)	195
6.2.4	Lenovo XClarity Bootable Media Creator (BoMC)	200
6.2.5	Installing an OS with the local USB port	200
6.3	Operating system installation examples	200
6.3.1	Installing Windows Server 2012 R2 and 2016	201
6.3.2	Installing SUSE Linux Enterprise Server 12 SP2	205
6.3.3	Installing Red Hat Enterprise Linux 7.3	211
6.3.4	Installing Lenovo Customized VMware ESXi 6.5	216
6.4	Installing an OS on the M.2 adapter	220

6.5 Booting from SAN	221
Chapter 7. Managing the hardware	223
7.1 Management tools overview	224
7.2 Lenovo XClarity Administrator	225
7.2.1 Discovering the SR950 server	226
7.2.2 Remote control	230
7.2.3 Hardware monitoring	232
7.2.4 Updating firmware	233
7.2.5 Operating system deployment	233
7.2.6 Call home service	233
7.3 Lenovo XClarity Integrators	236
7.4 Lenovo XClarity Energy Manager	240
7.4.1 Licensing	241
7.4.2 System requirements	242
7.4.3 Download and install	242
7.4.4 Launching Energy Manager	247
7.4.5 Setting up Energy Manager	249
7.4.6 Energy optimization with Energy Manager	258
7.5 Lenovo XClarity Controller management functions	263
7.5.1 XCC functions to diagnose and manage the server	263
7.5.2 Remote control	266
7.5.3 Powering the server on and off with XCC	269
7.6 Lenovo XClarity Essentials	271
7.7 LCD system information display panel	272
Related publications	273
Lenovo Press publications	273
Other publications and online resources	274
Notices	275
Trademarks	276

Preface

Lenovo® ThinkSystem™ SR950 is a high-performance server that is designed for your most demanding, mission-critical workloads. Suitable workloads include in-memory databases, large transactional databases, batch and real-time analytics, ERP, CRM, and virtualized server applications.

The powerful 4U ThinkSystem SR950 can grow from two to eight Intel Xeon Scalable Family processors, and with 96 DIMM sockets, supporting up to 12 TB of high-speed memory. The modular design of the SR950 aids upgrades and servicing with easy front or rear access to all major subsystems to maximize server availability.

The SR950 packs numerous fault-tolerant and high-availability features into a high-density, 4U rack-optimized design that reduces the space needed to support massive network computing operations and simplify servicing.

This book provides a general understanding of SR950 technology, components and management features. It also describes preparing and implementing the server, including power and cooling considerations, operating system installations, systems management, and firmware update tools.

This book is aimed at customers, Lenovo Business Partners, and Lenovo employees who want to understand the features and capabilities of the ThinkSystem SR950 and want to learn how to install and configure the server 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 Consulting IT Specialist in the Lenovo Enterprise Business Group in Morrisville, North Carolina in the USA. He manages residencies and produces pre-sale and post-sale technical publications for hardware and software topics that are related to System x®, Flex System™, and BladeCenter® servers. He has authored over 300 books and papers. Prior to working for Lenovo starting in 2014, David had worked for IBM both in the U.S. and Australia since 1989. 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.

Comments welcome

Your comments are important to us!

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

- ▶ Use the online feedback form found at the web page for this document:
<http://lenovopress.com/lp0746>
- ▶ Send your comments in an email to:
comments@lenovopress.com

Do you have the latest version?

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

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.

31 January 2019

New RAID adapter - 3.9, "Controllers for internal storage" on page 72:

- ▶ ThinkSystem RAID 730-8i 2GB Flash PCIe 12Gb Adapter, 4Y37A09722

New drive options - 3.10, "Internal drive options" on page 74:

- ▶ ThinkSystem 2.5" SS530 400GB Performance SAS 12Gb Hot Swap SSD
- ▶ ThinkSystem 2.5" SS530 800GB Performance SAS 12Gb Hot Swap SSD
- ▶ ThinkSystem 2.5" SS530 1.6TB Performance SAS 12Gb Hot Swap SSD
- ▶ ThinkSystem 2.5" SS530 3.2TB Performance SAS 12Gb Hot Swap SSD
- ▶ ThinkSystem U.2 Intel P4610 1.6TB Mainstream NVMe PCIe3.0 x4 Hot Swap SSD
- ▶ ThinkSystem U.2 Intel P4610 3.2TB Mainstream NVMe PCIe3.0 x4 Hot Swap SSD
- ▶ ThinkSystem U.2 Intel P4610 6.4TB Mainstream NVMe PCIe3.0 x4 Hot Swap SSD
- ▶ ThinkSystem U.2 Intel P4510 1.0TB Entry NVMe PCIe3.0 x4 Hot Swap SSD
- ▶ ThinkSystem U.2 Intel P4510 2.0TB Entry NVMe PCIe3.0 x4 Hot Swap SSD
- ▶ ThinkSystem U.2 Intel P4510 4.0TB Entry NVMe PCIe3.0 x4 Hot Swap SSD
- ▶ ThinkSystem U.2 Intel P4510 8.0TB Entry NVMe PCIe3.0 x4 Hot Swap SSD
- ▶ ThinkSystem 2.5" Intel S4610 3.84TB Mainstream SATA 6Gb Hot Swap SSD

New Flash Storage adapters - 3.16, "Flash Storage adapters" on page 88:

- ▶ ThinkSystem HHHL KCM51V 1.6TB Mainstream NVMe PCIe 3.0 x4 Flash Adapter
- ▶ ThinkSystem HHHL KCM51V 3.2TB Mainstream NVMe PCIe 3.0 x4 Flash Adapter
- ▶ ThinkSystem HHHL KCM51V 6.4TB Mainstream NVMe PCIe 3.0 x4 Flash Adapter

Updated the list of operating systems, Chapter 6, "Operating system installation" on page 183

Removed the following withdrawn options:

- ▶ ThinkSystem 2.5" 5100 1.92TB Mainstream SATA 6Gb Hot Swap SSD
- ▶ ThinkSystem 2.5" 5100 240GB Mainstream SATA 6Gb Hot Swap SSD
- ▶ ThinkSystem 2.5" 5100 480GB Mainstream SATA 6Gb Hot Swap SSD
- ▶ ThinkSystem 2.5" 5100 960GB Mainstream SATA 6Gb Hot Swap SSD
- ▶ ThinkSystem 2.5" Intel S4500 1.92TB Entry SATA 6Gb Hot Swap SSD
- ▶ ThinkSystem 2.5" Intel S4500 240GB Entry SATA 6Gb Hot Swap SSD
- ▶ ThinkSystem 2.5" Intel S4500 3.84TB Entry SATA 6Gb Hot Swap SSD
- ▶ ThinkSystem 2.5" Intel S4500 480GB Entry SATA 6Gb Hot Swap SSD
- ▶ ThinkSystem 2.5" Intel S4500 960GB Entry SATA 6Gb Hot Swap SSD
- ▶ ThinkSystem 2.5" Intel S4600 1.92TB Mainstream SATA 6Gb Hot Swap SSD
- ▶ ThinkSystem 2.5" Intel S4600 240GB Mainstream SATA 6Gb Hot Swap SSD
- ▶ ThinkSystem 2.5" Intel S4600 480GB Mainstream SATA 6Gb Hot Swap SSD
- ▶ ThinkSystem 2.5" Intel S4600 960GB Mainstream SATA 6Gb Hot Swap SSD
- ▶ ThinkSystem 2.5" PM1633a 3.84TB Capacity SAS 12Gb Hot Swap SSD
- ▶ ThinkSystem 2.5" PM1633a 7.68TB Capacity SAS 12Gb Hot Swap SSD
- ▶ ThinkSystem 2.5" PM1635a 3.2TB Mainstream SAS 12Gb Hot Swap SSD
- ▶ ThinkSystem 2.5" PM1635a 800GB Mainstream SAS 12Gb Hot Swap SSD

- ▶ ThinkSystem 2.5" PM1635a 1.6TB Mainstream SAS 12Gb Hot Swap SSD
- ▶ ThinkSystem 2.5" PM863a 240GB Entry SATA 6Gb Hot Swap SSD
- ▶ ThinkSystem 2.5" PM863a 480GB Entry SATA 6Gb Hot Swap SSD
- ▶ ThinkSystem HHHL Intel P4500 4.0TB Entry NVMe PCIe3.0 x4 Flash Adapter
- ▶ ThinkSystem HHHL Intel P4500 8.0TB Entry NVMe PCIe3.0 x4 Flash Adapter
- ▶ ThinkSystem U.2 Intel P4500 2.0TB Entry NVMe PCIe3.0 x4 Hot Swap SSD
- ▶ ThinkSystem U.2 PX04PMB 800GB Performance NVMe PCIe 3.0 x4 Hot Swap SSD
- ▶ ThinkSystem U.2 PX04PMB 960GB Mainstream NVMe PCIe 3.0 x4 Hot Swap SSD

23 October 2018

- ▶ New drive options - Table 3-23 on page 77
 - ThinkSystem U.2 KCM51V 800GB Mainstream NVMe PCIe 3.0 x4 Hot Swap SSD
 - ThinkSystem U.2 KCM51V 1.6TB Mainstream NVMe PCIe 3.0 x4 Hot Swap SSD
 - ThinkSystem U.2 KCM51V 3.2TB Mainstream NVMe PCIe 3.0 x4 Hot Swap SSD
 - ThinkSystem U.2 KCM51V 6.4TB Mainstream NVMe PCIe 3.0 x4 Hot Swap SSD

3 October 2018

- ▶ Removed the following internal drive options which are now withdrawn:
 - ThinkSystem 2.5" 5100 480GB Entry SATA 6Gb Hot Swap SSD, 4XB7A08502
 - ThinkSystem 2.5" 5100 960GB Entry SATA 6Gb Hot Swap SSD, 4XB7A08503
 - ThinkSystem 2.5" 5100 1.92TB Entry SATA 6Gb Hot Swap SSD, 4XB7A08504
 - ThinkSystem U.2 PX04PMB 1.6TB Performance NVMe PCIe 3.0 x4 Hot Swap SSD, 7XB7A05922
 - ThinkSystem U.2 PX04PMB 1.92TB Mainstream NVMe PCIe 3.0 x4 Hot Swap SSD, 7N47A00096

11 September 2018

New information

- ▶ New families of solid-state drives - 3.10, "Internal drive options" on page 74
 - ThinkSystem 2.5" 5200 Mainstream SATA 6Gb Hot Swap SSDs
 - ThinkSystem 2.5" Intel S4510 Entry SATA 6Gb Hot Swap SSDs
 - ThinkSystem 2.5" Intel S4610 Mainstream SATA 6Gb Hot Swap SSDs
 - ThinkSystem 2.5" PM1645 Mainstream SAS 12Gb Hot Swap SSDs
 - ThinkSystem 2.5" PM1643 Capacity SAS 12Gb Hot Swap SSDs
- ▶ New capacities for existing SSD families
 - ThinkSystem 2.5" PM883 7.68TB Entry SATA 6Gb Hot Swap SSD
 - ThinkSystem U.2 PM983 7.68TB Entry NVMe PCIe 3.0 x4 Hot Swap SSD

Changed information

- ▶ Updated existing network adapters with new product names - 3.13, "Network adapters" on page 83. See the following article listing the old and new names:

<https://lenovopress.com/lp0931-thinksystem-network-adapters-renamed>

Introduction

The ThinkSystem SR950 is Lenovo's flagship server suitable for mission-critical applications that need the most processing power possible in a single server. The powerful 4U ThinkSystem SR950 can expand from two to as many as eight Intel Xeon Scalable Family processors.

The ThinkSystem SR950 server delivers 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.

This server packs numerous fault-tolerant and high-availability features into a high-density, rack-optimized chassis where all major subsystem components are front and rear accessible. This design makes upgrading and servicing the system simple and easy.

With the server's flexibility in initial hardware configuration, depending on the desired number of processors, drive bays, PCIe slots, and upgradeability to an 8-socket configuration, it means that the system will yield the best return from your application investment.

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 8
- ▶ 1.4, "Storage versus in-memory data" on page 10
- ▶ 1.5, "Flash storage" on page 11
- ▶ 1.6, "Energy efficiency" on page 12
- ▶ 1.7, "Services offerings" on page 12
- ▶ 1.8, "About this book" on page 13

1.1 Target workloads

The SR950 server introduces a new level of fault tolerance with advanced RAS features that are implemented in hardware and software to support the most demanding mission-critical applications that require 24x7 operations.

This server provides users who are looking for the highest level of scalable performance, 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 Lenovo's 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 the SR950 system can be highly used, which yields the best return for the following database applications:

- ▶ SAP Business Suite
- ▶ Microsoft SQL Data Warehouse
- ▶ SAP HANA
- ▶ IBM Db2

The SR950 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 SR950.

The OLTP workload optimization goal for the SR950 system 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 10.

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.

The SR950 system 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 10).

The ThinkSystem SR950 technology helps to 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 ThinkSystem SR950 expands the virtualization footprint to the workhorse applications of enterprise IT, namely those performance-intensive databases.

The SR950 makes virtualization of mission-critical, performance-intensive workloads possible in the following ways:

- ▶ Support for an integrated hypervisor: The server supports booting from an operating system or hypervisor installed on an M.2 solid-state drive.
- ▶ Processor support: The latest Gold or Platinum level Intel Xeon Processors (formerly codename “Skylake”) support Intel Virtualization Technology (Intel VT) Flex Priority and Intel VT Flex migration. These processes now also support memory bandwidth allocation.
- ▶ Large VM support for enterprise applications:
 - Virtualized SAP HANA
 - VMware vCloud Suite
 - Microsoft Hyper-V Private Cloud

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.

The SR950 server provides low latency, extreme performance, and efficient transaction management to accommodate mixed workload requirements. The SR950 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

Lenovo ThinkSystem SR950 is designed for your most demanding, mission-critical workloads, such as in-memory databases, large transactional databases, batch and real-time analytics, ERP, CRM, and virtualized server workloads. The powerful 4U ThinkSystem SR950 can grow from two to eight Intel Xeon Processor Scalable Family CPUs, achieving up to 45% faster performance than the previous generation. The modular design of SR950 speeds upgrades and servicing with easy front and rear access to all major subsystems, to maximize server uptime.




Figure 1-1 Lenovo ThinkSystem SR950

1.2.1 Scalability and performance

The SR950 offers numerous features to boost performance, improve scalability, and reduce costs:

- ▶ Dense server offering where up to 8 processors and 96 DIMMs are housed in only 4U of rack space.
- ▶ Offers a simple upgrade path that enables the server to be initially configured with 2 processors and then later be upgraded to 8 processors simply by adding a compute tray and system boards.
- ▶ Supports the Gold and Platinum level processors in the Intel Xeon Processor Scalable Family, including the “top bin” highest-performing 205 W processors available from Intel.
- ▶ Up to eight processors, each with up to 28 cores and 56 threads (for a total of up to 448 threads) to maximize the concurrent execution of multithreaded applications.
- ▶ Intelligent and adaptive system performance with Intel Turbo Boost Technology 2.0 allows processor cores to run at maximum speeds during peak workloads by temporarily exceeding processor TDP.
- ▶ Intel Hyper-Threading Technology boosts performance for multithreaded applications by enabling simultaneous multithreading within each processor core, up to two threads per core.
- ▶ Intel Virtualization Technology integrates hardware-level virtualization hooks that allow operating system vendors to better use the hardware for virtualization workloads.
- ▶ Support for up to 96 TruDDR4™ DIMMs operating at 2666 MHz means you have the fastest available memory subsystem and memory capacity of up to 12 TB using 96x 128 GB 3DS RDIMMs.
- ▶ Up to 24x 2.5-inch HDDs or SSDs provide a flexible and scalable all-in-one platform to meet your increasing demands.
- ▶ Of the 24x drive bays, support for up to 12 NVMe PCIe SSDs maximizes drive I/O performance, in terms of throughput, bandwidth, IOPS, and latency.
- ▶ The use of solid-state drives (SSDs) instead of, or along with, mechanical hard disk drives (HDDs), can improve I/O performance. An SSD can support up to 100 times more I/O operations per second (IOPS) than a typical HDD.
- ▶ New high-speed RAID controllers provide 12 Gb SAS connectivity to the drive backplanes. A variety of RAID adapters are available, with cache up to 4 TB and support for 12 drives on a single controller.
- ▶ Supports a new Lenovo-patented M.2 adapter design for convenient operating system boot functions. Available M.2 adapters support either one M.2 drive, or two M.2 drives in a RAID 1 configuration, for the boot drive performance and reliability.
- ▶ The server has a dedicated flexible LAN-on-motherboard (LOM) slot offering a variety of 10 GbE or Gigabit Ethernet adapters that do not occupy one of the standard PCIe slots.
- ▶ A total of up to 17 rear-accessible PCIe slots: 14 general-purpose PCIe 3.0 slots plus slots dedicated to the LOM adapter and two ML2 adapters. In addition, the server has two internal slots for RAID or HBA adapters.
- ▶ The server offers PCI Express 3.0 I/O expansion capabilities that improve the theoretical maximum bandwidth by almost 100% (8 GT/s per link using 128b/130b encoding) compared to the previous generation of PCI Express 2.0 (5 GT/s per link using 8b/10b encoding).

1.2.2 Availability and serviceability

The SR950 is based on a modular service model where all components and options can be removed from the front or rear of the system, even parts that are located in the center of the machine such as fans, memory DIMMs, and processors.

Watch the following video to show how easy it is to service the SR950 server:

<https://lenovopress.com/lp0899-thinksystem-sr950-ease-of-service-video>

The SR950 provides many features to simplify serviceability and increase system uptime:

- ▶ The server offers Single Device Data Correction (SDDC, also known as Chipkill), Adaptive Double- Device Data Correction (ADDDC, also known as Redundant Bit Steering or RBS), memory mirroring, and memory rank sparing for redundancy in the event of a non-correctable memory failure.
- ▶ The server offers hot-swap drives, supporting RAID redundancy for data protection and greater system uptime.
- ▶ The Dual M.2 Boot Adapter uses RAID-1 so that the two M.2 drives installed are configured as a redundant pair.
- ▶ The server has up to four hot-swap N+N redundant power supplies to help keep the server operational in the event of a power supply or utility supply failure
- ▶ Each compute tray has 6 hot-swap N+1 redundant fans (a total of 12 fans when two compute trays are installed) to maximize availability for business-critical applications.
- ▶ The power source-independent light path diagnostics functionality provides individual light path LEDs that lead the technician to failed (or failing) components, which simplifies servicing, speeds up problem resolution, and helps increase system availability.
- ▶ The included LCD system information display panel provides greater diagnostic function by listing all error messages and VPD data needed for a service call, thereby speeding problem resolution and increasing system uptime.
- ▶ Proactive Platform Alerts (including PFA and SMART alerts): Processors, voltage regulators, memory, internal storage (SAS/SATA HDDs and SSDs, NVMe SSDs, M.2 storage, and flash storage adapters), fans, power supplies, RAID controllers, and server ambient and subcomponent temperatures. Alerts can be surfaced through the XClarity Controller to managers such as Lenovo XClarity™ Administrator, VMware vCenter, and Microsoft System Center. These proactive alerts let you take appropriate actions in advance of possible failure, thereby increasing server uptime and application availability.
- ▶ Solid-state drives (SSDs) offer more reliability than mechanical HDDs for greater uptime.
- ▶ The built-in XClarity Controller continuously monitors system parameters, triggers alerts, and performs recovery actions in case of failures to minimize downtime.
- ▶ Built-in diagnostics in UEFI, using Lenovo XClarity Provisioning Manager, speed up troubleshooting tasks to reduce service time.
- ▶ Lenovo XClarity Provisioning Manager supports diagnostic functions and collects service data to USB key drive or remote CIFS share folder for troubleshooting, to reduce servicing time.
- ▶ Autorestart in the event of a momentary loss of AC power (based on power policy setting in the XClarity Controller service processor)
- ▶ Support for the XClarity Administrator Mobile app running on a supported smartphone and connected to the server through the service-enabled USB port, enables additional local systems management functions.

- ▶ Three-year or one-year customer-replaceable unit and onsite limited warranty, 9 x 5 next business day. Optional service upgrades are available.

1.2.3 Manageability and security

Powerful systems management features simplify local and remote management of the SR950:

- ▶ The server includes an XClarity Controller (XCC) to monitor server availability and perform remote management. XCC Enterprise is included as standard, which enables remote KVM and remote media files (ISO and IMG image files), boot capture, and power capping.
- ▶ Lenovo XClarity Administrator offers comprehensive hardware management tools that help to increase uptime, reduce costs and improve productivity through advanced server management capabilities.
- ▶ New UEFI-based Lenovo XClarity Provisioning Manager, accessible from F1 during boot, provides system inventory information, graphical UEFI Setup, platform update function, RAID Setup wizard, operating system installation function, and diagnostic functions.
- ▶ Support for Lenovo XClarity Energy Manager which captures real-time power and temperature data from the server and provides automated controls to lower energy costs.
- ▶ Integrated Trusted Platform Module (TPM) 2.0 support enables advanced cryptographic functionality, such as digital signatures and remote attestation.
- ▶ Supports Secure Boot to ensure that only a digitally signed operating system can be used. Supported with HDDs and SSDs as well as M.2 drives in the M.2 Adapter.
- ▶ Industry-standard Advanced Encryption Standard (AES) NI support for faster, stronger encryption.
- ▶ Intel Execute Disable Bit functionality can prevent certain classes of malicious buffer overflow attacks when combined with a supported operating system.
- ▶ Intel Trusted Execution Technology provides enhanced security through hardware-based resistance to malicious software attacks, allowing an application to run in its own isolated space, protected from all other software running on a system.

1.2.4 Energy efficiency

The SR950 offers the following energy-efficiency features to save energy, reduce operational costs, and increase energy availability:

- ▶ The server can be used in environments up to ASHRAE A4 with ambient temperatures up to 45 °C.
- ▶ The server is Energy Star 2.1 compliant
- ▶ Energy-efficient planar components help lower operational costs.
- ▶ High-efficiency power supplies with 80 PLUS Platinum certifications
- ▶ Intel Intelligent Power Capability powers individual processor elements on and off as needed, to reduce energy draw.
- ▶ Low-voltage 1.2 V DDR4 memory offers energy savings compared to 1.35 V and 1.5 V DDR3 DIMMs.
- ▶ Solid-state drives (SSDs) consume as much as 80% less power than mechanical 2.5-inch HDDs.
- ▶ The server uses hexagonal ventilation holes, which can be grouped more densely than round holes, providing more efficient airflow through the system.

- Optional Lenovo XClarity Energy Manager provide advanced data center power notification, analysis, and policy-based management to help achieve lower heat output and reduced cooling needs.

1.3 Positioning

The Lenovo ThinkSystem SR950 server is Lenovo's flagship server. The SR950 includes various new features when compared to the previous server generation of X-Architecture. This includes support for the new Intel Xeon Scalable Family processors. These processors support up to 28 cores, core speeds up to 3.6 GHz, and TDP ratings up to 205W.

When compared to the previous generations of Enterprise servers; eX5™ and X6, the SR950 is a powerful, compact 4U design that can be upgraded from a 2-socket, to a 4-socket or to an 8-socket system without having to replace the server enclosure or upgrade to a physically larger design. The SR950 has a choice of multiple types of processing, I/O, storage and memory configurations to suit all types of application requirements. This server fills the demand for enterprise workloads with more performance than ever before, high availability, and advanced RAS features.

Table 1-1 and Table 1-2 on page 9 shows a high-level comparison between the Broadwell processor-based x3850 X6 and x3950 X6 and the ThinkSystem SR950.

Table 1-1 Comparison between the x3850 X6 and SR950

	x3850 X6	SR950 (4-socket)	Benefits
Processor	<ul style="list-style-type: none"> ► Intel Xeon E7-4800/8800 v4 ► Up to 165W ► Up to 24 cores 	<ul style="list-style-type: none"> ► Intel Xeon Gold (5100, 6100) and Platinum (8100 series) ► Up to 205W ► Up to 28 cores 	<ul style="list-style-type: none"> ► Up to 25-45% higher performance from previous generation
Memory	<ul style="list-style-type: none"> ► 96 DIMM slots ► 6 TB max memory ► DDR4 8GB/16GB/32GB RDIMM, 64GB LRDIMM 	<ul style="list-style-type: none"> ► 48 DIMM slots ► 6 TB max memory ► DDR4 - 16GB/32GB RDIMM, 64GB LRDIMM, 128 GB 3DS RDIMM 	<ul style="list-style-type: none"> ► Allows for a more dense 8S upgrade
Internal storage	<ul style="list-style-type: none"> ► 2.5" HDD – 8 SAS/SATA - HDD/SSD ► 1.8" SSD – 16 SSDs ► 2.5" NVMe – 4 SSDs ► Add-In PCI Flash ► Front access RAID 	<ul style="list-style-type: none"> ► 2.5" HDD – 24 SAS/SATA - HDD/SSD ► 2.5" NVMe – 12 SSDs ► M.2 drive for booting ► Add-In PCI Flash ► Front access RAID 	<ul style="list-style-type: none"> ► Increased SAS/SATA and NVMe high speed storage ► M.2 provides a low cost boot option
Ethernet	<ul style="list-style-type: none"> ► 1GbE Management Port ► x8 Mezz LOM Dedicated Socket, choice of ML2 cards 	<ul style="list-style-type: none"> ► 1GbE Management Port ► x16 Mezz LOM Dedicated Socket, choice of ML2 cards ► Intel LOM adapters 	<ul style="list-style-type: none"> ► Higher speed ML2 adapter options ► Intel LOM provides a low cost 10GbE choice

	x3850 X6	SR950 (4-socket)	Benefits
PCIe	<ul style="list-style-type: none"> ▶ 9 rear PCIe slots, plus 1 ML2 ▶ Primary I/O Book and choice of Half-Length IO and Full-Length IO Books ▶ Up to 6 x16 PCIe (with 2 full length IO Books) 	<ul style="list-style-type: none"> ▶ 13 rear PCIe slots, plus 1 ML2, plus LOM ▶ Single I/O Book with choice of PCIe risers, No Full-Length adapter or GPU support ▶ Up to 6 x16 PCIe 	<ul style="list-style-type: none"> ▶ More PCIe in a dense 4U design ▶ Increased x16 PCI adapter support
PSU	<ul style="list-style-type: none"> ▶ 750W, 900W, 1400W 	<ul style="list-style-type: none"> ▶ 1100W, 1600W 	<ul style="list-style-type: none"> ▶ Common Lenovo PSU
System management	<ul style="list-style-type: none"> ▶ XClarity, System Director 	<ul style="list-style-type: none"> ▶ XClarity 	<ul style="list-style-type: none"> ▶ Simple and improved systems management
Other	<ul style="list-style-type: none"> ▶ 4U Modular rack design ▶ Two chassis for 4S x3850 and 8S x3950 ▶ LCD Panel ▶ TPM 	<ul style="list-style-type: none"> ▶ 4U Modular rack design ▶ Single 4U chassis ▶ LCD Panel and mobile mgmt app ▶ TPM or TCM 	<ul style="list-style-type: none"> ▶ Single 4U system supporting 2S – 8S

Table 1-2 Comparison between the x3950 X6 and SR950

	x3950 X6	SR950 (8-socket)	Benefits
Processor	<ul style="list-style-type: none"> ▶ Intel Xeon E7-4800/8800 v4 ▶ Up to 165W ▶ Up to 24 cores 	<ul style="list-style-type: none"> ▶ Intel Xeon Gold (5100, 6100) and Platinum (8100 series) ▶ Up to 205W ▶ Up to 28 cores 	<ul style="list-style-type: none"> ▶ Up to 25-45% higher performance from previous generation
Memory	<ul style="list-style-type: none"> ▶ 192 DIMM slots ▶ 12 TB max memory ▶ DDR4 - 8GB/16GB/32GB RDIMM, 64GB LRDIMM 	<ul style="list-style-type: none"> ▶ 96 DIMM slots ▶ 12 TB max memory ▶ DDR4 - 16GB/32GB RDIMM, 64GB LRDIMM, 128GB 3DS RDIMM 	<ul style="list-style-type: none"> ▶ Allows for a more dense 8S upgrade
Internal storage	<ul style="list-style-type: none"> ▶ 2.5" HDD – 16 SAS/SATA - HDD/SSD ▶ 1.8" SSD – 32 SSDs ▶ 2.5" NVMe – 8 SSDs ▶ Add-In PCI Flash ▶ Front access RAID 	<ul style="list-style-type: none"> ▶ 2.5" HDD – 24 SAS/SATA - HDD/SSD ▶ 2.5" NVMe – 12 SSDs ▶ M.2 drive for booting ▶ Add-In PCI Flash ▶ Front access RAID 	<ul style="list-style-type: none"> ▶ Increased SAS/SATA and NVMe high speed storage ▶ M.2 provides a low cost boot option
Ethernet	<ul style="list-style-type: none"> ▶ Two 1GbE Management Port ▶ Two x8 Mezz LOM Dedicated Socket, choice of ML2 cards 	<ul style="list-style-type: none"> ▶ 1GbE Management Port ▶ Two x16 Mezz LOM Dedicated Socket, choice of ML2 ▶ Intel LOM adapters 	<ul style="list-style-type: none"> ▶ Higher speed ML2 adapter options ▶ Intel LOM provides a low cost 10GbE choice

	x3950 X6	SR950 (8-socket)	Benefits
PCIe	<ul style="list-style-type: none"> ▶ 18 rear PCIe slots, plus 2 ML2 ▶ Primary IO Book and choice of Half-Length IO and Full-Length IO Books ▶ Up to 10 x16 PCIe (with 4 full length IO Books) 	<ul style="list-style-type: none"> ▶ 14 rear PCIe slots, plus 2 ML2, plus LOM ▶ Single IO Book with choice of PCIe risers, No Full-Length adapter or GPU support ▶ Up to 11 x16 PCIe 	<ul style="list-style-type: none"> ▶ More PCIe in a dense 4U design ▶ Increased x16 PCI adapter support
PSU	<ul style="list-style-type: none"> ▶ 750W, 900W, 1400W 	<ul style="list-style-type: none"> ▶ 1100W, 1600W 	<ul style="list-style-type: none"> ▶ Common Lenovo PSU
System management	<ul style="list-style-type: none"> ▶ XClarity, System Director 	<ul style="list-style-type: none"> ▶ XClarity 	<ul style="list-style-type: none"> ▶ Simple and improved systems management
Other	<ul style="list-style-type: none"> ▶ 4U Modular rack design ▶ Two chassis for 4S x3850 and 8S x3950 ▶ LCD Panel ▶ TPM 	<ul style="list-style-type: none"> ▶ 4U Modular rack design ▶ Single 4U chassis ▶ LCD Panel and mobile mgmt app ▶ TPM or TCM 	<ul style="list-style-type: none"> ▶ Single 4U system supporting 2S – 8S

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 the ThinkSystem SR950 server combines extreme IOPS performance and low response time for transactional database workloads. The flash technologies that are used in the SR950 server includes 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. The SR950 supports up to 12 PCIe NVMe drives.

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 2.5.2, “PCIe NVMe drives” on page 41

- ▶ Flash Storage Adapters

These storage devices offer high-performance NAND technology in a standard PCIe adapter form factor. Up to 3.84 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 2.5.4, “Flash Storage Adapters” on page 43

- ▶ SSDs with SAS/SATA interfaces

The SR950 server supports 2.5-inch SSDs with 6 GB or 12 Gb SAS/SATA host connections. These drives include the 7.68 TB capacity SAS SSD. These drives provide high performance and capacity. Supported RAID controllers offer various RAID and performance options.

For more information, see 2.5.1, “SAS/SATA SSD technology” on page 40

These technologies allow the ThinkSystem SR950 server to deliver break-through performance for targeted workloads by offering lower latency and higher performance compared to traditional HDDs.

1.6 Energy efficiency

The ThinkSystem SR950 server offers 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 1100 W AC and 1600 W AC power supplies have 80 PLUS Platinum certification.
- ▶ Intel Intelligent Power Capability powers individual processor elements on and off as needed to reduce power draw.
- ▶ Low-voltage 1.2 V TruDDR4 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 Lenovo solutions and technologies and data center assessments. The Professional Services team offers the following benefits in support of Lenovo 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, the Professional Services team has the following offerings specifically for the SR950:

- ▶ 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
- ▶ Water cooled technology assessment, planning and installation

For more information about these services, see this website:

<http://shop.lenovo.com/us/en/systems/services/>
http://download.lenovo.com/servers_pdf/dcs_offerings_v3.0.0.pdf

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 the SR950 system.

For more information, contact the appropriate team that is listed in Table 1-3.

Table 1-3 Professional Services and Data Center Services contact details

Contact	Geography
Professional Services team	
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
Data Center Services team	
power@lenovo.com	Worldwide

1.8 About this book

In this book, readers get a general understanding of ThinkSystem SR950 technology 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 SR950 hardware and architecture.
- ▶ The remainder of the book describes preparing and implementing the SR950 server. We describe power and cooling considerations, operating system installations, systems management, and firmware update tools.

SR950 Technology

In this chapter, we give an overview of the technologies that are included in the Lenovo ThinkSystem SR950 server. We describe system architecture, chassis design with modular structure, and Intel Xeon Platinum (8100) and Gold (6100 and 5100) Scalable Family processors.

Then we describe the current memory and internal drive options and features of the server subsystem, including innovative memory-channel storage technology and PCIe NVMe solid-state drives (SSDs). We also describe other advanced technology in the server, including server configuration options.

This chapter includes the following topics:

- ▶ 2.1, “System design” on page 16
- ▶ 2.2, “System architecture” on page 26
- ▶ 2.3, “Processors” on page 28
- ▶ 2.4, “Memory” on page 36
- ▶ 2.5, “Internal storage” on page 40
- ▶ 2.6, “Unified Extensible Firmware Interface (UEFI)” on page 43
- ▶ 2.7, “xClarity Controller (XCC)” on page 44
- ▶ 2.8, “Scalability” on page 45

2.1 System design

The ThinkSystem SR950 is Lenovo's flagship enterprise rack server. This server is scalable up to 8-sockets using the latest Platinum (8100) and Gold (6100 and 5100) members of the Intel Xeon Scalable Family processors.

Figure 2-1 shows the ThinkSystem SR950 server.



Figure 2-1 Lenovo ThinkSystem SR950

There are three types of trays that make up the SR950:

- ▶ **Compute tray:** The processors and memory are located on the compute trays, the Upper Compute Tray and the Lower Compute Tray. Each compute tray holds up to four processors and 24 DIMMs on two system boards. The compute trays are accessible from the front of the server once the front bezel is removed.
- ▶ **Storage tray:** For some four-socket storage-rich configurations, a storage tray can be used in the upper tray area instead of a compute tray to add additional storage to the chassis. The storage tray is similarly accessible from the front of the server.
- ▶ **I/O tray:** The I/O tray is accessible from the rear of the server and houses all the PCIe slots. The I/O tray can be configured with multiple types of riser cards to suit the type of I/O cards being installed in the server.

These trays allow the server to be configured in multiple ways to achieve the desired level of processing power, I/O, storage and memory capacities.

The following subsections discuss each of these trays in detail.

2.1.1 Compute trays

The ThinkSystem SR950 can scale from 2-sockets to 4-sockets, all the way up to 8-sockets without the need to replace the server's enclosure or upgrade to a physically larger design. This is achieved through adding compute trays and system boards to the 4U chassis.

The server supports up to two Compute Trays. The compute trays are accessible from the front of the server. Figure 2-2 on page 17 shows the top compute tray being removed from the chassis.




Figure 2-2 ThinkSystem SR950 compute tray being removed (upper)

Figure 2-3 shows the lower compute tray being removed from the chassis.




Figure 2-3 ThinkSystem SR950 compute tray being removed (lower)

Figure 2-4 shows a top view of the lower and upper compute trays. The processors and memory are located towards the rear of the compute tray. The front of the compute tray holds the storage and long sliding tabs for pulling the out hotswap fans.




Figure 2-4 Upper (left) and lower (right) compute trays

Each compute tray contains:

- ▶ One or two compute system boards, each comprising:
 - Two processors
 - 24 DIMMs
- ▶ Six hot-swap fans, accessible from the front of the server even when the compute tray is installed
- ▶ 12x 2.5-inch hot-swap drive bays
- ▶ One dedicated PCIe 3.0 x8 slot reserved for a internal RAID adapter or SAS HBA for internal SAS/SATA drives

Figure 2-5 shows the rear of the compute tray with two compute system boards, where one of the system boards is being removed.




Figure 2-5 Compute tray with two system boards, one being removed

The compute system boards on its own is shown in Figure 2-6 when not installed in a Compute tray. Each system board holds two processors and 24 DIMM sockets, 12 per processor. The system board also has connectors for the NVMe ports used to connect to the PCIe NVMe drive bays at the front of the server.




Figure 2-6 Compute system board

The dedicated PCIe 3.0 x8 slot used with the RAID adapter or SAS HBA that connects to the drive bays in the compute tray is located between the fan chutes and the drive bays, as shown in Figure 2-7 on page 19.




Figure 2-7 Internal RAID adapter slot (lower compute tray)

2.1.2 Storage tray

The SR950 chassis can support up to 24 drives internally when both compute trays are installed (12 drives in each tray). However, for some four-socket storage-rich configurations with only one compute tray installed, a storage tray can be used in the upper tray area instead of an additional compute tray to provide the storage. The storage tray has no memory or processors on it.

Only one storage tray can be installed and it must be installed in the upper tray area. The lower tray area requires a compute tray to be installed for the processor(s) and memory.

A storage tray contains:

- ▶ NVMe ports (for connecting NVMe drives in the tray to the processors in the lower compute tray)
- ▶ Six hot-swap fans, accessible from the front of the server even when the storage tray is installed
- ▶ 12x 2.5-inch hot-swap drive bays (6 NVMe)
- ▶ One PCIe slot reserved for a RAID adapter for internal SAS/SATA drives (located under the drives)

Figure 2-8 on page 20 shows the storage tray pulled out of the system. The rear of the tray lacks the system board that contains the CPU and memory DIMMs




Figure 2-8 Storage tray

2.1.3 I/O tray

At the rear of the server is the I/O tray. Figure 2-9 displays the locations of the slots at the rear of the server that make up the I/O tray.




Figure 2-9 PCIe 3.0 slots at the rear of the SR950 server

The I/O tray can contain the following:

- ▶ 17 PCIe 3.0 slots depending on the riser cards installed in the tray. These are all accessible from the rear of the server
- ▶ One internal dedicated M.2 slot for an M.2 adapter located internally on the I/O tray

Figure 2-10 on page 21 shows the rear of the server where the I/O tray is being removed.



Figure 2-10 I/O tray being removed from the rear of the server

The I/O tray without riser cards and with riser cards installed is shown in Figure 2-11.




Figure 2-11 I/O tray without riser cards installed (left) and with riser cards installed (right)

PCIe Slots 1-4 and 10-15 are set on riser cards. There are multiple types of riser cards that can be installed in the riser slots (1 and 2) in the tray. This allows for flexible I/O configurations in the system. Figure 2-12 on page 22 shows the riser card options for riser card slots 1 and 2 in the I/O tray.




Figure 2-12 Riser card options for PCIe slots 1-4 and 10-15

Slots 16 and 17 are also set on a single riser card as seen in Figure 2-13.




Figure 2-13 Riser card option for PCIe slots 16 and 17

PCIe slots 5-9 are located on the I/O tray planar as seen in Figure 2-11 on page 21 which displays the I/O tray.

For additional information on the type of slots and the adapters that are supported in the system, refer to 3.12, “I/O expansion options” on page 78.

For recommended configurations of the SR950 system and its trays, refer to 3.2, “Server configurations” on page 52.

Additionally the 4U chassis supports up to four power supplies and up to 12 redundant hot-swap single-rotor fans, 6 in each compute tray.

2.1.4 Power supplies

The server supports up to four hot-swap power supplies (PSU) with capacities of 1100W and 1600W. Power supplies are N+N redundant, which means that if the server has four power supplies installed, then the server can continue operation even with two failed power supplies,

or if the utility service that feeds those two power supplies goes offline. The power supplies used in the SR950 are Lenovo Common Form Factor (CFF) power supplies.

Figure 2-14 shows the location of the SR950 power supplies.




Figure 2-14 Power supply unit for the SR950

For additional information on the power supplies refer to 3.19, “Power supplies” on page 89.

2.1.5 Fans

The SR950 uses a set of 60mm hot-swap single-rotor fans for system-wide cooling. Each compute tray or storage tray houses 6 hot-swap fans. The 6 fans in the compute tray are N+1 redundant, meaning that a single fan failure is tolerated in each compute tray.

Fans are located internal to the server immediately in front of the compute system boards, but are accessible from the front of the server by the long handle attached to the fan. As the handle is pulled out, the attached fan rotates from vertical orientation to horizontal orientation, allowing the fan to be removed from the system using the handle without removing the tray or powering off the system. Figure 2-15 shows the locations of the fans with some fans partially removed to show the mechanism.



Figure 2-15 Hot-swap fans in the SR950, front bezel removed

Figure 2-16 shows the lower system tray with the 6 fans some partially and fully removed.




Figure 2-16 Six fans installed in each compute tray (lower compute tray shown)

For additional information on the fans and cooling the system, refer to 3.18, “Cooling” on page 88.

2.1.6 Mid-Chassis interconnects

In the middle of the server there are mid-chassis interconnects. These interconnects provide the communication path between the compute and storage trays and the I/O tray and power supplies. The interconnects in the middle of the server can be seen in Figure 2-17.




Figure 2-17 Top view of the mid-chassis interconnects in the chassis

The interconnects behind the I/O tray can be seen in Figure 2-18 on page 25.




Figure 2-18 Interconnects at the back of the I/O tray

Figure 2-19 shows a front (left) and rear (right) view of the internal interconnects of the chassis when all trays are removed.




Figure 2-19 Front (left) and rear (right) view of the chassis showing internal chassis interconnects

2.2 System architecture

This section describes the architecture of the ThinkSystem SR950. Figure 2-20 shows the server at a high level. The system boards with processors and memory are installed in compute trays along with the drive bays and storage adapter. The I/O board houses all onboard PCIe slots as well as the slots for the riser cards that offer the remaining PCIe slots.

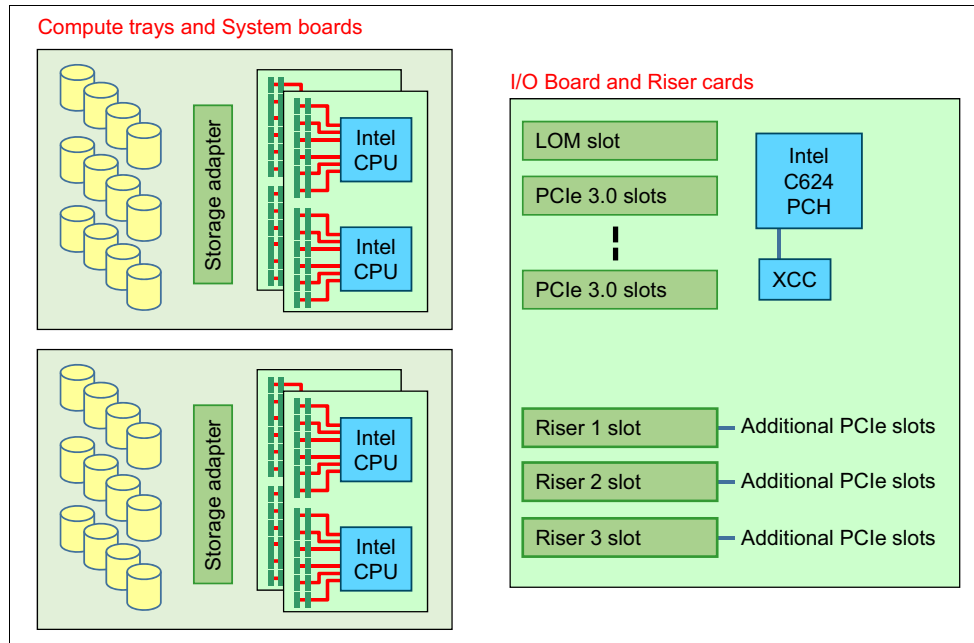


Figure 2-20 SR950 architecture

The specific connections between system boards, processors and PCIe slots are dependent on the server configuration. This is described in detail in 3.2, “Server configurations” on page 52. Specific processor-slot connections based on each configuration are described in 3.12.1, “Slot availability by server configuration” on page 81.

Figure 2-21 shows System board 1 where processors 1 and 2 are installed. The slots on the I/O board are connected directly to these processor (or indirectly through the Intel C624 Platform Controller Hub, PCH).

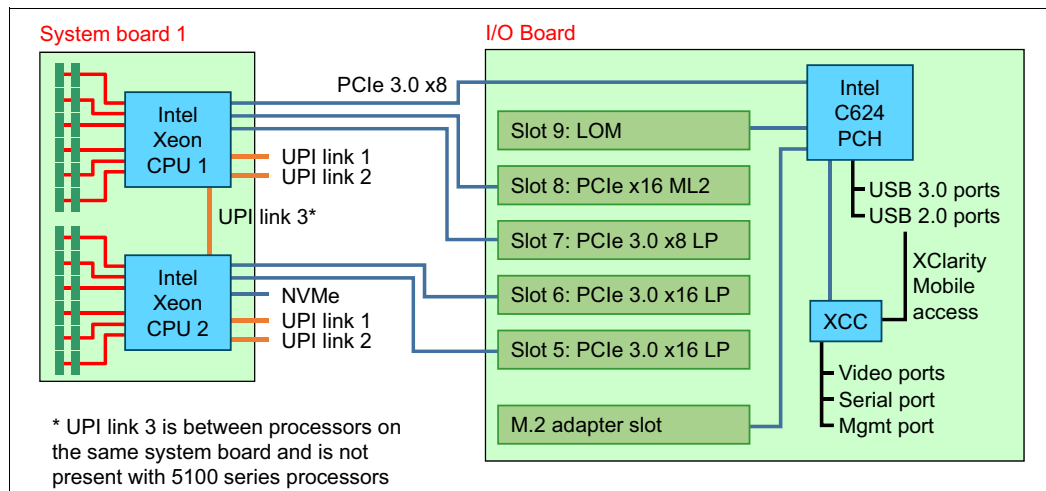


Figure 2-21 System board 1 connections to the I/O board

The processors within the chassis communicate via UltraPath Interconnect (UPI) ports. These ports allow high speed communication between the CPUs. UPI replaces the Intel QuickPath Interconnect (QPI). One UPI link connects the two processors on each system board (UPI 3) and the other two UPI links are used to connect the other processors in adjacent boards to form four-way and larger systems.

In four-processor configurations, the design of the UPI links is either in a mesh or a ring topology. When a processor has 3 UPI ports, it can be used in a 4-socket mesh configuration where all processors are directly connected to every other processor. As discussed in 2.3, “Processors” on page 28, the 5100 Series processors only have 2 UPI ports (ports 1 and 2) and therefore do not support mesh topology. These two UPI configurations are shown in Figure 2-22.

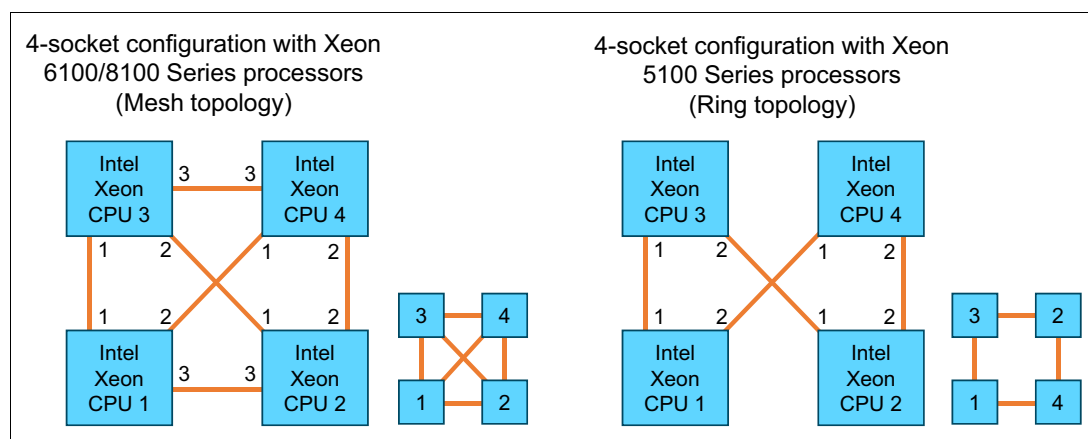


Figure 2-22 Mesh (left) and Ring (right) topologies for four-socket configurations

Ring topology with 8100 Series processors: As we discuss in 3.2, “Server configurations” on page 52, the “Four-Socket Upgradable / Storage Rich” configuration is a four-socket configuration that uses a Ring topology but uses 8100 Series processors.

The server supports an 8-socket configuration (8100 Series processors required). Figure 2-23 shows the UPI links used to connect eight processors.

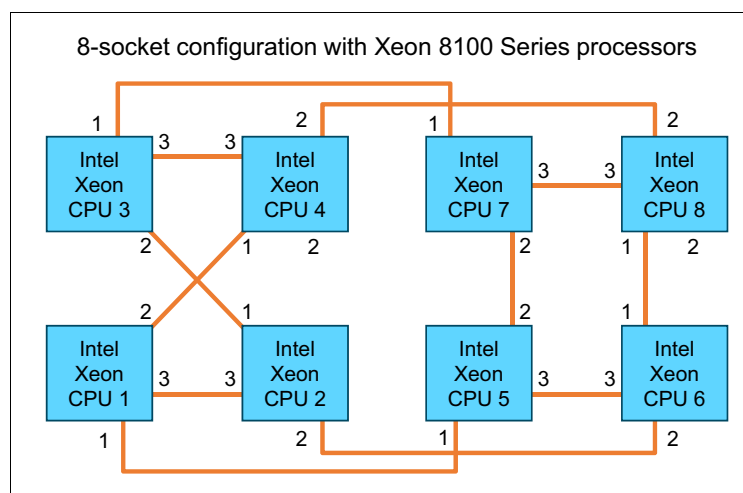


Figure 2-23 8-socket topology

Refer to Table 3-8 on page 62 for information on the number UPI links each processor option has for the SR950 system.

The next section discusses the technology of the processors used in the SR950 in more detail.

2.3 Processors

The new Intel Xeon Processor Scalable Family (formerly codename “Skylake-SP”, succeeding “Broadwell”) processors have been grouped into four functional levels or *shelves*:

- ▶ Bronze
- ▶ Silver
- ▶ Gold
- ▶ Platinum

The SR950 supports the Gold (5100 and 6100) and Platinum (8100) level processors, this includes the “top bin” highest performing 205 W processors. These processors feature the new Intel microarchitecture and provide a higher core count, higher frequencies, faster UPI links (the follow-on to QPI), more UPI links, faster DDR4 bus speeds, faster AVX-512 (instruction set architecture) and advanced RAS features.

Physically, the new Intel Skylake processor is larger in size than the previous generations of Intel processors. Figure 2-24 shows a size comparison of Skylake and Broadwell processors.




Figure 2-24 Broadwell (left) and Skylake (right) processors

2.3.1 Gold and Platinum product family

The following groups of the Intel Xeon Scalable Family that are used in the SR950 server include:

- ▶ The Intel Xeon Gold product family. This includes the 5100 and 6100 Series processors. This group supports four-processor configurations only. 5100 Series processors have two UPI links and 6100 Series processors have three UPI links.
- ▶ The Intel Xeon Platinum product family are 8100 Series processors. This group supports up to eight-socket configurations and they each have three UPI links.

The Gold and Platinum processors offers the following key features:

- ▶ Up to 2666 MHz memory bus speed
- ▶ 14nm process technology
- ▶ Up to 28 cores per processor
- ▶ UPI links at speeds up to 11.2 GT/s
- ▶ Intel Hyper-Threading Technology (2 threads per core)
- ▶ Intel AVX-512 (AVX2 in previous generations)
- ▶ DDR4 continued memory interface support, which brings greater performance and power efficiency
- ▶ Rebalanced Cache Hierarchy: Increased MLC 1.375 MB Last Level Cache per core
- ▶ Integrated Fabric: Intel Omni-Path Architecture
- ▶ Integrated PCIe 3.0 controller with 48 lanes per processor at 8 GT/s
- ▶ 1.5x memory bandwidth increase (6 channels vs. 4 in previous generation)
- ▶ Memory Technology: 6xDDR4 channels 2133, 2400, 2666 MT/s RDIMM, LRDIMM, Apache Pass
- ▶ Intel Virtualization Technology (VT-x and VT-d)
- ▶ Intel Turbo Boost Technology
- ▶ Improved performance for integer and floating point operations
- ▶ Virtualization improvements with regards to posted interrupts, page modification logging, and VM enter/exit latency reduction
- ▶ Intel AES-New instructions for accelerating of encryption
- ▶ Advanced UPI and memory reliability, availability, and serviceability (RAS) features
- ▶ Machine Check Architecture (MCA) recovery (non-running and running paths)
- ▶ Resource director technology (RDT): Cache monitoring technology, cache allocation technology, memory bandwidth monitoring

Security technologies include:

- ▶ Intel BIOS Guard 2.0
- ▶ Intel Boot Guard
- ▶ Secure Key
- ▶ Intel Trusted Execution Technology (TXT)
- ▶ Intel QuickAssist Technology (QAT)
- ▶ Intel Platform Trust Technology (PTT)
- ▶ Intel Memory Protection Extensions (MPX)
- ▶ Mode based execution control (XU/XS bits)

Table 2-1 compares the previous generation of the Intel Xeon processors to the new Scalable Family processors that are supported in SR950 systems.

Table 2-1 Processor comparisons

Feature	X6 family, Xeon E7 v3	X6 family, Xeon E7 v4	Intel Xeon Scalable Family
Processor family	Intel Xeon E7-8800 v3 Intel Xeon E7-4800 v3	Intel Xeon E7-8800 v4 Intel Xeon E7-4800 v4	Intel Xeon 8100 Intel Xeon 6100 Intel Xeon 5100
Processor codenames	Haswell EX	Broadwell EX	Skylake SP
Cores per CPU	Up to 18 cores	Up to 24 cores	Up to 28 cores
Last level cache	Up to 45 MB L3 cache	Up to 60 MB L3 cache	Up to 1.375 MB per core Up to 38.5 MB per processor
QPI / UPI Data rate	QPI: 3 links 9.6 GT/s max	QPI: 3 links 9.6 GT/s max	UPI: 2 or 3 links 10.4 GT/s max
CPU TDP rating	Up to 165 W	Up to 165 W	Up to 205 W
DIMM sockets	24 DDR3 DIMMs per CPU 24 DDR4 DIMMs per CPU	24 DDR4 DIMMs per CPU	12 DDR4 DIMMs per CPU
Maximum memory speeds	2133 MHz	2400 MHz	2666 MHz
PCIe technology	PCIe 3.0 (8 GTps)	PCIe 3.0 (8 GTps)	PCIe 3.0 (8 GTps)

2.3.2 Gold and Platinum features

Intel Xeon Scalable Family processors include a broad set of features and extensions. Many of these technologies are common with previous Intel Xeon generations, and some technologies are brand new.

The feature comparison between the 8100, 6100 and 5100 Series processors is highlighted in Table 2-2.

Table 2-2 Feature comparisons of the Platinum and Gold processors

Feature	5100 (Gold)	6100 (Gold)	8100 (Platinum)
Number of UPI ports to sockets	<ul style="list-style-type: none"> 2 socket - 2 UPI 4 socket - 2 UPI 	<ul style="list-style-type: none"> 2 socket - 2 UPI 2 socket - 3 UPI 4 socket - 2 UPI 4 socket - 3 UPI 	<ul style="list-style-type: none"> 2 socket - 2 UPI 2 socket - 3 UPI 4 socket - 2 UPI 4 socket - 3 UPI 8 socket - 3 UPI
Memory channel and speeds	<ul style="list-style-type: none"> 6-channel DDR4 2400 MHz 	<ul style="list-style-type: none"> 6-channel DDR4 2666 MHz 	<ul style="list-style-type: none"> 6-channel DDR4 2666 MHz
Number of UPI links and speed	<ul style="list-style-type: none"> 2 UPI links 10.4 GT/s 	<ul style="list-style-type: none"> 3 UPI links 10.4 GT/s 	<ul style="list-style-type: none"> 3 UPI links 10.4 GT/s
Turbo supported	<ul style="list-style-type: none"> Intel Turbo Boost 2.0 	<ul style="list-style-type: none"> Intel Turbo Boost 2.0 	<ul style="list-style-type: none"> Intel Turbo Boost 2.0
Hyper-Thread supported	<ul style="list-style-type: none"> Intel Hyper-Threading 	<ul style="list-style-type: none"> Intel Hyper-Threading 	<ul style="list-style-type: none"> Intel Hyper-Threading
Advanced Vector supported	<ul style="list-style-type: none"> Intel AVX-512 1x 512-bit FMA^a 	<ul style="list-style-type: none"> Intel AVX-512 2x 512-bit FMAs 	<ul style="list-style-type: none"> Intel AVX-512 2x 512-bit FMAs

Feature	5100 (Gold)	6100 (Gold)	8100 (Platinum)
Memory capacity	► 768GB	► 768GB ► 1.5TB ^b	► 768GB ► 1.5TB ^b
Number of PCIe lanes	► 48 lanes PCIe Gen3	► 48 lanes PCIe Gen3	► 48 lanes PCIe Gen3
RAS supported	► Advanced RAS	► Advanced RAS	► Advanced RAS

a. Note: Gold Processor 5122 will support 2666 DDR4 and 2 512-bit FMA units

b. 1.5TB support requires a processor with an M suffix, capable of supporting the larger RAM size. Refer to Table 3-8 on page 62 for additional information on supported processor models

Notable features of the Intel Xeon Scalable Family processors are described in the following subsections.

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) available in the Intel Xeon Scalable Family, 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 web page:

<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. The new Scalable Family of processors has had VT-x latency reductions and increased memory bandwidth allocation.

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 web page:

<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 web page:

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

Turbo Boost Technology 2.0

As with the previous generation of Intel Processors, the Intel Xeon Scalable 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 SR950 system. 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 web page:

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

UltraPath Interconnect (UPI)

The Intel Xeon Scalable family processor uses the new UltraPath Interconnects (UPI) to provide communication between processor to processor. UPI succeeds QuickPath Interconnect (QPI) technology which was used in previous generations of processors.

UPI brings new, faster coherent links (10.4 GT/s versus 9.6 GT/s with QPI) and greater message efficiency, increased bandwidth and performance over QPI, improved messaging efficiency allowing multiple requests per packet and improved power efficiency allowing

multiple requests per packet and improved power efficiency through lowering the unused UPI links (L0p) to a low-power state.

Each processor has either two or three UPI links depending on the model of the processor creating a high-speed, low latency path to the other processors. Figure 2-22 on page 27 and Figure 2-23 on page 27 show the UPI interconnections between the processors.

Table 3-8 on page 62 lists the number of UPI links that each processor model offers.

For more information about UPI, see this web page:

<http://software.intel.com/en-us/articles/intel-xeon-processor-scalable-family-technical-overview>

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 PCIe 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 web page:

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

RAS features

The Intel Xeon Scalable Family provides several new features as well as enhancements of some existing features associated with the RAS (Reliability, Availability, and Serviceability). Two levels of support are provided with the Scalable Family processors, Standard RAS and Advanced RAS, however all Gold and Platinum processors used in the SR950 have Advanced RAS features.

The following is a list of the Advanced RAS features:

- ▶ Local Machine Check Exception based recovery
Local Machine Check Exception (LMCE) improves the Machine Check Architecture (MCA) recovery-execution path event and increases the possibility of recovery. It increases the possibility of recovery by letting the thread that is consuming the uncorrected data handle the error without broadcasting the fatal MCERR events to the rest of threads in the system.
- ▶ SDDC +1, Adaptive DDC (MR) +1
SDDC can correct any single memory chip failure as well as multi-bit errors from any portion of a single memory chip. It can reconstruct memory contents on the fly, even in the event of the complete failure of one chip. This increases the longevity of the chip.
DDC enables a memory DIMM to continue operation even in the event of two sequential DRAM device corrections. The Enhanced Adaptive DDC (DDC+1) adds the capability to detect and correct up to two DRAM device corrections on top of DDC. DDC+1 is a new feature and provides additional uptime.
- ▶ Addressed range / Partial Memory Mirroring
Legacy memory mirroring is transparent to the OS; however, address range mirroring requires a firmware-OS interface for a user to specify the desired subset of memory to mirror. The OS system-managed memory mirroring allows selective ranges to mirror, increasing the integrity of the data at an efficient cost.

- ▶ Intel UPI Protocol Protection via 32 bit Rolling CRC check retry
Detection of transient data errors over Intel UPI interconnects, via 32 bit rolling CRC, and the ability to retry the failed transactions. This adds a level of protection reduces circuit-level errors, detects data errors across the system and limits the impact of errors.
- ▶ PCI Express ECRC
PCI Express end-to-end CRC checking (cyclic redundancy check), provides error-detecting code for data entering the system via PCI devices increasing system reliability.
- ▶ Advanced Error Detection and Correction
Advanced Error Detection and Correction (AEDC) checks to identify and correct errors the processor may encounter within its internal workings. AEDC will attempt to correct the fault by retrying the instruction. If the event is successful after retry it is considered a corrected event. If the event fails after retry, a fatal MCERR is logged and signaled. ADEC is limited to the integer unit in the processor making it self-contained within the processor itself.
- ▶ Reset error harvesting
Async Warm Reset based off a global reset (was previously DWR). This allows the system to reset but retain error logs in the machine check banks due to the warm reset; as opposed to a cold reset where error logs may not be retained.

Machine Check Architecture recovery

The Intel Xeon processor Scalable family also includes Enhanced Machine Check Architecture (eMCA) 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.

The new MCA recovery features of the Intel Xeon processor Scalable 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 Machine Check Architecture (eMCA)
Compared to the Machine Check Architecture (MCA) of previous generations, the Enhanced MCA (eMCA) 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.

Security improvements

The Intel Xeon processor Scalable family features the following important security improvements that help to protect systems from different types of security threats:

- ▶ Intel QuickAssist Technology (QAT)
Intel QuickAssist Technology (Intel QAT) accelerates and compresses cryptographic workloads by offloading the data to hardware capable of optimizing those functions. This makes it easier for developers to integrate built-in cryptographic accelerators into network, storage, and security applications. In the case of the Intel Xeon processor Scalable family, Intel QAT is integrated into the Intel C620 Platform Controller Hub (PCH) chip and offers

outstanding capabilities including 100 Gbs Crypto, 100Gbs Compression, 100kops RSA, and 2k Decrypt. This can benefit the solution in the following ways:

- Server: secure browsing, email, search, big-data analytics (Hadoop), secure multi-tenancy, IPsec, SSL/TLS, OpenSSL
- Networking: firewall, IDS/IPS, VPN, secure routing, Web proxy, WAN optimization, 3G/4G authentication
- Storage: real-time data compression, static data compression, secure storage.

► Intel Trusted Execution Technology (Intel TXT) with PPI-x

Intel TXT, Intel VT-x, and Intel VT-d is hardware-based techniques, with which you can isolate VMs and start VMs in a trusted environment only. In addition, malware-infected VMs cannot affect another VMs on the same host. New to the Scalable Family processors is Intel TXT with Physical Presence Interface-Extended (PPI-x).

► Intel Boot Guard

Boot Guard adds another level of protection to the server by performing a cryptographic Root of Trust for Measurement (RTM) of the early firmware platform storage device such as the trusted platform module or Intel Platform Trust Technology (Intel PTT). It can also cryptographically verify early firmware using OEM-provided policies. Unlike Intel TXT, Boot Guard doesn't have any software requirements; it is enabled at the factory, and it cannot be disabled. Boot Guard operates independently of Intel TXT but it is also compatible with it. Boot Guard reduces the chance of malware exploiting the hardware or software components.

► Intel BIOS Guard 2.0

BIOS Guard is an augmentation of existing chipset-based BIOS flash protection capabilities. The Intel Xeon processor Scalable family adds the fault tolerant boot block update capability. The BIOS flash is segregated into a protected and unprotected regions. These processors bypasses the top-swap feature and flash range register locks/protections, for explicitly enabled signed scripts, to facilitate the fault-tolerant boot block update. This feature protects the BIOS flash from modification without the platform manufacturer's authorization, as well as during BIOS updates. It can also help defend the platform from low-level DOS attacks.

► Intel Memory Protection Extensions (Intel MPX)

C/C++ pointer arithmetic is a convenient language construct often used to step through an array of data structures. If an iterative write operation does not take into consideration the bounds of the destination, adjacent memory locations may get corrupted. Such unintended modification of adjacent data is referred as a buffer overflow. Buffer overflows have been known to be exploited, causing denial-of-service (DoS) attacks and system crashes. Similarly, uncontrolled reads could reveal cryptographic keys and passwords. More sinister attacks, which do not immediately draw the attention of the user or system administrator, alter the code execution path such as modifying the return address in the stack frame to execute malicious code or script.

Intel's Execute Disable Bit and similar hardware features have blocked buffer overflow attacks that redirected the execution to malicious code stored as data. Intel MPX technology consists of new Intel architecture instructions and registers that compilers can use to check the bounds of a pointer at runtime before it is used.

► Mode Based Execution (MBE) Control (XU/XS bits)

MBE provides finer grain control on execute permissions to help protect the integrity of the system code from malicious changes. It provides additional refinement within the Extended Page Tables (EPT) by turning the Execute Enable (X) permission bit into two options:

- XU for user pages
- XS for supervisor pages

The CPU selects one or the other based on permission of the guest page and maintains an invariant for every page that does not allow it to be writable and supervisor-executable at the same time. A benefit of this feature is that a hypervisor can more reliably verify and enforce the integrity of kernel-level code. The value of the XU/XS bits is delivered through the hypervisor, this means that hypervisor support is necessary.

2.4 Memory

The SR950 uses Lenovo TruDDR4 memory operating at up to 2666 MHz. The server supports 12 DIMMs per processor, which corresponds to 48 DIMMs with four processors installed and 96 DIMMs when eight processors are installed. Each processor has six memory channels with two DIMMs per channel.

DDR4 is a memory standard that is supported by the Scalable family of processors. DDR4 memory modules can run at greater speeds than DDR3 DIMMs, operate at lower voltage, and are more energy-efficient than DDR3 modules.

Lenovo TruDDR4 memory modules are tested and tuned by Lenovo to maximize performance and reliability. They have ECC protection and support Chipkill technologies.

The SR950 supports three different types of DDR4 memory:

- ▶ Registered DIMMs (RDIMMs)
- ▶ Load Reduced DIMMs (LRDIMMs)
- ▶ Three-Dimensional Stacked RDIMMs (3DS RDIMMs)

Many end users configure server memory based on a capacity requirement. Historically, RDIMMs have been favored for lower capacity requirements and LRDIMMs have been selected for higher capacity requirements. A relatively new technology is three-dimension stacking (3DS) which stacks DRAM components using Through Silicon Vias (TSV) and allows for greater capacity DIMMs without increasing the load on the memory bus. The 128GB RDIMM listed above is a 3DS RDIMM.

All processors supported by the SR950 support all three memory types, however the use of more than six 128 GB 3DS RDIMMs per processor (that is, more than 768 GB per processor) requires a processor model with an “M” suffix.

Each processor has two Integrated Memory Controllers (IMC) and six DIMMs are connected to each IMC as shown in Figure 2-25 on page 37. Although the system will function with only one memory DIMM installed, it is recommended to install six memory DIMMs per processor to allow for three memory DIMMs per IMC (one DIMM per channel). This will ensure the best system performance.




Figure 2-25 Integrated Memory Controllers with attached DIMMs

For a discussion about memory performance, see the Lenovo Press paper, *Intel Xeon Scalable Family Balanced Memory Configurations*, available from:

<https://lenovopress.com/lp0742>

All DIMMs operate at a speed of 2666 MHz, both at 1 DIMM per channel and 2 DIMMs per channel. However, if the processor selected has a lower memory bus speed (e.g. 2400 MHz), then all DIMMs will operate at that lower speed. Memory will also operate at a lower speed based on power efficiency modes set in UEFI, as shown in Table 2-3.

Table 2-3 Memory speed as a function of Processor and server operating mode

Processor	Maximum Performance Mode	Efficiency Favors Performance Mode	Efficiency Favors Power Mode	Minimal Power Mode
Xeon 8100	2667 MHz	2667 MHz	2400 MHz	1867 MHz
Xeon 6100	2667 MHz	2667 MHz	2400 MHz	1867 MHz
Xeon 5100	2400 MHz	2400 MHz	2133 MHz	1867 MHz

2.4.1 Memory RAS modes

The following memory RAS modes are supported by the SR950 system:

- ▶ Independent channel mode
- ▶ Mirrored channel mode
- ▶ Rank sparing mode

Independent channel mode

In independent memory mode, memory channels can be populated with DIMMs in any order and you can populate all channels for each processor in any order with no matching requirements. However, specific DIMM population sequence will provide the best performance when operating in the supported memory modes.

Independent memory mode provides the highest and most efficient level of memory performance, but doesn't offer memory mirroring or rank sparing like the other modes offer.

When installing DIMMs in the server for Independent mode, all channels must run at the same interface frequency. Individual channels may run at different DIMM timings (RAS latency, CAS latency, and so forth).

Mirrored channel mode

Memory-mirroring mode provides full memory redundancy however this also reduces the usable system memory capacity in half. Memory channels are grouped in pairs with each channel receiving the same data. If a failure occurs, the memory controller seamlessly switches from the DIMMs on the primary channel to the DIMMs on the backup channel.

Mirrored Channel Mode must be populated identically with regards to size and organization. DIMM slot populations within a channel do not have to be identical, but the same DIMM slot location across Channel 0, Channel 1, and Channel 2 or between Channel 3, Channel 4 or Channel 5 must be populated identically. For example, you must install two identical 16 GB 2666 MHz DIMMs equally and symmetrically across the two memory channels to achieve 16 GB of mirrored memory.

Rank sparing mode

In memory-sparing mode, one memory rank serves as a spare for other ranks on the same channel in case they fail.

The spare rank is held in reserve and not used as active memory until a failure is indicated, with reserved capacity subtracted from the total available memory in the system. After an error threshold is surpassed in a system protected by memory sparing, the content of a failing rank of DIMMs is copied to the spare rank. The failing rank is then taken offline and the spare rank placed online for use as active memory in place of the failed rank.

Since the failover process involves copying of memory content, the level of memory redundancy provided by memory sparing is less than that provided by memory mirroring: memory mirroring is the preferred failure-protection choice for critical applications.

The spare rank(s) must have identical or larger memory capacity than any of the other ranks (sparing source ranks) on the same channel. After sparing, the sparing source rank will no longer be available as system memory. The system will lose 25% memory capacity.

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 rules on channel DIMM population and channel DIMM matching vary by the operational mode used and must be met at all times. For more information on the memory options and memory rules and population requirements for each operational mode, refer to 3.7, “Memory options” on page 64.

2.4.2 Memory protection features

The SR950 system supports the following memory protection technologies:

- ▶ Error-correcting code (ECC)
- ▶ Single Device Data Correction (SDDC), also known as Chipkill
- ▶ Adaptive Double-Device Data Correction (ADDDC), also known as Redundant Bit Steering or RBS

These are described in the following subsections.

Error-correcting code (ECC)

Error-correcting code (ECC) memory can detect and correct the most common kinds of internal data corruption and is used in scenarios where data corruption can not be tolerated. All SR950 DIMMs have ECC capability for data protection.

Typically, ECC memory maintains a memory system immune to single-bit errors: the data that is read from each word is always the same as the data that had been written to it.

Single Device Data Correction (Chipkill)

Single Device Data Correction (SDDC), also known as Chipkill, is an advanced form of ECC available on the SR950 server. SDDC protects the memory in the system from any single memory chip failure. It also protects against multi-bit errors from any portion of a single memory chip.

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

SDDC is used in all memory modes, however it is only supported with x4-based memory DIMMs; look for “x4” in the DIMM description.

Adaptive double-device data correction (RBS)

ADDDC (also known as Redundant Bit Steering or RBS) is deployed at runtime to dynamically map out a failing DRAM device and continue to provide SDDC coverage on the DIMM, translating to longer DIMM longevity. The operation occurs at the fine granularity of bank and/or rank to have minimal impact on the overall system performance.

With ADDDC, the memory subsystem is always configured to operate in performance mode. When the number of corrections on a DRAM device reaches the targeted threshold value, with help from the UEFI runtime code, the identified failing DRAM region is adaptively placed in lockstep mode where the identified failing region of the DRAM device is mapped out of ECC. Once in ADDDC, cache line ECC continues to cover single DRAM (x4) error detection and apply a correction algorithm to the nibble.

Each DDR4 channel supports one to two regions that can manage one or two faulty DRAMs, and, depending on the Xeon Scalable Family processor model, support is at the Bank level or at full Rank granularity:

- ▶ Adaptive Data Correction (ADC): This is offered on Bronze and Silver processors in the Intel Xeon Scalable Family. Adaptive data correction can only occur at the Bank granularity.
- ▶ Adaptive double-device data correction (ADDDC): This is offered on Gold and Platinum processors. ADDDC can occur at the Bank and Rank granularity, with additional hardware facilities for device map-out.

The SR950 server uses Gold and Platinum processors and so supports ADDDC at rank granularity.

The dynamic nature of the operation makes the performance implications of the lockstep operation on the system to be material only after the DRAM device is detected to be failing. The overall lockstep impact on system performance is now a function of the number of bad DRAM devices on the channel, with the worst-case scenario of two bad Ranks on every DDR4 channel.

2.5 Internal storage

In this section, we give an overview of the following technologies that are used in the SR950 system:

- ▶ 2.5.1, “SAS/SATA SSD technology” on page 40
- ▶ 2.5.2, “PCIe NVMe drives” on page 41
- ▶ 2.5.3, “M.2 drives” on page 42
- ▶ 2.5.4, “Flash Storage Adapters” on page 43

2.5.1 SAS/SATA SSD technology

The SR950 offers up to 24x 2.5-inch drive bays, depending on the server configuration selected as discussed in 3.2, “Server configurations” on page 52.

SSD solutions can provide the following benefits compared to traditional hard drives:

- ▶ 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 SR950 system delivers up to 240,000 IOPS per single drive. 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.

SAS and SATA are available in four different categories, depending on the endurance rating for the drives, expressed as the number of full-drive writes per day (DWPD) that the SSD supports:

- ▶ Enterprise Entry (<3 DWPD)
- ▶ Enterprise Mainstream (3-5 DWPD)
- ▶ Enterprise Performance (10+ DWPD)
- ▶ Enterprise Capacity (low DWPD but very high drive capacity)

Entry, Mainstream and Performance have relatively similar read and write IOPS performance, but the key difference between them is their endurance (or lifetime). SSDs have a finite number of program/erase (P/E) cycles so their expected usable life depends how long they can perform write operations. SSD write endurance is typically measured by the number of program/erase (P/E) cycles that the drive incurs over its lifetime, listed as the total bytes of written data (TBW) in the device specification. The value of DWPD is calculated based on the TBW value, the capacity of the drive, and an expected use of 5 years.

SAS and SATA drives are supported in either of the two available SR950 drive backplanes, shown in Figure 2-26 on page 41:

- ▶ SAS/SATA backplane: Supports four SAS or SATA SSD drives.
- ▶ AnyBay™ backplane: Supports four SAS/SATA drives, or two SAS/SATA drives and two NVMe drives.




Figure 2-26 SR950 drive backplanes

The SR950 supports up to six of these drive backplanes for a total of 24 drive bays. The advantage of the AnyBay backplane is that it allows you to start with SAS or SATA SSDs and install four of them per backplane, and then later as needed, you can replace two of the SAS/SATA drives with NVMe drives for higher performance.

For additional information on using SSDs in the SR950, refer to 3.8, “Internal storage” on page 66.

The supported drives are listed in 3.10, “Internal drive options” on page 74.

2.5.2 PCIe NVMe drives

With the use of Lenovo AnyBay backplanes, the SR950 can be configured to support up to 12 PCIe NVMe drives. Figure 2-27 on page 42 shows a photo of one of the supported NVMe PCIe SSD.



Figure 2-27 NVMe PCIe SSD

Non-Volatile Memory Express (NVMe) is new PCIe 3.0 high performance SSD technology that provides high I/O throughput and low latency. NVMe interfaces remove SAS/SATA bottlenecks and enables the full performance of contemporary NAND flash memory. Each NVMe PCI SSD has direct PCIe 3.0 x4 connection, which provides at least 2x more bandwidth and 2x less latency than SATA/SAS-based SSD solutions. NVMe drives are also optimized for heavy multi-threaded workloads by using internal parallelism and many other improvements, such as enlarged I/O queues.

NVMe SSDs are typically available in three different categories, depending on the endurance rating for the drives, expressed as the number of full-drive writes per day (DWPD) that the SSD supports:

- ▶ Enterprise Entry (<3 DWPD)
- ▶ Enterprise Mainstream (3-5 DWPD)
- ▶ Enterprise Performance (10+ DWPD)

The NVMe SSDs supported with the SR950 include a standard hot-swap drive tray, however these are PCIe devices and so only support informed hot removal and hot insertion, provided the operating system supports these functions.

NVMe SSDs are the same 2.5-inch form factor as SAS/SATA drives, however the use of NVMe drives requires the AnyBay backplane that includes PCIe connections. The AnyBay backplane, shown in Figure 2-26 on page 41, has a PCIe 3.0 x8 connector to support two PCIe NVMe drives each with a PCIe x4 host interface. These PCIe connections are routed to PCIe ports on the system boards of the SR950 and connect directly to the processors. The actual connections used depend on the hardware configuration of the server. See 3.2, “Server configurations” on page 52 for details.

2.5.3 M.2 drives

The server supports one or two M.2 form-factor drives for use as an operating system boot solution. With two M.2 drives configured, the drives are configured by default as a RAID-1 mirrored pair for redundancy.

The M.2 drives install into a Lenovo patented M.2 adapter which in turn is installed in a dedicated slot on the system board. Refer to 3.8.3, “M.2 drives” on page 70 for details of implementation.



Figure 2-28 Dual M.2 Boot Adapter and a 128 GB M.2 drive

2.5.4 Flash Storage Adapters

The SR950 also supports Flash Storage Adapters which are effectively an NVMe SSDs in the form factor of a PCIe adapter. We describe NVMe drives in 2.5.2, “PCIe NVMe drives” on page 41.

Similar to NVMe SSDs, the Flash Storage Adapters are available in different categories, depending on the endurance rating for the drives, expressed as the number of full-drive writes per day (DWPD) that the SSD supports. The SR950 supports:

- ▶ Enterprise Mainstream (3-5 DWPD)
- ▶ Enterprise Performance (10+ DWPD)

Up to 14 Flash Storage Adapters are supported in the SR950.

Figure 2-29 shows the NVMe Flash Storage Adapter.



Figure 2-29 ThinkSystem HHHH PX04PMC 3.84TB Mainstream NVMe PCIe 3.0 x4 Flash Adapter

2.6 Unified Extensible Firmware Interface (UEFI)

The SR950 uses the new Lenovo XClarity Provisioning Manager for UEFI-based system management, accessible from F1 during boot. 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.

UEFI includes the following features:

- ▶ System inventory information
- ▶ Graphical UEFI setup screens
- ▶ Platform updating capabilities
- ▶ RAID setup wizard
- ▶ Operating system installation
- ▶ Diagnostic functions
- ▶ Uses 64-bit code architecture, with no restrictions for code size.
- ▶ Offers a modular design; adapter vendors can add any functions in ROM.
- ▶ Integration with light path diagnostics for errors
- ▶ Event logs
- ▶ Remote out-of-band access
- ▶ In-band (OS) and out-of-band (remote) UEFI code updates

The UEFI functions are discussed further in Chapter 5, “Preparing the hardware” on page 119.

For more information about UEFI, see this website:

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

2.7 xClarity Controller (XCC)

The server includes the xClarity Controller management processor, the equivalent of IMM in System x servers. XCC helps you monitor server availability and perform remote management. It is based on the Pilot4 XE401 baseboard management controller (BMC) and uses a dual-core ARM Cortex A9 service processor.

XCC Enterprise is included as standard with the SR950, which enables remote KVM and remote media files (ISO and IMG image files), boot capture, and power capping.

The XCC 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)
- ▶ Serial port text console redirection
- ▶ System LED control (power, HDD, activity, alerts, and heartbeat)

XCC collects inventory on the following components:

- ▶ CPU
- ▶ DIMM
- ▶ Disk
- ▶ Power supply units
- ▶ Fan
- ▶ PCI cards
- ▶ System board
- ▶ Compute board 1 and 2
- ▶ System LOM
- ▶ System firmware

XCC provides the following features:

- ▶ An embedded web server, which gives you remote control from any standard web browser
- ▶ Shared or dedicated Ethernet port. Support for LAN over USB for in-band communications to the XCC.
- ▶ 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 that can be attached to e-mail alerts
- ▶ Boot Capture
- ▶ Remote mounting of ISO and IMG files
- ▶ Virtual console collaboration
- ▶ Power capping

XCC in the SR950 server supports remote management through the following interfaces:

- ▶ Intelligent Platform Management Interface (IPMI) 2.0
- ▶ Simple Network Management Protocol (SNMP) v3 (no SET commands; no SMNP v1)
- ▶ Common Information Model (CIM-XML)
- ▶ Representational State Transfer (REST) support
- ▶ Redfish support (DMTF compliant)
- ▶ Web browser - HTML 5-based browser interface (Java and ActiveX not required) using a responsive design (content optimized for device being used - laptop, tablet, phone) with NLS support

For additional information and using XCC, refer to Chapter 5, “Preparing the hardware” on page 119.

2.8 Scalability

The SR950 server has a flexible modular 4U design with which you can increase the server’s compute, memory and storage simply by adding either an additional compute tray or storage tray to the chassis.

Likewise, the SR950’s I/O can be upgraded and/or expanded simply by adding additional riser cards to the rear of the server to provide additional I/O slots. This design means upgrading the hardware without having to upgrade the chassis.

Figure 2-30 on page 46 shows an example of a tray being added to the chassis to increase CPU, memory and storage capacity of the system.



Figure 2-30 Top tray of the SR950 being added for additional CPU, memory and storage

The server supports five different configurations, depending on the desired number of processors, drive bays, PCIe slots, and upgradeability to an 8-socket configuration.

For additional information on these configurations and scalability of these configurations, refer to the 3.2, “Server configurations” on page 52

Product information

In this chapter, we describe the components of the Lenovo ThinkSystem SR950.

This chapter includes the following topics:

- ▶ 3.1, “Components and connectors” on page 48
- ▶ 3.2, “Server configurations” on page 52
- ▶ 3.3, “Server configuration upgrades” on page 56
- ▶ 3.4, “Standard specifications” on page 58
- ▶ 3.5, “Models” on page 60
- ▶ 3.6, “Processor options” on page 61
- ▶ 3.7, “Memory options” on page 64
- ▶ 3.8, “Internal storage” on page 66
- ▶ 3.9, “Controllers for internal storage” on page 72
- ▶ 3.10, “Internal drive options” on page 74
- ▶ 3.11, “Optical drives” on page 78
- ▶ 3.12, “I/O expansion options” on page 78
- ▶ 3.13, “Network adapters” on page 83
- ▶ 3.14, “Fibre Channel host bus adapters” on page 86
- ▶ 3.15, “SAS adapters for external storage” on page 86
- ▶ 3.16, “Flash Storage adapters” on page 88
- ▶ 3.17, “GPU adapters and coprocessors” on page 88
- ▶ 3.18, “Cooling” on page 88
- ▶ 3.19, “Power supplies” on page 89
- ▶ 3.20, “Integrated virtualization” on page 90
- ▶ 3.21, “Systems management” on page 90
- ▶ 3.22, “Security” on page 94
- ▶ 3.23, “Operating system support” on page 94

3.1 Components and connectors

Figure 3-1 shows the front of the SR950 server. The server supports up to 24x 2.5-inch hot-swap drive bays and all are accessible from the front of the server. 12 of these drive bays can be configured to support NVMe SSDs. Also in the front is the front operator panel complete with a pull-out LCD display for additional management functions. Behind the front bezel are the hot-swap fans.

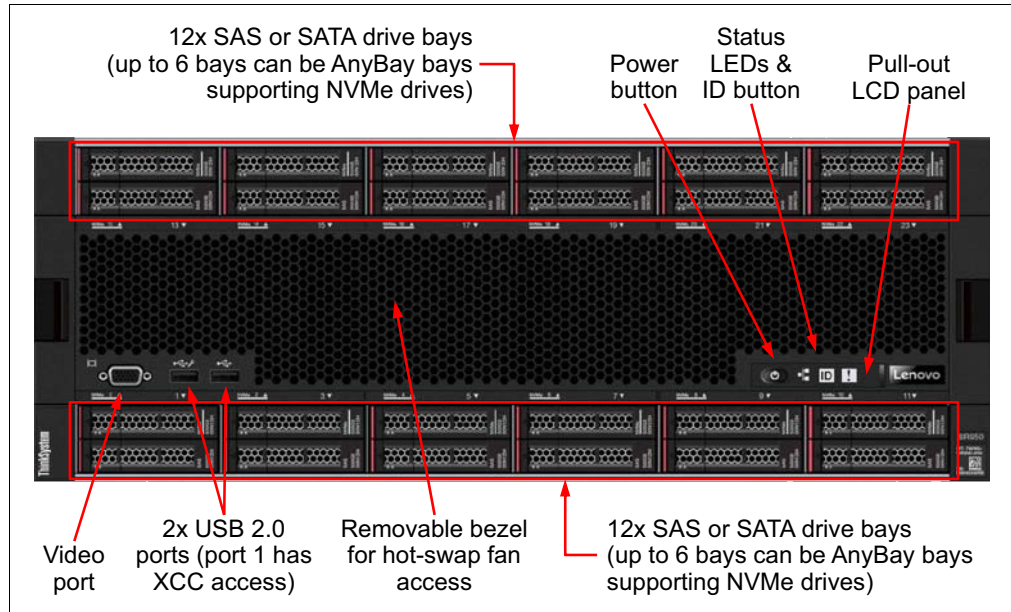


Figure 3-1 Front view of the Lenovo ThinkSystem SR950

Figure 3-2 on page 49 shows the rear of the SR950 server. The server has up to 17 PCIe slots, depending on the riser cards selected and the number of processors installed. In addition, the server has two internal slots for RAID adapters or HBAs to connect to the internal SAS/SATA drives.




Figure 3-2 Rear view of the Lenovo ThinkSystem SR950

The processors and other components are located on compute trays, the Upper Compute Tray and the Lower Compute Tray. Each compute tray holds up to four processors on two system boards. The compute trays are accessible from the front of the server, as shown in Figure 3-3.




Figure 3-3 Lower Compute Tray partially removed

The server has one or two compute trays. Each compute tray contains:

- ▶ One or two compute system boards, each comprising:
 - Two processors
 - 24 DIMMs
- ▶ Six hot-swap fans, accessible even when the compute tray is installed
- ▶ 12x 2.5-inch hot-swap drive bays
- ▶ One PCIe slot reserved for a RAID adapter for internal SAS/SATA drives

Figure 3-4 shows the rear of the compute tray with its two compute system boards, one of them removed.




Figure 3-4 Compute tray with two compute system boards

The compute system board is shown in Figure 3-5. Each system board holds two processors and 24 DIMM sockets, 12 per processor. The system board also has connections for the NVMe ports used to connect to the PCIe NVMe drive bays at the front of the server.




Figure 3-5 Compute system board

Processors installed in the compute system boards are numbered as shown in Figure 3-6 on page 51.




Figure 3-6 Processor numbering (viewed from the rear of the compute trays)

For some four-socket storage-rich configurations, a storage tray is used in the upper tray area instead of a compute tray. Figure 3-7 shows the key components of the storage tray.




Figure 3-7 Storage tray

All slots and ports at the rear of the server are housed in the I/O tray as shown in Figure 3-8 on page 52.




Figure 3-8 I/O tray

3.2 Server configurations

The SR950 supports five different configurations, depending on the desired number of processors, drive bays, PCIe slots, and upgradability to an 8-socket configuration.

Notes:

- ▶ The configurations listed here represent Lenovo's recommendations regarding the selection of compute system boards, I/O riser cards, and NVMe drive bays to maximize the available resources. You can select fewer drives and lower slot counts by using lower-slot-count riser cards to meet application requirements.
- ▶ The use of only two, three or six processors in configurations where supported will reduce the available I/O slots and NVMe drive bays. This is described in detail in 3.8, "Internal storage" on page 66 and 3.12, "I/O expansion options" on page 78.
- ▶ Configurations 1 and 2 support either two, three or four processors however the use 5100 Series processors requires that four processors be selected. Two or three 5100 Series processors are not supported in these configurations due to the way the processors are connected.
- ▶ Configurations 3 and 4 are designed to be upgradable to 8 processors. As a result, 8100 Series processors are required.

When building a server in the configurator, the feature codes in Table 3-1 are key to selecting the right configuration.

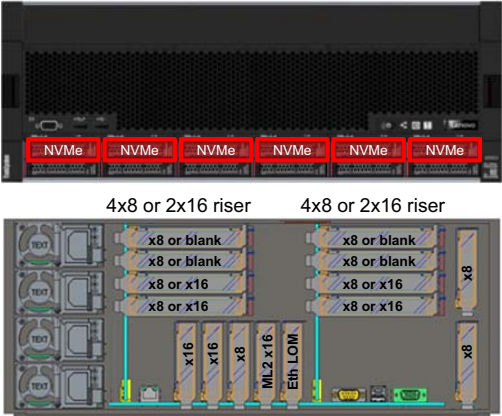
Table 3-1 Feature codes for selecting base server configurations

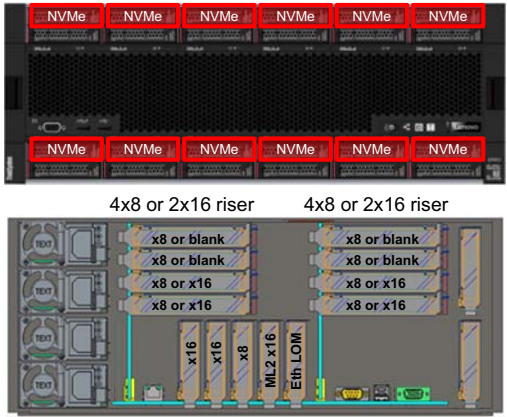
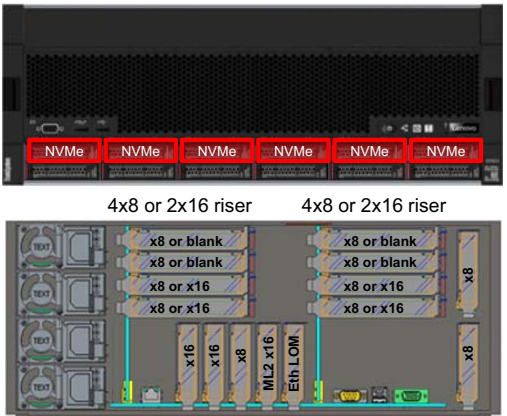
Part number	Feature code	Description
CTO only	AUNY	SR950 2S to 4S Base (Upgrade to 8S using the 4S Performance to 8S Upgrade Kit)
CTO only	AUKF	SR950 2S to 8S Base (Simple upgrade without needing the upgrade kit)

Part number	Feature code	Description
CTO only	AX3Y	12 Hard Drive Maximum System
CTO only	AX3Z	24 Hard Drive Maximum System

Upgrades: Configurations 1-4 can be upgraded to an 8-socket configuration (Configuration 5) as described in 3.3, “Server configuration upgrades” on page 56.

Table 3-2 Server configurations

 <p>The diagram illustrates the internal layout of the server. At the top is a 'Filler tray'. Below it is a 'Compute tray' containing six NVMe drives. Underneath the compute tray are two '4x8 or 2x16 riser' units. Below the risers are two rows of drive bays. The left row contains four bays, each labeled 'x8 or blank', 'x8 or blank', 'x8 or x16', and 'x8 or x16'. The right row contains four bays, each labeled 'x8 or blank', 'x8 or blank', 'x8 or x16', and 'x8 or x16'. To the right of these bays is a '2x8 riser' unit. At the bottom of the diagram are labels for 'ML2 x16', 'Eth LOM', and 'x8'.</p>	<p>Configuration 1. Four-Socket Performance</p> <p>The highest-performing 4-socket server, at the lowest cost when up to 12 drives is enough and without the need for a simple upgrade.</p> <ul style="list-style-type: none"> ► Supports 2, 3 or 4 processors (if using 5100 Series processors, 4 processors must be selected) ► 4 sockets configured in mesh topology for best performance ► 12 drive bays, 6 of which support NVMe drives (4 NVMe with 3 processors and 2 NVMe with 2 processors) ► 15 rear PCIe slots with 4 processors (10 rear slots with 3 processors, 6 rear slots with 2 processors) ► Filler installed in the upper tray area ► Can be upgraded to 8S but requires the 4S Performance to 8S Upgrade Kit, additional components, and Lenovo hardware installation. ► Select feature codes AUNY and AX3Y
---	--

 <p>Storage tray</p> <p>Compute tray</p> <p>4x8 or 2x16 riser 4x8 or 2x16 riser</p> <p>No riser</p>	<h3>Configuration 2. Four-Socket Performance/Storage Rich</h3> <p>The highest-performing 4-socket server, when needing more than 12 drives without the need for a simple upgrade.</p> <ul style="list-style-type: none"> ▶ Supports 2, 3 or 4 processors (if using 5100 Series processors, 4 processors must be selected) ▶ 4 processors configured in mesh topology for best performance ▶ 24 drives, 12 of which support NVMe drives (8 NVMe with 3 processors, 4 NVMe with 2 processors) ▶ 13 rear PCIe slots with 4 processors (9 rear slots with 3 processors, 5 rear slots with 2 processors) ▶ Storage tray installed in the upper tray area ▶ Can be upgraded to 8S but requires the 4S Performance to 8S Upgrade Kit, additional components, and Lenovo hardware installation. ▶ Select feature codes AUNY and AX3Z
 <p>Filler tray</p> <p>Compute tray</p> <p>4x8 or 2x16 riser 4x8 or 2x16 riser</p> <p>2x8 riser</p>	<h3>Configuration 3. Four-Socket Upgradable</h3> <p>The lowest-cost 4-socket server capable of a simple upgrade to 8 socket and when 12 drive bays with 4 processors is enough storage.</p> <ul style="list-style-type: none"> ▶ 2 or 4 processors, must be 8100 Series processors ▶ 4 processors in a ring topology ▶ 12 drive bays including up to 6 NVMe (2 NVMe with 2 processors) ▶ Up to 15 rear PCIe slots with 4 processors (6 rear slots with 2 processors) ▶ Filler installed in the upper tray area ▶ Upgradable to 8-socket with an additional compute tray and two system boards ▶ Once upgraded, the system will have 24 drive bays ▶ Select feature codes AUKF and AX3Y

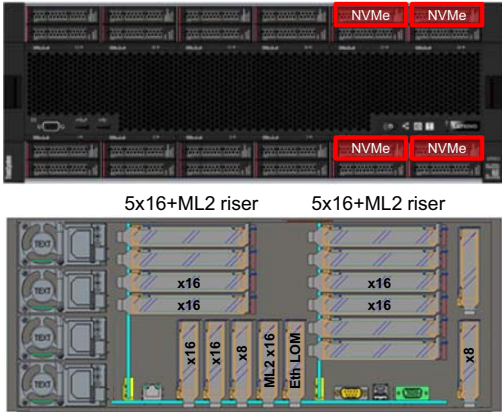
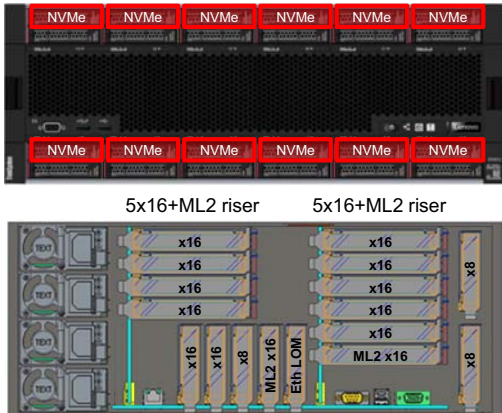
	<p>Configuration 4. Four-Socket Upgradable / Storage Rich</p> <p>4-socket server capable of a simple upgrade to 8-socket and needing more than 12 drive bays with 4 processors.</p> <ul style="list-style-type: none"> ► Requires 4 processors, must be 8100 Series ► 4 processors in a ring topology ► 24 drive bays (including up to 4 NVMe) ► 10 rear PCIe slots ► Two compute trays, with one system board in each tray ► Upgradable to 8-socket with two additional system boards ► Once upgraded, the server will support 12 NVMe drives. ► Select feature codes AUKF and AX3Z
	<p>Configuration 5. Eight-Socket Storage Rich</p> <p>Full featured 6- or 8-socket server with up to 24 drive bays and the maximum number of PCIe x16 slots.</p> <ul style="list-style-type: none"> ► Requires 6 or 8 processors, must be 8100 Series ► 24 drive bays, 12 of which support NVMe drives (8 NVMe with 6 processors) ► 17 rear PCIe slots (13 rear slots with 6 processors installed) ► Two compute trays each with two system boards ► Select feature codes AUKF and AX3Z

Table 3-3 summarizes the configurations.

Table 3-3 SR950 server configurations

Configuration	Supported processors	Lower Tray	Upper Tray	Processor Topology	Maximum storage	Maximum PCIe slots	Upgrade to 8S?
4S Performance	2 or 4 Any processor ^a	Compute	Filler	4S Mesh ^b	12 drives (6 NVMe)	15	Yes ^c
4S Performance/ Storage Rich	4 Any processor	Compute	Storage	4S Mesh ^b	24 drives (12 NVMe)	13	Yes ^c
4S Upgradable	2 or 4 8100 Series only	Compute	Filler	4S Ring	12 drives (6 NVMe)	15	Yes ^d
4S Upgradable/ Storage Rich	4 8100 Series only	Compute 1 board	Compute 1 board	4S Ring	24 drives (4 NVMe)	10	Yes ^d
8S Storage Rich	8 8100 Series only	Compute	Compute	8S	24 drives (12 NVMe)	17	Not applic.

- a. If using 5100 Series processors, then 4 processors must be selected. Using two 5100 Series processors is not supported.
- b. Mesh topology requires a processor with 3 UPI connections; Processors such as the 5100 series with only 2 UPI connections will operate in Ring topology. See the table in 3.6, "Processor options" on page 61.
- c. Upgrade requires the ThinkSystem SR950 4S Performance to 8S Upgrade Kit, 4TA7A10969. See 3.3, "Server configuration upgrades" on page 56 below.
- d. After upgrade, maximum storage will be 24 drives (12 NVMe) and up to 17 total PCIe slots available (may require replacement riser cards)

For details about PCIe slot availability based on each server configuration, see 3.12.1, "Slot availability by server configuration" on page 81.

3.3 Server configuration upgrades

This section describes the system upgrades you can perform in the field. The upgrades are based on the configurations that are described in 3.2, "Server configurations" on page 52.

Key components in the upgrades are listed in Table 3-4.

Table 3-4 System upgrades

Part number	Feature code	Description
4TA7A10969	Various	ThinkSystem SR950 4S Performance to 8S Upgrade Kit
7XG7A03955	AUN0	ThinkSystem SR950 2-CPU, 24 DIMM Compute System Board
7XG7A03956	Various	ThinkSystem SR950 Upper Compute Tray Kit

4S Performance to 8S Upgrade Kit, 4TA7A10969, includes the following components:

- ▶ 4 mid-plane UPI & sideband interconnects

The Compute System Board option, 7XG7A03955, includes the following components:

- ▶ One compute system board, without processors or memory

The Upper Compute Tray Kit, 7XG7A03956, includes the following components:

- ▶ 1x compute tray with 12 drive bays and 1 slot for storage adapter
- ▶ 1x compute system board
- ▶ 3x 4-bay drive bay fillers

Note: The Upper Compute Tray Kit does not include drive backplanes

Table 3-5 describes the before and after configurations and the components you will need to order to achieve the desired configuration.

Table 3-5 Upgrades to Server configurations (See Table 3-2 on page 53 for configuration definitions)

Initial configuration	Desired upgraded configuration	Parts to order and install
Configuration 1 (Four-Socket Performance) 2 processors installed	Configuration 1 (Four-Socket Performance) 4 processors installed	<ul style="list-style-type: none"> ▶ 1x Compute System Board, 7XG7A03955 ▶ 2x processor options (match installed) ▶ Memory DIMMs (matching installed is recommended)

Initial configuration	Desired upgraded configuration	Parts to order and install
Configuration 1 (Four-Socket Performance) 3 processors installed	Configuration 1 (Four-Socket Performance) 4 processors installed	<ul style="list-style-type: none"> ▶ 1x processor options (match installed) ▶ Memory DIMMs (matching installed is recommended)
Configuration 1 (Four-Socket Performance) 4 processors installed (including SAP HANA models and configurations)	Configuration 5 (Eight-Socket Storage Rich) 8 processors installed	<ul style="list-style-type: none"> ▶ 1x 4S Performance to 8S Upgrade Kit, 4TA7A10969 ▶ 1x Upper Compute Tray Kit, 7XG7A03956^a ▶ 1x Compute System Board, 7XG7A03955 ▶ 4x 81xx processor options (match installed)^b ▶ Memory DIMMs (matching installed is recommended) ▶ Drive backplanes for upper compute tray ▶ RAID adapter or HBA for upper compute tray ▶ Drives for upper compute tray ▶ Additional/replacement PCIe riser cards ▶ Lenovo Basic Hardware Installation services
Configuration 2 (Four-Socket Performance/Storage Rich) 4 processors installed (including customized SAP HANA models)	Configuration 5 (Eight-Socket Storage Rich) 8 processors installed	
Configuration 2 (Four-Socket Performance/Storage Rich) 2 processors installed	Configuration 2 (Four-Socket Performance/Storage Rich) 4 processors installed	<ul style="list-style-type: none"> ▶ 1x Compute System Board, 7XG7A03955 ▶ 2x processor options (match installed) ▶ Memory DIMMs (matching installed is recommended)
Configuration 2 (Four-Socket Performance/Storage Rich) 3 processors installed	Configuration 2 (Four-Socket Performance/Storage Rich) 4 processors installed	<ul style="list-style-type: none"> ▶ 1x processor options (match installed) ▶ Memory DIMMs (matching installed is recommended)
Configuration 3 (Four-Socket Upgradable) 2 processors installed	Configuration 3 (Four-Socket Upgradable) 4 processors installed	<ul style="list-style-type: none"> ▶ 1x Compute System Board, 7XG7A03955 ▶ 2x processor options (match installed) ▶ Memory DIMMs (matching installed is recommended)
Configuration 3 (Four-Socket Upgradable) 4 processors installed	Configuration 5 (Eight-Socket Storage Rich) 8 processors installed	<ul style="list-style-type: none"> ▶ 1x Upper Compute Tray Kit, 7XG7A03956^a ▶ 1x Compute System Board, 7XG7A03955 ▶ 4x 81xx processor options (match installed)^b ▶ Memory DIMMs (matching installed is recommended) ▶ Drive backplanes for upper compute tray ▶ RAID adapter or HBA for upper compute tray ▶ Drives for upper compute tray ▶ Additional/replacement PCIe riser cards
Configuration 4 (Four-Socket Upgradable / Storage Rich) 4 processors installed	Configuration 5 (Eight-Socket Storage Rich) 8 processors installed	<ul style="list-style-type: none"> ▶ 2x Compute System Board, 7XG7A03955 ▶ 4x 81xx processor options (match installed)^b ▶ Memory DIMMs (matching installed is recommended) ▶ Additional/replacement PCIe riser cards
Configuration 5 (Eight-Socket Storage Rich) 6 processors installed	Configuration 5 (Eight-Socket Storage Rich) 8 processors installed	<ul style="list-style-type: none"> ▶ 2x processor options (match installed) ▶ Memory DIMMs (matching installed is recommended)

a. The Upper Compute Tray Kit includes 1x Compute System Board

b. 8-socket configurations must use Intel Xeon Platinum (81xx) processors

Configuration notes:

- ▶ The added processors must match the processors already installed.

- ▶ When upgrading to Configuration 5 (8-socket configuration), Platinum-level (81xx) processors must be used, as listed in 3.6, “Processor options” on page 61. Gold processors do not support 8-socket configurations.
- ▶ For best performance, each processor added should match the memory configuration of existing processors (number & size of DIMMs)
- ▶ The ThinkSystem SR950 4S Performance to 8S Upgrade Kit, 4TA7A10969, contains the mid-chassis interconnects needed for an 8-socket configuration, replacing the existing interconnects. Lenovo installation is highly recommended due to the complexity of this upgrade. Installation cost will vary depending on location and scope of work.

3.4 Standard specifications

Table 3-6 lists the standard specifications for the SR950 server.

Table 3-6

Components	Specification
Machine type	7X11 - 1-year warranty 7X12 - 3-year warranty 7X13 - SAP HANA configurations with 3-year warranty
Form factor	4U rack
Processor	Up to eight Intel Xeon Processor Scalable Family of processors, either Gold or Platinum level processors (formerly codename “Skylake”). Supports processors with up to 28 cores, core speeds up to 3.6 GHz, and TDP ratings up to 205W. Processors and memory are located on system boards installed in compute trays: 2 processors per system board, 2 system boards per compute tray, and 2 compute trays in the server.
Chipset	Intel C624 chipset (formerly codename “Lewisburg”)
Memory	Up to 96 DIMM sockets (12 DIMMs per processor, installed on each system board). Lenovo TruDDR4 RDIMMs, LRDIMMs and 3DS RDIMMs are supported. Each processor has 6 memory channels, with 2 DIMMs per channel
Memory maximums	With RDIMMs: Up to 3 TB with 96x 32 GB RDIMMs and eight processors With LRDIMMs: Up to 6 TB with 96x 64 GB LRDIMMs and eight processors With 3DS RDIMMs: Up to 12 TB with 96x 128 GB 3DS RDIMMs and eight processors Note: Support of more than 768 GB per processor (more than 3 TB for a 4-processor system or 6 TB for an 8-processor system) requires M-suffix processors with 1.5 TB memory support. See 3.6, “Processor options” on page 61 for information.
Memory protection	ECC, Chipkill, Adaptive double-device data correction (ADDDC), memory mirroring, and memory rank sparing.
Disk drive bays	Up to 24x 2.5" hot-swap SAS/SATA drive bays. With AnyBay backplanes, 12 bays can be configured to support PCIe NVMe drives.
Maximum internal storage	Up to 184.3 TB with 24x 7.68 TB 2.5-inch SAS SSDs or Up to 48 TB with 24x 2 TB 2.5-inch NL SAS HDDs

Components	Specification
Storage controller	<ul style="list-style-type: none"> ▶ No embedded RAID <ul style="list-style-type: none"> – 12 Gb SAS/SATA RAID adapters: <ul style="list-style-type: none"> – RAID 530i (cacheless) supports RAID 0, 1, 10, 5, 50 – RAID 730-8i with 1GB cache supports RAID 0, 1, 10, 5, 50 – RAID 730-8i with 2GB flash-backed cache supports RAID 0, 1, 10, 5, 50, 6, 60 – RAID 930-8i with 2GB flash-backed cache supports RAID 0, 1, 10, 5, 50, 6, 60 – RAID 930-16i with 4GB flash-backed cache supports RAID 0, 1, 10, 5, 50, 6, 60 ▶ 12 Gb SAS/SATA non-RAID: 430-8i and 430-16i HBAs
Optical and tape bays	No internal bays.
Network interfaces	Integrated 10Gb Ethernet controller, based on the Intel Ethernet Connection X722. The server supports 1 of 6 available LOM adapters, providing either 2- or 4-port Gigabit Ethernet, 2- or 4-port 10 GbE with SFP+ connections or 2- or 4-port 10 GbE with RJ45 connections. One port can optionally be shared with the XClarity Controller (XCC) management processor for Wake-on-LAN and NC-SI support. 1 or 2 ML2 adapter slots for additional network adapters.
PCI Expansion slots	<p>Up to 17 PCIe 3.0 rear-accessible slots:</p> <ul style="list-style-type: none"> ▶ Slots 1-4 on a riser, one of: <ul style="list-style-type: none"> – Riser with 4x PCIe 3.0 x16 slots, all full-height half-length, all 75W – Riser with 2x PCIe 3.0 x16 slots, all full-height half-length, both 75W – Riser with 4x PCIe 3.0 x8 slots, full-height half-length, all 75W ▶ Slot 5: PCIe 3.0 x16, low-profile, 75W ▶ Slot 6: PCIe 3.0 x16, low-profile, 75W ▶ Slot 7: PCIe 3.0 x8, low-profile, 75W ▶ Slot 8: PCIe 3.0 x16 ML2, low-profile (supports NC-SI) ▶ Slot 9: PCIe 3.0 x8 LOM adapter slot, low-profile (supports NC-SI) ▶ Slots 10-15 on a riser, one of: <ul style="list-style-type: none"> – Riser 5x16+ML2: <ul style="list-style-type: none"> • Five PCIe 3.0 x16 slots, all full-height half-length, all 75W • One PCIe 3.0 x16 ML2 slot, full-height half-length, 30W (no NC-SI support) – Riser 2x16: Two PCIe 3.0 x16 slots, all full-height half-length, both 75W – Riser 4x8: Four PCIe 3.0 x8 slots, full-height half-length, all 75W ▶ Slot 16-17 on a riser: <ul style="list-style-type: none"> – PCIe 3.0 x8 low profile, 75W – PCIe 3.0 x8 low profile, 75W <p>Additional internal PCIe slots:</p> <ul style="list-style-type: none"> ▶ One or two PCIe 3.0 x8 slots for storage adapters ▶ M.2 adapter slot
Ports	Front: two USB 2.0 ports (one also for XCC local management), and one DB-15 VGA video port. Rear: two USB 3.0, one DB-15 VGA video, one DB-9 serial, and one 1 GbE RJ-45 systems management. The serial port can be shared with the XCC for serial redirection functions. Only one VGA port can be used at a time.
Cooling	Up to 12 redundant hot-swap single-rotor fans, 6 in each compute tray. Each tray has N+1 fan redundancy.

Components	Specification
Power supply	Up to four redundant hot-swap 1100 W or 1600 W power supplies (80 PLUS Platinum certified), supporting 200-240 V AC input. The 1100 W power supply also supports 100-127V AC input supply with a reduced capacity of 1050 W. No support for 240 V DC.
Hot-swap parts	Drives, power supplies, and fans.
Video	G200 graphics with 16 MB memory with 2D hardware accelerator, integrated into the XClarity Controller. Maximum resolution is 1920x1200 32bpp at 60Hz.
Security features	Power-on password, administrator's password, Trusted Platform Module (TPM), supporting TPM 1.2 or TPM 2.0. In China only, optional Trusted Cryptographic Module (TCM).
Systems management	XClarity Controller embedded management, XClarity Administrator centralized infrastructure delivery, XClarity Integrator plugins, and XClarity Energy Manager centralized server power management. XClarity Controller Enterprise functionality standard with the SR950. Light path diagnostics for local failure detection and reporting, with LCD diagnostic pull-out panel.
Operating systems	Microsoft Windows Server 2012 R2, 2016; RHEL 6 & 7 x64; SLES 11 & 12 x64; VMware ESXi 6.0, 6.5. See 3.23, "Operating system support" on page 94 for details.
Limited warranty	One year or three-year (model dependent) 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: 4-hour or 2-hour response time, 6-hour fix time, 1-year or 2-year warranty extension, software support for Lenovo hardware and some third-party applications.
Dimensions	Height: 173 mm (6.8 in.); depth EIA flange to rear: 800.0 mm (31.5 in.), depth overall: 835.0 mm (32.9 in.); width across top cover: 447.0 mm (17.6 in.), width to EIA flanges: 482.0 mm (19.0 in.).
Weight	Minimum configuration: 32.6 kg (71.9 lb), Maximum configuration: 58.7 kg (129.4 lb)

The server is shipped with the following items:

- ▶ Rail kit
- ▶ Documentation flyer
- ▶ Power cords (model and region dependent)

3.5 Models

For the latest models, see the Lenovo Press product guide:

- ▶ Standard models

<https://lenovopress.com/lp0647-thinksystem-sr950-server#models>

- ▶ SAP HANA configurations

<https://lenovopress.com/lp0647-thinksystem-sr950-server#sap-hana-configurations>

- ▶ SAP HANA models

<https://lenovopress.com/lp0647-thinksystem-sr950-server#sap-hana-models>

3.6 Processor options

The SR950 supports the Gold and Platinum level processors in the Intel Xeon Processor Scalable Family, including the “top bin” highest performing 205 W processors available from Intel.

The server supports up to four or eight processors, depending on the configuration selected, however processors are all installed within the 4U rack enclosure:

- ▶ For configurations that support up to eight processors, the server supports 2, 4, 6 or 8 processors
- ▶ For configurations of up to four processors (lower cost configuration, no upgrade path to eight processors), the server supports 2, 3 or 4 processors

As described in 3.1, “Components and connectors” on page 48, processors and memory are located on compute system boards - two processors per board. Two system boards are installed in a compute tray, and there are two compute trays in the server.

For configurations with only one compute tray, you can upgrade to add a second compute tray. You can also add additional system boards as needed. The ordering information for these upgrades is listed in Table 3-7.

Table 3-7 Compute tray and system board upgrades

Part number	Feature code	Description
7XG7A03955	AUN0	ThinkSystem SR950 2-CPU, 24 DIMM Compute System Board
7XG7A03956	Various	ThinkSystem SR950 Upper Compute Tray Kit

The Compute System Board option, 7XG7A03955, includes the following components:

- ▶ One compute system board, without processors or memory

The Upper Compute Tray Kit, 7XG7A03956, includes the following components:

- ▶ 1x compute tray with 12 drive bays and 1 slot for storage adapter
- ▶ 1x compute system board
- ▶ 3x 4-bay drive bay fillers

Note: the Upper Compute Tray Kit does not include drive backplanes

Table 3-8 on page 62 lists the supported processors. The table indicates which processors support 8-socket configurations and whether the processors support memory capacities greater than 768 GB per processor.

All processors listed support:

- ▶ Hyper-Threading Technology
- ▶ Turbo Boost Technology 2.0

Table 3-8 on page 62 also lists the number of UltraPath Interconnect (UPI) ports each processor has. UPI ports are used to connect processors together for high-speed communications.

- ▶ When a processor has 3 UPI ports, it can be used in a 4-socket mesh configuration where all processors are directly connected to every other processor. 5100 Series processors do not support mesh topology, since they only have 2 UPI ports.

- When a processor has only 2 UPI ports, it can only be used in a 4-socket ring configuration where all processors are connected to the 2 adjacent processors and the fourth processor is only connected indirectly via another processor.

The mesh and ring topologies are shown in Figure 3-9.




Figure 3-9 Mesh and ring processor topologies

Table 3-8 Processor options

Part number	Feature code	Intel Xeon processor	Suitable for 8-socket	Memory speed	Supports >768 GB per CPU	L3 cache	UPI links	AVX-512 units
4XG7A08857	AX8M	Gold 5115 10C 85W 2.4GHz	No	2400 MHz	No	13.75 MB	2	1
4XG7A09097	B20S	Gold 5117 14C 105W 2.0GHz	No	2400 MHz	No	19.25 MB	2	1
4XG7A08845	AX7D	Gold 5118 12C 105W 2.3GHz	No	2400 MHz	No	16.5 MB	2	1
4XG7A08864	AX7F	Gold 5119T 14C 85W 1.9GHz	No	2400 MHz	No	19.25 MB	2	1
4XG7A08832	AX7C	Gold 5120 14C 105W 2.2GHz	No	2400 MHz	No	19.25 MB	2	1
4XG7A08883	AX7E	Gold 5120T 14C 105W 2.2GHz	No	2400 MHz	No	19.25 MB	2	1
4XG7A08849	AX70	Gold 5122 4C 105W 3.6GHz	No	2666 MHz ^b	No	16.5 MB ^a	2	2 ^b
4XG7A08835	AWEX	Gold 6126 12C 125W 2.6GHz	No	2666 MHz	No	19.25 MB ^a	3	2
4XG7A08853	AX73	Gold 6126T 12C 125W 2.6GHz	No	2666 MHz	No	19.25 MB ^a	3	2
4XG7A08846	AX6Z	Gold 6128 6C 115W 3.4GHz	No	2666 MHz	No	19.25 MB ^a	3	2
4XG7A08833	AX6D	Gold 6130 16C 125W 2.1GHz	No	2666 MHz	No	22 MB	3	2
4XG7A08851	AX72	Gold 6130T 16C 125W 2.1GHz	No	2666 MHz	No	22 MB	3	2
4XG7A08838	AX6U	Gold 6132 14C 140W 2.6GHz	No	2666 MHz	No	19.25 MB	3	2

Part number	Feature code	Intel Xeon processor	Suitable for 8-socket	Memory speed	Supports >768 GB per CPU	L3 cache	UPI links	AVX-512 units
4XG7A08847	AX6Y	Gold 6134 8C 130W 3.2GHz	No	2666 MHz	No	24.75 MB ^a	3	2
4XG7A08863	AX7A	Gold 6134M 8C 130W 3.2GHz	No	2666 MHz	Yes	24.75 MB ^a	3	2
4XG7A08848	AX6W	Gold 6136 12C 150W 3.0GHz	No	2666 MHz	No	24.75 MB ^a	3	2
4XG7A08834	AX6Q	Gold 6138 20C 125W 2.0GHz	No	2666 MHz	No	27.5 MB	3	2
4XG7A08852	AX71	Gold 6138T 20C 125W 2.0GHz	No	2666 MHz	No	27.5 MB	3	2
4XG7A08836	AX6R	Gold 6140 18C 140W 2.3GHz	No	2666 MHz	No	24.75 MB	3	2
4XG7A08858	AX79	Gold 6140M 18C 140W 2.3GHz	No	2666 MHz	Yes	24.75 MB	3	2
4XG7A08839	AX6E	Gold 6142 16C 150W 2.6GHz	No	2666 MHz	No	22 MB	3	2
4XG7A08859	AX78	Gold 6142M 16C 150W 2.6GHz	No	2666 MHz	Yes	22 MB	3	2
7XG7A05003	AX6X	Gold 6144 8C 150W 3.5GHz	No	2666 MHz	No	24.75 MB ^a	3	2
7XG7A05004	AX6V	Gold 6146 12C 165W 3.2GHz	No	2666 MHz	No	24.75 MB ^a	3	2
4XG7A08840	AWEW	Gold 6148 20C 150W 2.4GHz	No	2666 MHz	No	27.5 MB	3	2
7XG7A05005	AX6T	Gold 6150 18C 165W 2.7GHz	No	2666 MHz	No	24.75 MB	3	2
4XG7A08837	AX6P	Gold 6152 22C 140W 2.1GHz	No	2666 MHz	No	30.25 MB	3	2
7XG7A03957	AX6S	Gold 6154 18C 200W 3.0GHz	No	2666 MHz	No	24.75 MB	3	2
4XG7A08854	AX6L	Platinum 8153 16C 125W 2.0GHz	Yes	2666 MHz	No	22 MB	3	2
4XG7A08850	AWEV	Platinum 8156 4C 105W 3.6GHz	Yes	2666 MHz	No	16.5 MB ^a	3	2
4XG7A08855	AX7B	Platinum 8158 12C 150W 3.0GHz	Yes	2666 MHz	No	24.75 MB ^a	3	2
4XG7A08841	AWGJ	Platinum 8160 24C 150W 2.1GHz	Yes	2666 MHz	No	33 MB	3	2
4XG7A08860	AX77	Platinum 8160M 24C 150W 2.1GHz	Yes	2666 MHz	Yes	33 MB	3	2

Part number	Feature code	Intel Xeon processor	Suitable for 8-socket	Memory speed	Supports >768 GB per CPU	L3 cache	UPI links	AVX-512 units
4XG7A08856	AX6N	Platinum 8160T 24C 150W 2.1GHz	Yes	2666 MHz	No	33 MB	3	2
4XG7A08842	AX6K	Platinum 8164 26C 150W 2.0GHz	Yes	2666 MHz	No	35.75 MB	3	2
7XG7A03958	AX6M	Platinum 8168 24C 205W 2.7GHz	Yes	2666 MHz	No	33 MB	3	2
4XG7A08843	AX6J	Platinum 8170 26C 165W 2.1GHz	Yes	2666 MHz	No	35.75 MB	3	2
4XG7A08862	AX76	Platinum 8170M 26C 165W 2.1GHz	Yes	2666 MHz	Yes	35.75 MB	3	2
4XG7A08844	AX6H	Platinum 8176 28C 165W 2.1GHz	Yes	2666 MHz	No	38.5 MB	3	2
4XG7A08861	AX75	Platinum 8176M 28C 165W 2.1GHz	Yes	2666 MHz	Yes	38.5 MB	3	2
7XG7A03959	AX6G	Platinum 8180 28C 205W 2.5GHz	Yes	2666 MHz	No	38.5 MB	3	2
7XG7A03960	AX74	Platinum 8180M 28C 205W 2.5GHz	Yes	2666 MHz	Yes	38.5 MB	3	2

a. L3 cache is 1.375 MB per core except with the processors indicated with this footnote

b. All Gold 5000-level processors, except the 5122, support 2400 MHz memory speeds and have one AVX-512 512-bit FMA units. The 5122 processor supports 2666 MHz and has two FMA units

3.7 Memory options

The SR950 uses Lenovo TruDDR4 memory operating at up to 2666 MHz. The server supports 12 DIMMs per processor, which corresponds to 48 DIMMs with four processors installed and 96 DIMMs when eight processors are installed. Each processor has six memory channels with two DIMMs per channel.

The SR950 supports three different types of DDR4 memory:

- ▶ Registered DIMMs (RDIMMs)
- ▶ Load Reduced DIMMs (LRDIMMs)
- ▶ Three-Dimensional Stacked RDIMMs (3DS RDIMMs)

The maximum memory possible for each DIMM type is listed in Table 3-9.

Table 3-9 Maximum memory configurations by DIMM type

Memory type ^a	Maximum memory possible
RDIMMs	3TB (96 x 32GB)
LRDIMMs	6TB (96 x 64GB)
3DS RDIMMs	12TB (96 x 128GB)

a. RDIMMs (SR, DR), LRDIMMs or 3DS DIMMs cannot be mixed.

As described in 3.1, “Components and connectors” on page 48, the memory is installed on compute system boards and there are two system boards per compute tray and one or two compute trays per server, depending on the configuration.

All DIMMs operate at a speed of 2666 MHz, both at 1 DIMM per channel and 2 DIMMs per channel. However, if the processor selected has a lower memory bus speed (e.g. 2400 MHz), then all DIMMs will operate at that lower speed.

Table 3-10 lists the memory options that are available for SR950.

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 to maximize performance and reliability. From a service and support standpoint, Lenovo TruDDR4 memory automatically assumes the system warranty, and Lenovo provides service and support worldwide.

Table 3-10 Memory Options

Part number	Feature code	Description	Maximum supported
RDIMMs			
7X77A01302	AUNB	ThinkSystem 16GB TruDDR4 2666 MHz (1Rx4 1.2V) RDIMM	96 (12 per processor)
7X77A01303	AUNC	ThinkSystem 16GB TruDDR4 2666 MHz (2Rx8 1.2V) RDIMM	96 (12 per processor)
7X77A01304	AUND	ThinkSystem 32GB TruDDR4 2666 MHz (2Rx4 1.2V) RDIMM	96 (12 per processor)
LRDIMMs			
7X77A01305	AUNE	ThinkSystem 64GB TruDDR4 2666 MHz (4Rx4 1.2V) LRDIMM	96 (12 per processor)
3DS RDIMMs			
4ZC7A08716	AUW5	ThinkSystem 64GB TruDDR4 2666MHz (4Rx4, 1.2V) 3DS RDIMM	96 (12 per processor)
7X77A01307	AUNF	ThinkSystem 128GB TruDDR4 2666 MHz (8Rx4 1.2V) 3DS RDIMM	96 (12 per processor)

The following rules apply when selecting the memory configuration:

- ▶ The server supports RDIMMs, LRDIMMs and 3DS RDIMMs. UDIMMs are not supported.
- ▶ Mixing RDIMMs and LRDIMMs is not supported.
- ▶ Mixing 3DS RDIMMs with either RDIMMs or LRDIMMs is not supported.
- ▶ Mixing x4 and x8 DIMMs is supported.

For best performance, consider the following:

- ▶ Populate memory DIMMs in quantities of 6 or 12 per processor, so that all memory channels are used.
- ▶ When mixing 16 GB and 32 GB DIMMs in the same configuration, use 16GB 2Rx8 DIMMs instead of 16 GB 1Rx4 DIMMs for better performance.
- ▶ Populate memory channels so they all have the same total memory capacity
- ▶ Ensure all memory controllers on a processor socket have the same DIMM configuration.
- ▶ All processor sockets on the same physical server should have the same DIMM configuration.

The following memory protection technologies are supported:

- ▶ ECC
- ▶ Chipkill (for x4-based memory DIMMs; look for “x4” in the DIMM description)
- ▶ Adaptive double-device data correction (ADDDC)
- ▶ Memory mirroring
- ▶ Memory rank sparing

If memory channel mirroring is used, then DIMMs must be installed in pairs or sets of three (minimum of one pair or set of three per processor), and all DIMMs in the pair or set of three must be identical in type and size. 50% of the installed capacity is available to the operating system.

If memory rank sparing is used, then 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 largest rank in the channel will be automatically selected as the spare rank. The amount of memory available to the operating system depends on the number, capacity and rank counts of the DIMMs installed.

For details about the recommended DIMM population sequence for each memory mode, see the *ThinkSystem SR950 Memory Population Reference*, available from:

http://thinksystem.lenovofiles.com/help/topic/7X12/pdf_files.html

3.8 Internal storage

The SR950 offers up to 24x 2.5-inch drive bays, depending on the server configuration selected (as described in 3.2, “Server configurations” on page 52). The server also supports either one or two M.2 drives, installed in an M.2 adapter.

In this section:

- ▶ 3.8.1, “Backplanes and drive bays” on page 66
- ▶ 3.8.2, “Adapters and cabling” on page 69
- ▶ 3.8.3, “M.2 drives” on page 70
- ▶ 3.8.4, “SED encryption key management with ISKLM” on page 71

3.8.1 Backplanes and drive bays

Drives are all located at the front of the server, 12 drives at the front of the upper tray and 12 at the front of the lower tray. Drives are connect to 4-drive backplanes in a 2-by-2 configuration, as shown in Figure 3-10 on page 67. All drive bays are 2.5-inch form factor.




Figure 3-10 Location of backplanes

There are two different backplanes available for the SR950:

- ▶ SAS/SATA backplane: Supports 4 SAS or SATA drives
- ▶ AnyBay backplane:
 - Top two drives support SAS, SATA, or NVMe interface drives (Lenovo AnyBay)
 - Bottom two drive bays support SAS or SATA drives only

Regular 2.5-inch SAS/SATA drive bays support only SAS or SATA drives; however, the Lenovo AnyBay drive bay design allows a choice of SATA, SAS, or U.2 (NVMe) PCIe drives. This design enables the flexibility to configure some of the bays with high-performance PCIe SSDs while still using the other bays for high-capacity HDDs, which is an ideal solution for storage-tiering.

Table 3-11 lists the two drive backplanes available.

Table 3-11 Drive backplanes

Part number	Feature code	Description	Maximum supported
7XB7A03966	AUN6	ThinkSystem 2.5" SAS/SATA 2x2 Bay Backplane Kit	6
7XB7A03965	AUN5	ThinkSystem 2.5" SAS/SATA/NVMe 2x2 Bay Backplane Kit	6

If the server is configured with all AnyBay backplanes, then the server supports up to 12 NVMe drives, depending on the server configuration. The NVMe-capable drive bays are the even-numbered bays as shown in Figure 3-11.

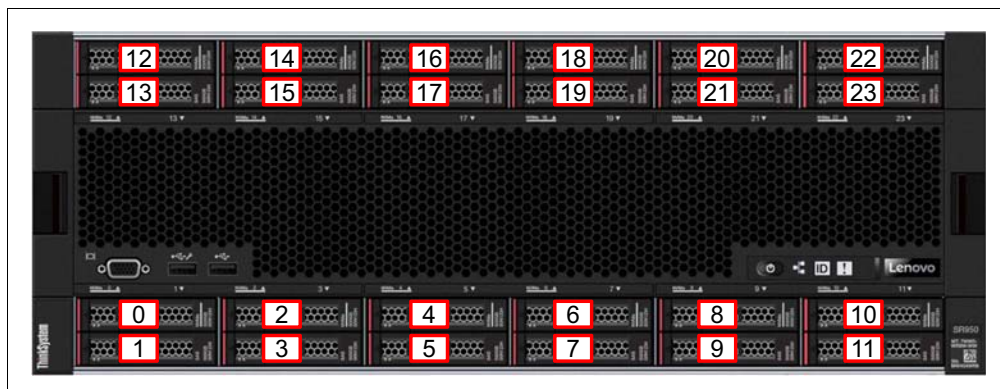


Figure 3-11 Drive bay numbering

The quantity and location of the AnyBay drive bays depends on the configuration, as listed in Table 3-12.

See 3.2, “Server configurations” on page 52 for details about the five server configurations.

Table 3-12 Drive bays and NVMe support for each SR950 configuration

Configuration	Maximum storage	Location of NVMe drive bays (see Figure 3-11 for bay numbering)
Configuration 1. Four-Socket Performance		
2 processors installed	12 drives (2 NVMe)	Bays 8 and 10
3 processors installed	12 drives (4 NVMe)	Bays 4, 6, 8, 10
4 processors installed	12 drives (6 NVMe)	Bays 0, 2, 4, 6, 8, 10
Configuration 2. Four-Socket Performance/ Storage Rich		
2 processors installed	24 drives (4 NVMe)	Bays 8, 10 and 20, 22
3 processors installed	24 drives (8 NVMe)	Bays 4, 6, 8, 10, and 16, 18, 20, 22
4 processors installed	24 drives (12 NVMe)	Bays 0, 2, 4, 6, 8, 10, 12, 14, 16, 18, 20, 22
Configuration 3. Four-Socket Upgradable		
2 processors installed	12 drives (2 NVMe)	Bays 8 and 10
4 processors installed	12 drives (6 NVMe)	Bays 0, 2, 4, 6, 8, 10
Configuration 4. Four-Socket Upgradable / Storage Rich		
4 processors installed	24 drives (4 NVMe)	Bays 8, 10 and 20, 22
Configuration 5. Eight-Socket Storage Rich		
6 processors installed	24 drives (8 NVMe)	Bays 0, 2, 4, 6, 8, 10 and 12, 14
8 processors installed	24 drives (12 NVMe)	Bays 0, 2, 4, 6, 8, 10, 12, 14, 16, 18, 20, 22

Backplane connections are as follows:

- The 4-drive SAS/SATA backplane has one SAS/SATA connector and connects to a supported RAID controller or SAS HBA installed in the dedicated storage adapter slot.
- The 4-drive AnyBay backplane has one SAS/SATA connector and connects to a supported RAID controller or SAS HBA plus one NVMe connector to connect to NVMe ports on one of the system boards in a compute tray or directly on the storage tray.

The SAS/SATA backplane is shown in Figure 3-12.




Figure 3-12 SAS/SATA backplane

3.8.2 Adapters and cabling

Each compute tray or storage tray includes one dedicated PCIe slot for a RAID adapter or SAS HBA.

Tip: The adapters are described in detail in 3.9, “Controllers for internal storage” on page 72.

SAS/SATA backplanes are connected to the RAID adapter or HBA installed in the dedicated slot in that tray. AnyBay backplanes (NVMe support on 2 of the 4 drive bays) are connected to the RAID adapter or HBA or to an NVMe port adjacent to processor on a system board.

Table 3-13 lists which processor each SAS/SATA storage adapter and which NVMe port each backplane is connected to.

Processor numbering is shown in Figure 3-6 on page 51.

Table 3-13 Processor connectivity for storage adapters and NVMe ports

Configuration	Slot location	4S Performance	4S Performance/ Storage Rich	4S Upgradeable	4S Upgradeable/ Storage Rich	8S Storage Rich
Processors installed		Two (1,2) or 3 three (1,2,4) or four (1,2,3,4)	Two (1,2) or 3 three (1,2,4) or four (1,2,3,4)	Two (1,2) or Four (1,2,3,4)	Four (1,2,5,6)	Six (1,2,3,4,5,7) or eight (1-8)
Storage adapter slot	Upper tray	NC	1	NC	5	5
Storage adapter slot	Lower tray	1	1	1	1	1
NVMe on backplane 4	Upper tray	NC	3	NC	NC	7
NVMe on backplane 5		NC	4	NC	NC	8
NVMe on backplane 6		NC	2	NC	6	6
NVMe on backplane 1	Lower tray	3 ^a	3	3 ^a	NC	3
NVMe on backplane 2		4 ^a	4	4 ^a	NC	4
NVMe on backplane 3		2	2	2	2	2

a. See note below:

Note: In 2, 3 and 6 processor configurations, the following backplanes cannot be AnyBay backplanes:

- In 2-processor configurations, processors 3 and 4 are not installed; this means drive backplanes 1 and 2 cannot be AnyBay backplanes.
- In 3-processor configurations, processor 3 is not installed; this means drive backplanes 1 and 4 cannot be AnyBay backplanes
- In 6-processor configurations, processors 6 & 8 are not installed; this means drive backplanes 5 and 6 cannot be AnyBay backplanes

3.8.3 M.2 drives

The server supports one or two M.2 form-factor SATA drives for use as an operating system boot solution. With two M.2 drives configured, the drives are configured by default as a RAID-1 mirrored pair for redundancy.

The M.2 drives install into an M.2 adapter which in turn is installed in a dedicated slot on the system board. See the internal view of the server in 3.1, “Components and connectors” on page 48 for the location of the M.2 slot.

There are two M.2 adapters supported as listed in

Table 3-14 M.2 adapters (enablement kits)

Part number	Feature code	Description	Maximum supported
7Y37A01092	AUMU	ThinkSystem M.2 Enablement Kit (contains the Single M.2 Boot Adapter; supports 1 drive)	1
7Y37A01093	AUMV	ThinkSystem M.2 with Mirroring Enablement Kit (contains the Dual M.2 Boot Adapter, supports 1 or 2 drives)	1

A Dual M.2 Boot Adapter with one 128GB M.2 drive partially inserted is shown in Figure 3-13. The second M.2 drive is installed on the other side of the adapter.



Figure 3-13 Dual M.2 Boot Adapter and a 128 GB M.2 drive

Features of the Dual M.2 Boot Adapter:

- ▶ PCIe 2.0 x2 host interface (connects to the PCH)
- ▶ Based on the Marvell 88SE9230 6 Gbps SATA controller
- ▶ Supports two 6 Gbps SATA M.2 drives (it is not supported to have only one drive installed)
- ▶ Supports 3 different physical sizes of M.2 drives: 42mm (2242), 60mm (2260) and 80mm (2280)
- ▶ RAID functionality provided by the M.2 adapter
- ▶ RAID 1 by default; also supports RAID 0 and JBOD
- ▶ UEFI-based settings to enable/disable RAID mode and to review inventory
- ▶ Adapter and drive firmware update using Lenovo firmware tools
- ▶ Management via I2C interface

M.2 naming convention: 2242, 2260 and 2280 are the industry terms for the M.2 drive dimensions. For example, 2280 corresponds to a drive that is 22mm wide and 80mm long.

Features of the Single M.2 Boot Adapter:

- ▶ 6 Gbps SATA host interface (connects to the PCH)
- ▶ Supports one 6 Gbps SATA M.2 drive
- ▶ Supports 3 different physical sizes of M.2 drives: 42mm (2242), 60mm (2260) and 80mm (2280)
- ▶ Drive firmware update using Lenovo firmware tools
- ▶ Management via I2C interface
- ▶ VPD reporting of adapter inventory

The Single M.2 Boot Adapter is shown in Figure 3-14 with the 32GB M.2 drive installed.



Figure 3-14 Single M.2 Boot Adapter and a 32 GB M.2 drive

Supported drives are listed in 3.10, “Internal drive options” on page 74.

3.8.4 SED encryption key management with ISKLM

The server supports self-encrypting drives (SEDs) and FIPS 140-2 drives as described in the Internal drive options section. To effectively manage a large deployment of these drives in Lenovo servers, IBM Security Key Lifecycle Manager (SKLM) offers a centralized key management solution. A Lenovo Feature on Demand (FoD) upgrade is used to enable this SKLM support in the management processor of the server.

Table 3-15 lists the part numbers and feature codes for the upgrades.

Table 3-15 FoD upgrades for SKLM support

Part number	Feature code	Description
Security Key Lifecycle Manager - FoD (United States, Canada, Asia Pacific, and Japan)		
00D9998	A5U1	SKLM for System x/ThinkSystem w/SEDs - FoD per Install with 1 year S&S
00D9999	AS6C	SKLM for System x/ThinkSystem w/SEDs - FoD per Install with 3 year S&S
Security Key Lifecycle Manager - FoD (Latin America, Europe, Middle East, and Africa)		
00FP648	A5U1	SKLM for System x/ThinkSystem w/SEDs - FoD per Install with 1 year S&S
00FP649	AS6C	SKLM for System x/ThinkSystem w/SEDs - FoD per Install with 3 year S&S

3.9 Controllers for internal storage

The SR950 supports internal SAS and SATA drives with the addition of a RAID adapter or HBA. The adapter is installed in a dedicated slot in each compute tray. For NVMe support, the server uses NVMe ports that are on the compute system boards or on the storage tray if one is configured. Table 3-16 lists the supported adapters.

Tip: Unlike System x RAID adapters, no upgrades are available for these RAID adapters. All supported functions are included with the base part numbers

Table 3-16 Controllers for internal storage

Part number	Feature code	Description	Number of ports	Maximum supported
SAS/SATA HBA				
7Y37A01088	AUNL	ThinkSystem 430-8i SAS/SATA 12Gb HBA	8	2
7Y37A01089	AUNM	ThinkSystem 430-16i SAS/SATA 12Gb HBA	12 ^a	2
RAID adapters				
7Y37A01082	AUNG	ThinkSystem RAID 530-8i PCIe 12Gb Adapter	8	2
7Y37A01083 ^b	AUNH ^b	ThinkSystem RAID 730-8i 1GB Cache PCIe 12Gb Adapter	8	2
4Y37A09722 ^c	B4RQ ^c	ThinkSystem RAID 730-8i 2GB Flash PCIe 12Gb Adapter	8	2
7Y37A01084	AUNJ	ThinkSystem RAID 930-8i 2GB Flash PCIe 12Gb Adapter	8	2
7Y37A01085	AUNK	ThinkSystem RAID 930-16i 4GB Flash PCIe 12Gb Adapter	12 ^a	2
4Y37A09721	B31E	ThinkSystem RAID 930-16i 8GB Flash PCIe 12Gb Adapter	12 ^a	2

a. The 430-16i HBA and RAID 930-16i adapter support 12 drives when used in the SR950

b. The RAID 730-8i 1GB Cache adapter is not available in USA and Canada

c. The RAID 730-8i 2GB Flash adapter is only available in Asia Pacific countries and in China.

Table 3-18 on page 73 compares the functions of the RAID storage adapters.

Table 3-17 Comparison of internal RAID storage controllers

Feature	RAID 530-8i	RAID 730-8i 1GB	RAID 730-8i 2GB	RAID 930-8i	RAID 930-16i
Adapter type	RAID controller	RAID controller	RAID controller	RAID controller	RAID controller
Part number	7Y37A01082	7Y37A01083	4Y37A09722	7Y37A01084	7Y37A01085 4Y37A09721
Form factor	PCIe HHHL	PCIe low profile	PCIe low profile	PCIe HHHL	PCIe HHHL
Controller chip	LSI SAS3408	LSI SAS3108	LSI SAS3108	LSI SAS3508	LSI SAS3516
Host interface	PCIe 3.0 x8	PCIe 3.0 x8	PCIe 3.0 x8	PCIe 3.0 x8	PCIe 3.0 x8
Port interface	12 Gb SAS	12 Gb SAS	12 Gb SAS	12 Gb SAS	12 Gb SAS
Number of ports	8	8	8	8	12 ^a
Port connectors	2x Mini-SAS HD x4 (SFF-8643)	2x Mini-SAS HD x4 (SFF-8643)	2x Mini-SAS HD x4 (SFF-8643)	2x Mini-SAS HD x4 (SFF-8643)	4x Mini-SAS HD x4 (SFF-8643)

Feature	RAID 530-8i	RAID 730-8i 1GB	RAID 730-8i 2GB	RAID 930-8i	RAID 930-16i
Drive interface	SAS, SATA	SAS, SATA	SAS, SATA	SAS, SATA	SAS, SATA
Drive type	HDD, SED, SSD	HDD, SSD	HDD, SED, SSD	HDD, SED, SSD	HDD, SED, SSD
Hot-swap drives	Yes	Yes	Yes	Yes	Yes
Max devices	8	8	8	8	12 ^b
RAID levels	0, 1, 10, 5, 50	0, 1, 10, 5, 50	0, 1, 10, 5, 50, 6, 60	0, 1, 10, 5, 50, 6, 60	0, 1, 10, 5, 50, 6, 60
JBOD mode	Yes	Yes	Yes	Yes	Yes
Cache	None	1GB (Standard)	2GB (Standard)	2GB (Standard)	4GB or 8GB (Standard)
CacheVault cache protection	No	No	Yes (Flash)	Yes (Flash)	Yes (Flash)
Performance Accelerator (FastPath)	Yes	No	Yes	Yes	Yes
SSD Caching (CacheCade Pro 2.0)	No	No	No	No	No
SED support ^b	Yes	Yes	Yes	Yes	Yes

a. With the SR950, the 430-16i HBA and RAID 930-16i support up to 12 drives.

b. SED (self-encrypting drive) support of the SAS HBAs is by using software on the server (SED commands are passed through the HBA to the drives). SED support by RAID controllers is provided using the built-in MegaRAID SafeStore functionality of the adapter.

Table 3-18 compares the functions of the supported internal HBAs.

Table 3-18 Comparison of HBA internal storage controllers

Feature	430-8i	430-16i
Adapter type	HBA	HBA
Part number	7Y37A01088	7Y37A01089
Form factor	PCIe low profile	PCIe low profile
Controller chip	LSI SAS3408	LSI SAS3416
Host interface	PCIe 3.0x8	PCIe 3.0x8
Port interface	12 Gb SAS	12 Gb SAS
Number of ports	8	16
Port connectors	2x Mini-SAS HD x4 (SFF-8643)	4x Mini-SAS HD x4 (SFF-8643)
Drive interface	SAS, SATA	SAS, SATA
Drive type	HDD, SSD, SED ^a	HDD, SSD, SED ^a
Hot-swap drives	Yes	Yes
Max devices	8	12 ^b

Feature	430-8i	430-16i
RAID levels	No RAID	No RAID
JBOD mode	Yes	Yes
Cache	No	No
CacheVault cache protection	No	No
Performance Accelerator (FastPath)	No	No
SSD Caching (CacheCade Pro 2.0)	No	No
SED support ^a	Yes	Yes

a. SED (self-encrypting drive) support of the SAS HBAs is by using software on the server (SED commands are passed through the HBA to the drives). SED support by RAID controllers is provided using the built-in MegaRAID SafeStore functionality of the adapter.

b. With the SR950, the 430-16i HBA and RAID 930-16i support up to 12 drives.

3.10 Internal drive options

The following tables list the hard disk drive and solid-state drive options for the internal disk storage of the server.

- ▶ Table 3-19: 2.5-inch hot-swap 12 Gb SAS HDDs
- ▶ Table 3-21 on page 75: 2.5-inch hot-swap 12 Gb SAS SSDs
- ▶ Table 3-22 on page 76: 2.5-inch hot-swap 6 Gb SAS/SATA SSDs
- ▶ Table 3-23 on page 77: 2.5-inch U.2 NVMe SSDs
- ▶ Table 3-24 on page 78: M.2 drives

Tip: M.2 drive support is listed in 3.8.3, “M.2 drives” on page 70.

Table 3-19 2.5-inch hot-swap 12 Gb SAS HDDs

Part number	Feature	Description	Maximum supported
2.5-inch hot-swap HDDs - 12 Gb SAS 10K			
7XB7A00024	AULY	ThinkSystem 2.5" 300GB 10K SAS 12Gb Hot Swap 512n HDD	24
7XB7A00025	AULZ	ThinkSystem 2.5" 600GB 10K SAS 12Gb Hot Swap 512n HDD	24
7XB7A00026	AUM0	ThinkSystem 2.5" 900GB 10K SAS 12Gb Hot Swap 512n HDD	24
7XB7A00027	AUM1	ThinkSystem 2.5" 1.2TB 10K SAS 12Gb Hot Swap 512n HDD	24
7XB7A00028	AUM2	ThinkSystem 2.5" 1.8TB 10K SAS 12Gb Hot Swap 512e HDD	24
7XB7A00069	B0YS	ThinkSystem 2.5" 2.4TB 10K SAS 12Gb Hot Swap 512e HDD	24
2.5-inch hot-swap HDDs - 12 Gb SAS 15K			
7XB7A00021	AULV	ThinkSystem 2.5" 300GB 15K SAS 12Gb Hot Swap 512n HDD	24
7XB7A00022	AULW	ThinkSystem 2.5" 600GB 15K SAS 12Gb Hot Swap 512n HDD	24
7XB7A00023	AULX	ThinkSystem 2.5" 900GB 15K SAS 12Gb Hot Swap 512e HDD	24
2.5-inch hot-swap HDDs - 12 Gb NL SAS			

Part number	Feature	Description	Maximum supported
7XB7A00034	AUM6	ThinkSystem 2.5" 1TB 7.2K SAS 12Gb Hot Swap 512n HDD	24
7XB7A00035	AUM7	ThinkSystem 2.5" 2TB 7.2K SAS 12Gb Hot Swap 512n HDD	24
2.5-inch hot-swap SED HDDs - 12 Gb SAS 10K			
7XB7A00030	AUM4	ThinkSystem 2.5" 300GB 10K SAS 12Gb Hot Swap 512n HDD SED	24
7XB7A00031	AUM5	ThinkSystem 2.5" 600GB 10K SAS 12Gb Hot Swap 512n HDD SED	24
7XB7A00070	B0YV	ThinkSystem 2.5" 2.4TB 10K SAS 12Gb Hot Swap 512e HDD FIPS	24

Table 3-20 2.5-inch hot-swap 6 Gb SAS/SATA HDDs

Part number	Feature	Description	Maximum supported
2.5-inch hot-swap HDDs - 6 Gb NL SATA			
7XB7A00036	AUUE	ThinkSystem 2.5" 1TB 7.2K SATA 6Gb Hot Swap 512n HDD	24
7XB7A00037	AUUJ	ThinkSystem 2.5" 2TB 7.2K SATA 6Gb Hot Swap 512e HDD	24

Table 3-21 2.5-inch hot-swap 12 Gb SAS SSDs

Part number	Feature	Description	Maximum supported
2.5-inch hot-swap SSDs - 12 Gb SAS - Enterprise Capacity			
4XB7A13646	B4A6	ThinkSystem 2.5" PM1643 7.68TB Capacity SAS 12Gb Hot Swap SSD	24
4XB7A13645	B4A7	ThinkSystem 2.5" PM1643 3.84TB Capacity SAS 12Gb Hot Swap SSD	24
2.5-inch hot-swap SSDs - 12 Gb SAS - Enterprise Performance (10+ DWPD)			
4XB7A10219	B4Y4	ThinkSystem 2.5" SS530 400GB Performance SAS 12Gb Hot Swap SSD	24
4XB7A10230	B4Y5	ThinkSystem 2.5" SS530 800GB Performance SAS 12Gb Hot Swap SSD	24
4XB7A10231	B4Y6	ThinkSystem 2.5" SS530 1.6TB Performance SAS 12Gb Hot Swap SSD	24
4XB7A10232	B4Y7	ThinkSystem 2.5" SS530 3.2TB Performance SAS 12Gb Hot Swap SSD	24
7N47A00124	AUMG	ThinkSystem 2.5" HUSMM32 400GB Performance SAS 12Gb Hot Swap SSD	24
7N47A00125	AUMH	ThinkSystem 2.5" HUSMM32 800GB Performance SAS 12Gb Hot Swap SSD	24
7N47A00126	AVRB	ThinkSystem 2.5" HUSMM32 1.6TB Performance SAS 12Gb Hot Swap SSD	24
2.5-inch hot-swap SSDs - 12 Gb SAS - Enterprise Mainstream (3-5 DWPD)			
4XB7A13653	B4A0	ThinkSystem 2.5" PM1645 800GB Mainstream SAS 12Gb Hot Swap SSD	24
4XB7A13654	B4A1	ThinkSystem 2.5" PM1645 1.6TB Mainstream SAS 12Gb Hot Swap SSD	24
4XB7A13655	B4A2	ThinkSystem 2.5" PM1645 3.2TB Mainstream SAS 12Gb Hot Swap SSD	24
7N47A00117	AUMC	ThinkSystem 2.5" PM1635a 400GB Mainstream SAS 12Gb Hot Swap SSD	24
2.5-inch hot-swap SED SSDs - 12 Gb SAS - Enterprise Performance (10+ DWPD)			
7SD7A05754	B11P	ThinkSystem 2.5" HUSMM32 400GB Performance SAS 12Gb HS SSD FIPS	24

Part number	Feature	Description	Maximum supported
7SD7A05753	B11Q	ThinkSystem 2.5" HUSMM32 800GB Performance SAS 12Gb HS SSD FIPS	24
7SD7A05752	B11R	ThinkSystem 2.5" HUSMM32 1.6TB Performance SAS 12Gb HS SSD FIPS	24

Table 3-22 2.5-inch hot-swap 6 Gb SAS/SATA SSDs

Part number	Feature	Description	Maximum supported
2.5-inch hot-swap SSDs - 6 Gb SATA - Enterprise Mainstream (3-5 DWPD)			
4XB7A13633	B49L	ThinkSystem 2.5" Intel S4610 240GB Mainstream SATA 6Gb Hot Swap SSD	24
4XB7A13634	B49M	ThinkSystem 2.5" Intel S4610 480GB Mainstream SATA 6Gb Hot Swap SSD	24
4XB7A13635	B49N	ThinkSystem 2.5" Intel S4610 960GB Mainstream SATA 6Gb Hot Swap SSD	24
4XB7A13636	B49P	ThinkSystem 2.5" Intel S4610 1.92TB Mainstream SATA 6Gb Hot Swap SSD	24
4XB7A13637	B49Q	ThinkSystem 2.5" Intel S4610 3.84TB Mainstream SATA 6Gb Hot Swap SSD	24
4XB7A10237	B488	ThinkSystem 2.5" 5200 240GB Mainstream SATA 6Gb Hot Swap SSD	24
4XB7A10238	B489	ThinkSystem 2.5" 5200 480GB Mainstream SATA 6Gb Hot Swap SSD	24
4XB7A10239	B48A	ThinkSystem 2.5" 5200 960GB Mainstream SATA 6Gb Hot Swap SSD	24
4XB7A10240	B48B	ThinkSystem 2.5" 5200 1.92TB Mainstream SATA 6Gb Hot Swap SSD	24
4XB7A10241	B48C	ThinkSystem 2.5" 5200 3.84TB Mainstream SATA 6Gb Hot Swap SSD	24
7SD7A05761	B110	ThinkSystem 2.5" 5100 3.84TB Mainstream SATA 6Gb Hot Swap SSD	24
2.5-inch hot-swap SSDs - 6 Gb SATA - Enterprise Entry (<3 DWPD)			
4XB7A10247	B498	ThinkSystem 2.5" Intel S4510 240GB Entry SATA 6Gb Hot Swap SSD	24
4XB7A10248	B499	ThinkSystem 2.5" Intel S4510 480GB Entry SATA 6Gb Hot Swap SSD	24
4XB7A10249	B49A	ThinkSystem 2.5" Intel S4510 960GB Entry SATA 6Gb Hot Swap SSD	24
4XB7A13622	B49B	ThinkSystem 2.5" Intel S4510 1.92TB Entry SATA 6Gb Hot Swap SSD	24
4XB7A13623	B49C	ThinkSystem 2.5" Intel S4510 3.84TB Entry SATA 6Gb Hot Swap SSD	24
4XB7A10195	B34H	ThinkSystem 2.5" PM883 240GB Entry SATA 6Gb Hot Swap SSD	24
4XB7A10196	B34J	ThinkSystem 2.5" PM883 480GB Entry SATA 6Gb Hot Swap SSD	24
4XB7A10197	B34K	ThinkSystem 2.5" PM883 960GB Entry SATA 6Gb Hot Swap SSD	24
4XB7A10198	B34L	ThinkSystem 2.5" PM883 1.92TB Entry SATA 6Gb Hot Swap SSD	24
4XB7A10199	B34M	ThinkSystem 2.5" PM883 3.84TB Entry SATA 6Gb Hot Swap SSD	24
4XB7A10200	B4D2	ThinkSystem 2.5" PM883 7.68TB Entry SATA 6Gb Hot Swap SSD	24
4XB7A10153	B2X2	ThinkSystem 2.5" 5200 480GB Entry SATA 6Gb Hot Swap SSD	24
4XB7A10154	B2X3	ThinkSystem 2.5" 5200 960GB Entry SATA 6Gb Hot Swap SSD	24
4XB7A10155	B2X4	ThinkSystem 2.5" 5200 1.92TB Entry SATA 6Gb Hot Swap SSD	24
4XB7A10156	B2X5	ThinkSystem 2.5" 5200 3.84TB Entry SATA 6Gb Hot Swap SSD	24

Part number	Feature	Description	Maximum supported
4XB7A10157	B2X6	ThinkSystem 2.5" 5200 7.68TB Entry SATA 6Gb Hot Swap SSD	24
4XB7A08505	B10R	ThinkSystem 2.5" 5100 3.84TB Entry SATA 6Gb Hot Swap SSD	24

Table 3-23 2.5-inch U.2 NVMe

Part number	Feature	Description	Maximum supported
2.5-inch SSDs - NVMe - Enterprise Performance (10+ DWPD)			
7N47A00081	AUMJ	ThinkSystem U.2 Intel P4800X 375GB Performance NVMe PCIe 3.0 x4 Hot Swap SSD	12
7N47A00083	B2ZJ	ThinkSystem U.2 Intel P4800X 750GB Performance NVMe PCIe 3.0 x4 Hot Swap SSD	12
2.5-inch SSDs - NVMe - Enterprise Mainstream (3-5 DWPD)			
4XB7A13936	B589	ThinkSystem U.2 Intel P4610 1.6TB Mainstream NVMe PCIe3.0 x4 Hot Swap SSD	12
4XB7A13937	B58A	ThinkSystem U.2 Intel P4610 3.2TB Mainstream NVMe PCIe3.0 x4 Hot Swap SSD	12
4XB7A13938	B58B	ThinkSystem U.2 Intel P4610 6.4TB Mainstream NVMe PCIe3.0 x4 Hot Swap SSD	12
4XB7A08516	B21W	ThinkSystem U.2 KCM51V 800GB Mainstream NVMe PCIe 3.0 x4 Hot Swap SSD	12
4XB7A08517	B21X	ThinkSystem U.2 KCM51V 1.6TB Mainstream NVMe PCIe 3.0 x4 Hot Swap SSD	12
4XB7A08518	B21Y	ThinkSystem U.2 KCM51V 3.2TB Mainstream NVMe PCIe 3.0 x4 Hot Swap SSD	12
4XB7A08519	B2XJ	ThinkSystem U.2 KCM51V 6.4TB Mainstream NVMe PCIe 3.0 x4 Hot Swap SSD	12
7SD7A05772	B11J	ThinkSystem U.2 Intel P4600 1.6TB Mainstream NVMe PCIe3.0 x4 Hot Swap SSD	12
7SD7A05771	B11K	ThinkSystem U.2 Intel P4600 3.2TB Mainstream NVMe PCIe3.0 x4 Hot Swap SSD	12
7SD7A05770	B11L	ThinkSystem U.2 Intel P4600 6.4TB Mainstream NVMe PCIe3.0 x4 Hot Swap SSD	12
2.5-inch SSDs - NVMe - Enterprise Entry (<3 DWPD)			
4XB7A10202	B58F	ThinkSystem U.2 Intel P4510 1.0TB Entry NVMe PCIe3.0 x4 Hot Swap SSD	12
4XB7A10204	B58G	ThinkSystem U.2 Intel P4510 2.0TB Entry NVMe PCIe3.0 x4 Hot Swap SSD	12
4XB7A10205	B58H	ThinkSystem U.2 Intel P4510 4.0TB Entry NVMe PCIe3.0 x4 Hot Swap SSD	12
4XB7A08513	B58J	ThinkSystem U.2 Intel P4510 8.0TB Entry NVMe PCIe3.0 x4 Hot Swap SSD	12
4XB7A10175	B34N	ThinkSystem U.2 PM983 1.92TB Entry NVMe PCIe 3.0 x4 Hot Swap SSD	12
4XB7A10176	B34P	ThinkSystem U.2 PM983 3.84TB Entry NVMe PCIe 3.0 x4 Hot Swap SSD	12

Part number	Feature	Description	Maximum supported
7N47A00984	AUV0	ThinkSystem U.2 PM963 1.92TB Entry NVMe PCIe 3.0 x4 Hot Swap SSD	12
7N47A00985	AUUU	ThinkSystem U.2 PM963 3.84TB Entry NVMe PCIe 3.0 x4 Hot Swap SSD	12
4XB7A10177	B4D3	ThinkSystem U.2 PM983 7.68TB Entry NVMe PCIe 3.0 x4 Hot Swap SSD	12
7SD7A05779	B11C	ThinkSystem U.2 Intel P4500 1.0TB Entry NVMe PCIe3.0 x4 Hot Swap SSD	12
7SD7A05777	B11E	ThinkSystem U.2 Intel P4500 4.0TB Entry NVMe PCIe3.0 x4 Hot Swap SSD	12

Note: NVMe PCIe SSDs support informed hot removal and hot insertion, provided the operating system supports PCIe SSD hot--swap.

Table 3-24 M.2 drives

Part number	Feature	Description	Maximum supported
M.2 SSDs - 6 Gb SATA - Enterprise Entry (<3 DWPD)			
7N47A00129	AUUL	ThinkSystem M.2 CV1 32GB SATA 6Gbps Non-Hot-Swap SSD	2
7N47A00130	AUUV	ThinkSystem M.2 CV3 128GB SATA 6Gbps Non-Hot-Swap SSD	2
7SD7A05703	B11V	ThinkSystem M.2 5100 480GB SATA 6Gbps Non-Hot-Swap SSD	2

3.11 Optical drives

The server does not support an internal optical drive. An external USB optical drive is available, listed Table 3-25.

Table 3-25 External optical drive

Part number	Feature code	Description
7XA7A05926	AVV8	ThinkSystem External USB DVD RW Optical Disk Drive

3.12 I/O expansion options

The server supports the following PCIe slots, depending on the server configuration selected:

- ▶ Slots 1-17 are PCIe 3.0 slots accessible from the rear of the server
- ▶ One internal dedicated M.2 slot for an M.2 adapter located internally on the I/O tray
- ▶ Two internal dedicated PCIe 3.0 x8 slots for storage adapters, located internally in each compute tray or storage tray

Figure 3-15 shows the locations of the slots at the rear of the server.




Figure 3-15 PCIe 3.0 slots at the rear of the server

The 17 PCIe 3.0 rear-accessible slots are as follows:

- ▶ Slots 1-4 are via Riser Slot 1, one riser card of:
 - Riser with 4x PCIe 3.0 x16 slots, all full-height half-length, all 75W
 - Riser with 2x PCIe 3.0 x16 slots, all full-height half-length, both 75W
 - Riser with 4x PCIe 3.0 x8 slots, full-height half-length, all 75W
- ▶ Slot 5: PCIe 3.0 x16, low profile, 75W
- ▶ Slot 6: PCIe 3.0 x16, low profile, 75W
- ▶ Slot 7: PCIe 3.0 x8, low profile, 75W
- ▶ Slot 8: PCIe 3.0 x16 ML2, low profile (supports NC-SI)
- ▶ Slot 9: PCIe 3.0 x8 LOM adapter slot, low profile (supports NC-SI)
- ▶ Slots 10-15 are via Riser Slot 2, one riser card of:
 - Riser 5x16+ML2
 - Five PCIe 3.0 x16 slots, all full-height half-length, all 75W
 - One PCIe 3.0 x16 ML2 slot, full-height half-length, 30W (no NC-SI support)
 - Riser 2x16: Two PCIe 3.0 x16 slots, all full-height half-length, both 75W
 - Riser 4x8: Four PCIe 3.0 x8 slots, full-height half-length, all 75W
- ▶ Slot 16-17 on a riser:
 - PCIe 3.0 x8 low profile, 75W
 - PCIe 3.0 x8 low profile, 75W

Additional internal PCIe slots:

- ▶ One or two PCIe 3.0 x8 slots for storage adapters, one in each compute tray
- ▶ M.2 adapter slot

NC-SI (Network Controller Sideband Interface or NCSI) support means the primary port on the adapter installed in that slot can be configured as shared between XCC out-of-band management traffic and production Ethernet traffic.

For riser slots 1 and 2, the server supports the following three riser cards. These are shown in Figure 3-16 on page 80.

- ▶ Riser 5x16+ML2:
 - Five PCIe 3.0 x16 slots, all full-height half-length, all 75W
 - One PCIe 3.0 x16 ML2 slot, full-height half-length, 30W (no NC-SI support)

Note: When the 5x16+ML2 riser is installed in riser slot 1, the bottom two slots (including the ML2) are not accessible and cannot be used.

- ▶ Riser 2x16: Two PCIe 3.0 x16 slots, all full-height half-length, both 75W
- ▶ Riser 4x8: Four PCIe 3.0 x8 slots, full-height half-length, all 75W

Note: The ML2 slot on the riser card does not support NC-SI.




Figure 3-16 Riser cards for riser slots 1 and 2

Slot numbering is as follows:

- ▶ When installed in riser slot 1:
 - Riser 5x16+ML2 becomes slots 1-4 (bottom two slots inaccessible)
 - Riser 2x16 becomes slots 3-4
 - Riser 4x8 becomes slots 1-4
- ▶ When installed in riser slot 2:
 - Riser 5x16+ML2 becomes slots 10-15
 - Riser 2x16 becomes slots 12-13
 - Riser 4x8 becomes slots 10-13

Note: When the 5x16+ML2 riser card is used in Riser slot 1, only the top 4 slots (all x16) are accessible. The lower x16 and x16 ML2 slots cannot physically be accessed or used.

Slots 16 and 17 are mounted on a single 2x8 riser card and are as follows:

- ▶ Slot 16: PCIe 3.0 x8, full-height half-length, 75W
- ▶ Slot 17: PCIe 3.0 x8, full-height half-length, 75W

The 2x8 riser card is shown in Figure 3-17.




Figure 3-17 2x8 riser card

Ordering information for the riser cards is listed in Table 3-26.

Table 3-26 Riser card ordering information

Part number	Feature code	Description
Riser cards for riser slots 1 and 2		
7XC7A03963	AUN3	ThinkSystem SR950 up to (5) x16 + x16 ML2 PCIe Riser
7XC7A03961	AUN1	ThinkSystem SR950 (2) x16 PCIe Riser
7XC7A03962	AUN2	ThinkSystem SR950 (4) x8 PCIe Riser
Riser card for 2x8 riser slot		
7XC7A03964	AUN4	ThinkSystem SR950 (2) x8 PCIe Riser

3.12.1 Slot availability by server configuration

The riser cards required for each server configuration are listed in the following table. See 3.2, “Server configurations” on page 52 for details.

Table 3-27 Riser card requirements

Riser	4S Performance	4S Performance/ Storage Rich	4S Upgradeable	4S Upgradeable/ Storage Rich	8S Storage Rich
Riser 1	4x8 or 2x16	4x8 or 2x16	4x8 or 2x16	5x16+ML2 ^a	5x16+ML2 ^a
Riser 2	4x8 or 2x16	4x8 or 2x16	4x8 or 2x16	5x16+ML2 ^a	5x16+ML2 ^a
2x8 Riser	2x8	None	2x8	2x8	2x8

a. If desired, the 4x8 or 2x16 riser cards can be used instead of the 5x16+ML2 riser, with reduced slot availability

Based on these riser card requirements, the available slots for each server configuration are highlighted with a grey background in Table 3-28. The table also lists the processor that each PCIe slot is connected to, based on the server configuration selected. In the table, NC means not connected.

Processor numbering is shown in Figure 3-6 on page 51.

Table 3-28 Processor connectivity for PCIe slots (NC = not connected)

	Slot location	4S Performance	4S Performance/ Storage Rich	4S Upgradable	4S Upgradable/ Storage Rich	8S Storage Rich
Processors installed		Two (1,2) or Three (1,2,4) or Four (1,2,3,4)	Two (1,2) or Three (1,2,4) or Four (1,2,3,4)	Two (1,2) or Four (1,2,3,4)	Four (1,2,5,6)	Six (1,2,3,4,5,7) or eight (1-8)
Slot 1	Riser 1	4 ^a (NC for 2x16) ^b	4 ^a (NC for 2x16) ^b	4 ^a (NC for 2x16) ^b	NC	8 ^a
Slot 2		4 ^a (NC for 2x16) ^b	4 ^a (NC for 2x16) ^b	4 ^a (NC for 2x16) ^b	NC	8 ^a
Slot 3		4 ^a	4 ^a	4 ^a	6	6 ^a
Slot 4		4 ^a	4 ^a	4 ^a	6	6 ^a
Slot 5	I/O Tray	2	2	2	2	2
Slot 6	I/O Tray	2	2	2	2	2
Slot 7	I/O Tray	1	1	1	1	1
Slot 8 (ML2)	I/O Tray	1	1	1	1	1
Slot 9 (LOM)	I/O Tray	1 (PCH)	1 (PCH)	1 (PCH)	1 (PCH)	1 (PCH)
Slot 10	Riser 2	3 ^a	3 ^a	3 ^a	NC	7
Slot 11		3 ^a	3 ^a	3 ^a	NC	7
Slot 12		3 ^a	3 ^a	3 ^a	5	5
Slot 13		3 ^a	3 ^a	3 ^a	5	5
Slot 14		NC	NC	NC	NC	3
Slot 15 (ML2)		NC	NC	NC	NC	3
Slot 16	2x8 Riser	3 ^a	NC	3 ^a	NC	3
Slot 17		1	NC	1	1	1
M.2 slot	I/O Tray	1 (PCH)	1 (PCH)	1 (PCH)	1 (PCH)	1 (PCH)
Storage adapter slot	Upper tray	NC	1	NC	5	5
Storage adapter slot	Lower tray	1	1	1	1	1

a. See note below

b. Slots 1 and 2 are not connected if the 2x16 riser card is used in riser slot 1

Note: In 2, 3 and 6 processor configurations, the following slots are not available:

- ▶ In 2-processor configurations, processors 3 and 4 are not installed; this means slots 1-4, slots 10-13, and slot 16 are not connected
- ▶ In 3-processor configurations, processor 3 is not installed; this means slots 10-13, and slot 16 are not connected
- ▶ In 6-processor configurations, processors 6 & 8 are not installed; this means slots 1-4 are not connected

3.13 Network adapters

The SR950 has an integrated 10Gb Ethernet controller, based on the Intel Ethernet Connection X722. The X722 has the following features:

- ▶ Integrated into the Intel PCH chipset
- ▶ Supports 1GbE or 10GbE connections, depending on the selected LOM Phy adapter
- ▶ Offers VXLAN/NVGRE hardware offloads
- ▶ Supports VMDq and SR-IOV for advanced virtualization
- ▶ Supports iWarp RDMA

To provide connectivity to the X722, the SR950 supports six different LOM (LAN on Motherboard) cards which provide the Ethernet PHY function. Only 1 LOM card can be installed. See Figure 3-2 on page 49 for the location of the LOM slot.

The LOM card supports NC-SI to allow the network connection to be shared between the operating system and the XClarity Controller (XCC) management module. The LOM also supports Wake-on-LAN (WOL).

For more information about the X722 controller, see the Lenovo Press product guide:

<http://lenovopress.com/lp0654-intel-x722-integrated-controller>

Table 3-29 lists the supported LOM cards.

Note: None of the LOM cards listed in the table support speeds lower than 1 Gbps.

Table 3-29 Supported LOM adapters

Part number	Feature code	Description	Ports	Slots supported	Maximum supported
Gigabit Ethernet					
7ZT7A00544	AUKG	ThinkSystem 1Gb 2-port RJ45 LOM	2x RJ45	7	1
7ZT7A00545	AUKH	ThinkSystem 1Gb 4-port RJ45 LOM	4x RJ45	7	1
10 Gb Ethernet					
7ZT7A00546	AUKJ	ThinkSystem 10Gb 2-port SFP+ LOM	2x SFP+ bays	7	1
7ZT7A00547	AUKK	ThinkSystem 10Gb 4-port SFP+ LOM	4x SFP+ bays	7	1
7ZT7A00548	AUKL	ThinkSystem 10Gb 2-port Base-T LOM	2x RJ45 (10GBASE-T)	7	1
7ZT7A00549	AUKM	ThinkSystem 10Gb 4-port Base-T LOM	4x RJ45 (10GBASE-T)	7	1

The SR950 also support one ML2 card for 4-socket configurations (slot 8). For 8-socket servers using the 5x16 ML2 riser, the server has two ML2 slots - slot 8 and slot 15. See 3.12, "I/O expansion options" on page 78 for details.

NC-SI support: The ML2 card in slot 8 supports NC-SI however the ML2 card in slot 15 does not support NC-SI.

Table 3-30 lists the supported ML2 adapters.

Table 3-30 Supported ML2 adapters

Part number	Feature code	Description	Bus width	Slots supported	Maximum supported ^a
Gigabit Ethernet					
7ZT7A00536	AUKW	ThinkSystem Intel I350-T4 ML2 1Gb 4-Port RJ45 Ethernet Adapter	PCIe 3.0 x8	8, 15	1 / 2
10 Gb Ethernet					
00JY940	ATRH	Intel X710-DA2 ML2 2x10GbE SFP+ Adapter	PCIe 3.0 x8	8, 15	1 / 2
7ZT7A00497	AUKQ	ThinkSystem Broadcom 57416 10GBASE-T 2-Port ML2 Ethernet Adapter	PCIe 3.0 x8	8, 15	1 / 2
00AG560 ^b	AT7U	Emulex VFA5.2 ML2 Dual Port 10GbE SFP+ Adapter	PCIe 3.0 x8	8, 15	1 / 2
01CV770	AU7Z	Emulex VFA5.2 ML2 2x10 GbE SFP+ Adapter and FCoE/iSCSI SW	PCIe 3.0 x8	8, 15	1 / 2
25 Gb Ethernet					
7ZT7A00507	AUKU	ThinkSystem Mellanox ConnectX-4 Lx 10/25GbE SFP28 2-Port ML2 Ethernet Adapter	PCIe 3.0 x8	8, 15	1 / 2
00MN990	ATZR	Mellanox ConnectX-4 Lx 10/25GbE SFP28 1-port ML2 Adapter	PCIe 3.0 x8	8, 15	1 / 2
InfiniBand					
7ZT7A00501	AUKR	ThinkSystem Mellanox ConnectX-3 Pro ML2 FDR 2-Port QSFP VPI Adapter	PCIe 3.0 x8	8, 15	1 / 2

a. Two ML2 adapters are supported with an 8-socket configuration if the 5x16+ML2 riser is installed in riser slot 2.

b. ThinkSystem servers do not support Features on Demand, so the Emulex VFA5.2 ML2 Dual Port 10GbE SFP+ Adapter, 00AG560, cannot be upgraded to FCoE support. If you need FCoE or iSCSI support, use 01CV770 instead.

Table 3-31 lists additional supported network adapters that can be installed in the regular PCIe slots.

Table 3-31 Supported PCIe Network Adapters

Part number	Feature code	Description	Slots supported	Maximum supported
Gigabit Ethernet				
7ZT7A00482	AUZX	ThinkSystem Broadcom 5720 1GbE RJ45 2-Port PCIe Ethernet Adapter	1-7, 10-14, 16-17	14
7ZT7A00484	AUZV	ThinkSystem Broadcom 5719 1GbE RJ45 4-Port PCIe Ethernet Adapter	1-7, 10-14, 16-17	14
7ZT7A00533	AUZZ	ThinkSystem I350-F1 PCIe 1Gb 1-Port SFP Ethernet Adapter	1-7, 10-14, 16-17	14
7ZT7A00534	AUZY	ThinkSystem I350-T2 PCIe 1Gb 2-Port RJ45 Ethernet Adapter	1-7, 10-14, 16-17	14
7ZT7A00535	AUZW	ThinkSystem I350-T4 PCIe 1Gb 4-Port RJ45 Ethernet Adapter	1-7, 10-14, 16-17	14
10 Gb Ethernet - 10GBase-T				

Part number	Feature code	Description	Slots supported	Maximum supported
7ZT7A00496	AUKP	ThinkSystem Broadcom 57416 10GBASE-T 2-Port PCIe Ethernet Adapter	1-7, 10-14, 16-17	14
00MM860	ATPX	Intel X550-T2 Dual Port 10GBase-T Adapter	1-7, 10-14, 16-17	14
4XC7A08225	B31G	ThinkSystem QLogic QL41134 PCIe 10Gb 4-Port Base-T Ethernet Adapter	1-7, 10-14, 16-17	14
10 Gb Ethernet - SFP+				
7ZT7A00537	AUKX	ThinkSystem X710-DA2 PCIe 10Gb 2-Port SFP+ Ethernet Adapter	1-7, 10-14, 16-17	14
7ZT7A00493	AUKN	ThinkSystem Emulex OCE14104B-NX PCIe 10Gb 4-Port SFP+ Ethernet Adapter	1-4, 10-14	9
00AG570 ^a	AT7S	Emulex VFA5.2 2x10 GbE SFP+ PCIe Adapter	1-7, 10-14, 16-17	14
00AG580	AT7T	Emulex VFA5.2 2x10 GbE SFP+ Adapter and FCoE/iSCSI SW	1-7, 10-14, 16-17	14
25 Gb Ethernet				
4XC7A08228	B21R	ThinkSystem QLogic QL41262 10/25GbE SFP28 2-Port PCIe Ethernet Adapter	1-7, 10-14, 16-17	14
7XC7A05523	B0WY	ThinkSystem Intel XXV710-DA2 10/25GbE SFP28 2-Port PCIe Ethernet Adapter	1-7, 10-14, 16-17	14
7ZT7A00505	AUKS	ThinkSystem Broadcom 57412 25GbE SFP28 1-Port PCIe Ethernet Adapter	1-7, 10-14, 16-17	14
01GR250	AUAJ	Mellanox ConnectX-4 Lx 10/25GbE SFP28 2-port PCIe Ethernet Adapter	1-7, 10-14, 16-17	14
4XC7A08229	B31C	ThinkSystem Mellanox ConnectX-5 Ex 25/40GbE 2-port Low-Latency Adapter	1-6, 10-14	11
40 Gb Ethernet				
00MM950	ATRN	Mellanox ConnectX-4 Lx 1x40GbE QSFP+ Adapter	1-7, 10-14, 16-17	14
4XC7A08229	B31C	ThinkSystem Mellanox ConnectX-5 Ex 25/40GbE 2-port Low-Latency Adapter	1-6, 10-14	11
100 Gb Ethernet				
00KH924	ASWQ	Mellanox ConnectX-4 1x100GbE/EDR IB QSFP28 VPI Adapter	1-6, 10-14	11
00MM960	ATRP	Mellanox ConnectX-4 2x100GbE/EDR IB QSFP28 VPI Adapter	1-6, 10-14	11
InfiniBand				
7ZT7A00500	AUVG	ThinkSystem Mellanox ConnectX-4 PCIe FDR 2-Port QSFP VPI Adapter	1-7, 10-14, 16-17	14
Omni-Path Architecture				
00WE023	AU0A	Intel OPA 100 Series Single-port PCIe 3.0 x8 HFA	1-7, 10-14, 16-17	14
00WE027	AU0B	Intel OPA 100 Series Single-port PCIe 3.0 x16 HFA	1-6, 10-14	11

- a. ThinkSystem servers do not support Features on Demand, so the Emulex VFA5.2 2x10 GbE SFP+ PCIe Adapter, 00AG570, cannot be upgraded to FCoE support. If you need FCoE or iSCSI support, use 00AG580 instead

For more information, including the transceivers and cables that each adapter supports, see the list of Lenovo Press Product Guides in the Networking adapters category:

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

3.14 Fibre Channel host bus adapters

Table 3-32 lists the Fibre Channel HBAs supported by SR950.

Table 3-32 Fibre Channel HBAs

Part number	Feature code	Description	Slots supported	Maximum supported
32 Gb Fibre Channel HBAs				
7ZT7A00516	AUNS	ThinkSystem QLogic QLE2740 PCIe 32Gb 1-Port SFP+ Fibre Channel Adapter	1-7, 10-14, 16-17	14
7ZT7A00518	AUNU	ThinkSystem QLogic QLE2742 PCIe 32Gb 2-Port SFP+ Fibre Channel Adapter	1-7, 10-14, 16-17	14
7ZT7A00519	AUNV	ThinkSystem Emulex LPe32002-M2-L PCIe 32Gb 2-Port SFP+ Fibre Channel Adapter	1-7, 10-14, 16-17	14
7ZT7A00517	AUNT	ThinkSystem Emulex LPe32000-M2-L PCIe 32Gb 1-Port SFP+ Fibre Channel Adapter	1-7, 10-14, 16-17	14
16 Gb Fibre Channel HBAs				
01CV750	ATZB	QLogic 16Gb Enhanced Gen5 FC Single-port HBA	1-7, 10-14, 16-17	14
01CV760	ATZC	QLogic 16Gb Enhanced Gen5 FC Dual-port HBA	1-7, 10-14, 16-17	14
01CV830	ATZU	Emulex 16Gb Gen6 FC Single-port HBA	1-7, 10-14, 16-17	14
01CV840	ATZV	Emulex 16Gb Gen6 FC Dual-port HBA	1-7, 10-14, 16-17	14

For more information, see the list of Lenovo Press Product Guides in the Host bus adapters category: <http://lenovopress.com/servers/options/hba>

3.15 SAS adapters for external storage

Table 3-33 lists SAS HBAs and RAID adapters supported by SR950 server for use with external storage.

Table 3-33 Adapters for external storage

Part number	Feature code	Description	Slots supported	Maximum supported
SAS HBAs				
7Y37A01090	AUNR	ThinkSystem 430-8e SAS/SATA 12Gb HBA	1-7, 10-14, 16-17	14
7Y37A01091	AUNN	ThinkSystem 430-16e SAS/SATA 12Gb HBA	1-7, 10-14, 16-17	14

Part number	Feature code	Description	Slots supported	Maximum supported
External RAID adapters				
7Y37A01087	AUNQ	ThinkSystem RAID 930-8e 4GB Flash PCIe 12Gb Adapter	1-7, 10-14, 16-17	Four ^a

a. Maximum is 4 adapters because there are only 4 supercap holders in the server, mounted on the side of the two riser cards.

For more information, see the list of Lenovo Press Product Guides in the Host bus adapters and RAID adapters categories:

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

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

Table 3-34 compares the specifications of the external SAS HBAs and RAID adapters.

Table 3-34 Comparison of external storage adapters

Feature	430-8e	430-16e	930-8e
Adapter type	HBA	HBA	External RAID
Part number	7Y37A01090	7Y37A01091	7Y37A01087
Controller chip	LSI SAS3408	LSI SAS3416	LSI SAS3508
Broadcom equivalent	HBA 9400-8e	HBA 9400-16e	MegaRAID 9480-8e
Host interface	PCIe 3.0x8	PCIe 3.0x8	PCIe 3.0x8
Port interface	12 Gb SAS	12 Gb SAS	12 Gb SAS
Number of ports	8	16	8
Port connectors	2x Mini-SAS HD SFF8644	4x Mini-SAS HD SFF8644	2x Mini-SAS HD SFF8644
Drive interface	SAS/SATA	SAS/SATA	SAS/SATA
Drive type	HDD/SSD/SED ^a	HDD/SSD/SED ^a	HDD,SED,SSD
Hot-swap drives	Yes	Yes	Yes
Maximum devices	512 (planned: 1024)	512 (planned: 1024)	64 (planned: 216)
RAID levels	None	None	0/1/10/5/50/6/60
JBOD mode	Yes	Yes	Yes
Cache	None	None	4GB (Standard)
CacheVault cache protection	None	None	Yes (Flash)
Performance Accelerator (FastPath)	No	No	Yes
SSD Caching (CacheCade Pro 2.0)	No	No	No
SED support*	Yes	Yes	Yes

a. SED support of the SAS HBAs is by using software on the server (SED commands are passed through the HBA to the drives). SED support by RAID controllers is provided using the built-in MegaRAID SafeStore functionality of the adapter.

3.16 Flash Storage adapters

The server supports the PCIe Flash Storage adapters listed in Table 3-35.

Table 3-35 Flash storage adapters

Part number	Feature code	Description	Slots supported	Maximum supported
Mainstream NVMe PCIe Adapters - Optimized for mixed-intensive application workloads with an endurance of 3-5 DWPD.				
4XB7A08520	B32L	ThinkSystem HHHH KCM51V 1.6TB Mainstream NVMe PCIe 3.0 x4 Flash Adapter	1-7, 10-14, 16-17	14
4XB7A08521	B32M	ThinkSystem HHHH KCM51V 3.2TB Mainstream NVMe PCIe 3.0 x4 Flash Adapter	1-7, 10-14, 16-17	14
4XB7A08522	B32N	ThinkSystem HHHH KCM51V 6.4TB Mainstream NVMe PCIe 3.0 x4 Flash Adapter	1-7, 10-14, 16-17	14
7SD7A05769	B11X	ThinkSystem HHHH Intel P4600 2.0TB Mainstream NVMe PCIe3.0 x4 Flash Adapter	1-7, 10-14, 16-17	14
7SD7A05768	B11Y	ThinkSystem HHHH Intel P4600 4.0TB Mainstream NVMe PCIe3.0 x4 Flash Adapter	1-7, 10-14, 16-17	14
7N47A00097	AUUP	ThinkSystem HHHH PX04PMC 1.92TB Mainstream NVMe PCIe 3.0 x4 Flash Adapter	1-7, 10-14, 16-17	14
7N47A00098	AUVY	ThinkSystem HHHH PX04PMC 3.84TB Mainstream NVMe PCIe 3.0 x4 Flash Adapter	1-7, 10-14, 16-17	14
Performance NVMe PCIe Adapters - Optimized for write-intensive application workloads with an endurance of 10+ DWPD.				
7XB7A05925	AWG8	ThinkSystem HHHH PX04PMC 1.6TB Performance NVMe PCIe 3.0 x4 Flash Adapter	1-7, 10-14, 16-17	14
7XB7A05924	AWG9	ThinkSystem HHHH PX04PMC 3.2TB Performance NVMe PCIe 3.0 x4 Flash Adapter	1-7, 10-14, 16-17	14

For details about these adapters, see the Lenovo Press product guides in the Flash Adapters category: <http://lenovopress.com/servers/options/ssdadapter>

3.17 GPU adapters and coprocessors

Full-length adapter cards such as graphics processing units (GPUs) do not physically fit in the server and are not supported.

3.18 Cooling

The SR950 uses a set of 60mm hot-swap single-rotor fans for system-wide cooling. Each compute tray or storage tray houses 6 hot-swap fans. The 6 fans in the compute tray are N+1 redundant, meaning that a single fan failure is tolerated in each compute tray.

Fans are located internal to the server immediately in front of the compute system boards, but are accessible from the front of the server by the long handle attached to the fan. As the

handle is pulled out, the attached fan rotates from vertical orientation to horizontal orientation, allowing the fan to be removed from the system using the handle, shown in Figure 3-18.




Figure 3-18 Six fans installed in each compute tray (lower compute tray shown)

3.19 Power supplies

The server supports up to four hot-swap power supplies. Power supplies are N+N redundant, which means that if the server has four power supplies installed, then the server can continue operation even with two failed power supplies, or if the utility service that feeds those two power supplies goes offline.

All power supplies installed in a server must have the same capacity (for example all 1600W).

Table 3-36 Power supply options

Part number	Feature code	Description	Maximum supported	110V AC	220V AC	240V DC China Only
7N67A00885	AVNV	1100W Platinum (230V/115V) Power Supply	4	Yes	Yes	No
7N67A00886	AUPT	1600W Platinum (230V) Power Supply	4	No	Yes	No

The 1100W power supply is auto-sensing and supports both 110V AC (100-127V 50/60 Hz) and 220V AC (200-240V 50/60 Hz) power. The 1600 W power supply only supports 220V AC power. 240V DC is not supported.

Note: At 100-127 V AC input, the 1100W power supply is limited to 1050W.

Power supply options do not include a power cord. For models of the SR950, the inclusion of a power cords is model dependent. Configure-to-order models can be configured without a power cord if desired. For a list of power cords for the system, refer to 4.11.1, “Server-to-PDU power cord options” on page 114.

Use the Lenovo Capacity Planner to determine exactly what power your server needs:

<https://datacentersupport.lenovo.com/us/en/products/solutions-and-software/software/lenovo-capacity-planner/solutions/HT504651>

3.20 Integrated virtualization

The server supports booting from an operating system or hypervisor installed on an M.2 solid-state drive. See 3.8.3, “M.2 drives” on page 70 for details and the list of available options.

You can download supported VMware vSphere hypervisor images from the following web page and load it on the M.2 drive using the instructions provided:

https://vmware.lenovo.com/content/custom_iso/6.5/

3.21 Systems management

The server contains an integrated service processor, XClarity Controller (XCC), which provides advanced service-processor control, monitoring, and alerting functions. The XCC is based on the Pilot4 XE401 baseboard management controller (BMC) using a dual-core ARM Cortex A9 service processor.

This section describes the hardware functions that the server provides for systems management. For additional systems management capabilities, see Chapter 7, “Managing the hardware” on page 223.

3.21.1 Local management

The SR950 offers the front operator panel showing key LED status indicators, as shown in Figure 3-19.




Figure 3-19 Front operator panel (LCD display panel closed)

The front operator panel pulls out (see the release latch in Figure 3-19) to show an LCD display panel and buttons, as shown in Figure 3-20 on page 91. The LCD display panel allows quick access to system status, firmware, network, and health information.



Figure 3-20 LCD display panel

The LCD display and the function buttons give you access to the following information:

- ▶ Error messages
- ▶ System VPD: machine type & mode, serial number, UUID string
- ▶ System firmware levels: UEFI and XCC firmware
- ▶ XCC network information: hostname, MAC address, IP address, DNS addresses
- ▶ Environmental data: Ambient temperature, CPU temperature, AC input voltage, estimated power consumption
- ▶ System reset action

If an environmental condition exceeds a threshold or if a system component fails, XCC lights LEDs inside the server to help you diagnose the problem and find the failing part. The SR950 has fault LEDs next to the following components:

- ▶ Each processor
- ▶ Each memory DIMM
- ▶ Each system fan
- ▶ Each power supply
- ▶ Each rear-facing PCIe slot
- ▶ Each drive bay

For local console use, the SR950 has both a front and rear VGA port. Only one can be active at a time. If displays are connected to both VGA ports, the front VGA port takes precedence.

3.21.2 System status with XClarity Mobile

The XClarity Mobile app now includes a tethering function where you can connect your Android or iOS device to the server via USB to see the status of the server.

The steps to connect the mobile device are as follows:

1. Enable USB Management on the server, by holding down the ID button for 3 seconds (or pressing the dedicated USB management button if one is present)
2. Connect the mobile device via a USB cable to the server's USB port with the management symbol USB Management symbol
3. In iOS or Android settings, enable Personal Hotspot or USB Tethering

4. Launch the Lenovo XClarity Mobile app

Once connected you can see the following information:

- ▶ Server status including error logs (read only, no login required)
- ▶ Server management functions (XClarity login credentials required)

For additional information on XClarity Mobile, refer to 5.6, “Tethering with XClarity Mobile” on page 178.

3.21.3 Remote management

Remote server management is provided through industry-standard interfaces:

- ▶ Intelligent Platform Management Interface (IPMI) Version 2.0
- ▶ Simple Network Management Protocol (SNMP) Version 3 (no SET commands; no SMNP v1)
- ▶ Common Information Model (CIM-XML)
- ▶ Representational State Transfer (REST) support
- ▶ Redfish support (DMTF compliant)
- ▶ Web browser - HTML 5-based browser interface (Java and ActiveX not required) using a responsive design (content optimized for the device being used: laptop, tablet, phone) with NLS support

A virtual presence (remote control and remote media) capability also comes standard in the server for remote server management.

The remote control functions include the following:

- ▶ Remotely viewing video with graphics resolutions up to 1600x1200 at 75 Hz with up to 23 bits per pixel, regardless of the system state
- ▶ Remotely accessing the server using the keyboard and mouse from a remote client
- ▶ Capturing blue-screen errors
- ▶ International keyboard mapping support
- ▶ LDAP-based authentication
- ▶ Optionally, with the XCC Enterprise license upgrade, mapping ISO and diskette IMG image files as virtual drives that are available for use by the server
- ▶ Boot Capture
- ▶ Remote mounting of ISO and IMG files
- ▶ Virtual console collaboration
- ▶ Power capping
- ▶ License for XClarity Energy Manager

For additional information on using the XCC remote function, refer to 5.2, “Preparing the hardware with XClarity Controller” on page 142.

3.21.4 Lenovo XClarity Administrator

Lenovo XClarity Administrator is a centralized resource management solution designed to reduce complexity, speed response, and enhance the availability of Lenovo systems and solutions.

Lenovo XClarity Administrator provides agent-free hardware management for ThinkSystem servers, in addition to ThinkServer®, System x, and Flex System servers. The administration dashboard is based on HTML 5 and allows fast location of resources so tasks can be run quickly.

Because Lenovo XClarity Administrator does not require any agent software to be installed on the managed endpoints, there are no CPU cycles spent on agent execution, and no memory is used, which means that up to 1GB of RAM and 1 - 2% CPU usage is saved, compared to a typical managed system where an agent is required.

Lenovo XClarity Administrator provides full management function to ThinkSystem servers, including the following:

- ▶ Discovery
- ▶ Inventory
- ▶ Monitoring and alerting
- ▶ Call home
- ▶ Centralized user management
- ▶ Cryptography modes, server certificates, and encapsulation
- ▶ Configuration patterns
- ▶ Operating system deployment
- ▶ Firmware updates

For more information about Lenovo XClarity Administrator, including ordering part numbers, see the Lenovo XClarity Administrator Product Guide:

<http://lenovopress.com/tips1200-lenovo-xclarity-administrator>

3.21.5 Lenovo XClarity Essentials

Lenovo offers the following XClarity Essentials software tools that can help you set up, use, and maintain the server at no additional cost:

- ▶ Lenovo Essentials OneCLI

OneCLI is a collection of server management tools that uses a command line interface program to manage firmware, hardware, and operating systems. It provides functions to collect full system health information (including health status), configure system setting, and update system firmware and drivers.

- ▶ Lenovo Essentials UpdateXpress

The UpdateXpress tool is a standalone GUI application for firmware and device driver updates that enables you to maintain your server firmware and device drivers up-to-date and help you avoid unnecessary server outages. The tool acquires and deploys individual updates and UpdateXpress System Packs™ (UXSPs) which are integration-tested bundles.

- ▶ Lenovo Essentials Bootable Media Creator™

The Bootable Media Creator (BOMC) tool is used to create bootable media for offline firmware update.

For more information and downloads, visit the Lenovo XClarity Essentials web page:

<http://support.lenovo.com/us/en/documents/LNV0-center>

3.21.6 Lenovo XClarity Energy Manager

Lenovo XClarity Energy Manager is an agent-free, web-based console that provides power management for ThinkSystem servers as well as System x and ThinkServer systems. It enables server density and data center capacity to be increased through the use of power capping.

Lenovo XClarity Energy Manager is a licensed product. A single-node XClarity Energy Manager license is included with the XClarity Controller Enterprise (XCC Enterprise) upgrade as described in the 3.21.3, "Remote management" on page 92. If your server does not have the XCC Enterprise upgrade, Energy Manager licenses can be ordered as shown in Table 3-37.

Table 3-37 Lenovo XClarity Energy Manager

Description	Part number	Maximum supported
4L40E51621	Lenovo XClarity Energy Manager Node License	1 node
4L40E51622	Lenovo XClarity Energy Manager Node License Pack	5 nodes
4L40E51623	Lenovo XClarity Energy Manager Node License Pack	50 nodes

3.22 Security

The server offers the following security features:

- ▶ Administrator and power-on password
- ▶ Trusted Platform Module (TPM) supporting both TPM 1.2 and TPM 2.0
- ▶ Optional Trusted Cryptographic Module (TCM), available only in China

The TCM module, available only for China customers, is installed in a dedicated socket on the I/O tray, as shown in Figure 3-8 on page 52. Ordering information is shown in Table 3-38.

Table 3-38 Security features

Part number	Feature code	Description
None ^a	AVKE	ThinkSystem Trusted Cryptographic Module (China customers only)

a. The Trusted Cryptographic Module (TCM) for China customers is not available as a field upgrade. The component is CTO on pre-configured models only.

3.23 Operating system support

The SR950 supports the following operating systems

- ▶ Microsoft Windows Server 2012 R2
- ▶ Microsoft Windows Server 2016
- ▶ Microsoft Windows Server, version 1709
- ▶ Microsoft Windows Server, version 1803

- ▶ Red Hat Enterprise Linux 6.10 x64
- ▶ Red Hat Enterprise Linux 6.9 x64
- ▶ Red Hat Enterprise Linux 7.3
- ▶ Red Hat Enterprise Linux 7.4
- ▶ Red Hat Enterprise Linux 7.5
- ▶ SUSE Linux Enterprise Server 11 x64 SP4
- ▶ SUSE Linux Enterprise Server 12 SP2
- ▶ SUSE Linux Enterprise Server 12 SP3
- ▶ SUSE Linux Enterprise Server 15
- ▶ SUSE Linux Enterprise Server 15 Xen
- ▶ VMware ESXi 6.0 U3
- ▶ VMware ESXi 6.5
- ▶ VMware ESXi 6.5 U1
- ▶ VMware ESXi 6.5 U2
- ▶ VMware ESXi 6.7

Note: Support for RHEL 6.9 on the SR950 is limited to up to 4 processors. Configurations greater than 4 processors are not supported.

For the latest information about the specific versions and service levels that are supported and any other prerequisites, see the Operating System Interoperability Guide:

<https://lenovopress.com/osig#servers=sr950-7x11-7x12-7x13>

For configure-to-order configurations, the server can be preloaded with VMware ESXi installed on M.2 cards. Ordering information is listed in Table 3-39.

Table 3-39 VMware ESXi preload

Part number	Feature code	Description
CTO only	AXFT	VMware ESXi 6.5 (factory installed)
CTO only	AXFS	VMware ESXi 6.0 U3 (factory installed)

Infrastructure planning

The Lenovo ThinkSystem SR950 is an enterprise-class Intel processor-based servers for mission-critical applications. This server 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 98
- ▶ 4.2, “Operating environment” on page 99
- ▶ 4.3, “Rack selection and rack options” on page 101
- ▶ 4.4, “Rack installation” on page 103
- ▶ 4.5, “Floor clearance” on page 104
- ▶ 4.6, “Using Rear Door Heat eXchanger (RDHX)” on page 105
- ▶ 4.7, “Power guidelines” on page 107
- ▶ 4.8, “Power monitoring” on page 109
- ▶ 4.9, “Cooling considerations” on page 112
- ▶ 4.10, “Uninterruptible power supply units” on page 113
- ▶ 4.11, “PDU and line cord selection” on page 114

4.1 Physical and electrical specifications

The SR950 physical and electrical specifications are listed below.

Dimensions and weight:

- ▶ Height: 173 mm (6.8 in.)
- ▶ Depth:
 - EIA flange to rear: 800.0 mm (31.5 in.)
 - Overall: 835.0 mm (32.9 in.)
- ▶ Width:
 - Top cover: 447.0 mm (17.6 in.)
 - EIA flange: 482.0 mm (19.0 in.)
- ▶ Weight:
 - Minimum configuration: 32.6 kg (71.9 lb.)
 - Maximum configuration: 58.7 kg (129.4 lb.)

Electrical:

- ▶ Models with 1600 W AC power supplies:
 - 200 - 240 (nominal) V AC; 50 Hz or 60 Hz; 9.0 A
 - Input kilovolt-amperes (kVA) (approximately):
 - Minimum configuration: 0.16 kVA
 - Maximum configuration: 7.2 kVA
- ▶ Models with 1100 W AC power supplies:
 - 100 - 127 (nominal) V AC; 50 Hz or 60 Hz; 11.9 A
 - 200 - 240 (nominal) V AC; 50 Hz or 60 Hz; 6.2 A
 - Input kilovolt-amperes (kVA) (approximately):
 - Minimum configuration: 0.12 kVA
 - Maximum configuration: 4.9 kVA

The server is supported operating in the following environments:

- ▶ ASHRAE A2:
 - Temperature: 10°C to 35°C
 - Humidity: 20% to 80%
 - Altitude 0-3000 m (10,000 ft), derated 1°C per 300 m above 950 m
- ▶ ASHRAE A3:
 - Temperature: 5°C to 40°C
 - Humidity: 8% to 85%
 - Altitude 0-3000 m (10,000 ft), derated 1°C per 175 m above 950 m
- ▶ ASHRAE A4:
 - Temperature: 5°C to 45°C
 - Humidity: 8% to 90%
 - Altitude 0-3000 m (10,000 ft), derated 1°C per 125 m above 950 m

Non-operating environment support:

- ▶ Without packaging:
 - Temperature: -10 °C to 60 °C
 - Humidity: 8 to 90%

- Storage with packaging:
 - Temperature: -40 °C to 70 °C
 - Humidity: 8 to 90%

The server generates the following heat:

- Thermal/heat output:
 - Minimum configuration: 938 BTU/hr, 275 W
 - Typical configuration: 3412 BTU/hr, 1000 W
 - Full configuration: 11,970 BTU/hr, 3508 W

4.2 Operating environment

The ThinkSystem SR950 server is designed to ensure thermal compliance to ASHRAE A4 ambient environments, however certain processor models and other components will limit support to only ASHRAE A2 conditions, as shown in Table 4-1. Further, some high-end processors (as indicated with an * in the table) when used in an eight-socket configuration may experience a slight drop in performance under extreme workloads when the ambient temperature is above 30°C.

The following restrictions also apply regarding ASHRAE A3 and A4 support:

- NVMe drives and adapters are supported up to 35°C only

Table 4-1 ASHRAE support by processor

Processor	ASHRAE A2 support	ASHRAE A3 support	ASHRAE A4 support
Intel Xeon Gold 5115 10C 85W 2.4GHz Processor	Yes	Yes	Yes
Intel Xeon Gold 5118 12C 105W 2.3GHz Processor	Yes	Yes	Yes
Intel Xeon Gold 5120 14C 105W 2.2GHz Processor	Yes	Yes	Yes
Intel Xeon Gold 5120T 14C 105W 2.2GHz Processor	Yes	Yes	Yes
Intel Xeon Gold 5122 4C 105W 3.6GHz Processor	Yes	Yes	Yes
Intel Xeon Gold 6126 12C 125W 2.6GHz Processor	Yes	Yes	Yes
Intel Xeon Gold 6126T 12C 125W 2.6GHz Processor	Yes	Yes	Yes
Intel Xeon Gold 6128 6C 115W 3.4GHz Processor	Yes	Yes	Yes
Intel Xeon Gold 6130 16C 125W 2.1GHz Processor	Yes	Yes	Yes
Intel Xeon Gold 6130T 16C 125W 2.1GHz Processor	Yes	Yes	Yes
Intel Xeon Gold 6132 14C 140W 2.6GHz Processor	Yes	Yes	Yes
Intel Xeon Gold 6134 8C 130W 3.2GHz Processor	Yes	Yes	Yes
Intel Xeon Gold 6134M 8C 130W 3.2GHz Processor	Yes	Yes	Yes
Intel Xeon Gold 6136 12C 150W 3.0GHz Processor	Yes	No	No
Intel Xeon Gold 6138 20C 125W 2.0GHz Processor	Yes	Yes	Yes
Intel Xeon Gold 6138T 20C 125W 2.0GHz Processor	Yes	Yes	Yes
Intel Xeon Gold 6140 18C 140W 2.3GHz Processor	Yes	Yes	Yes

Processor	ASHRAE A2 support	ASHRAE A3 support	ASHRAE A4 support
Intel Xeon Gold 6140M 18C 140W 2.3GHz Processor	Yes	Yes	Yes
Intel Xeon Gold 6142 16C 150W 2.6GHz Processor	Yes	No	No
Intel Xeon Gold 6142M 16C 150W 2.6GHz Processor	Yes	No	No
Intel Xeon Gold 6148 20C 150W 2.4GHz Processor	Yes	No	No
Intel Xeon Gold 6150 18C 165W 2.7GHz Processor	Yes	No	No
Intel Xeon Gold 6152 22C 140W 2.1GHz Processor	Yes	Yes	Yes
Intel Xeon Gold 6154 18C 200W 3.0GHz Processor	Yes		
Intel Xeon Platinum 8153 16C 125W 2.0GHz Processor	Yes	Yes	Yes
Intel Xeon Platinum 8156 4C 105W 3.6GHz Processor	Yes	Yes	Yes
Intel Xeon Platinum 8158 12C 150W 3.0GHz Processor	Yes	No	No
Intel Xeon Platinum 8160 24C 150W 2.1GHz Processor	Yes	No	No
Intel Xeon Platinum 8160M 24C 150W 2.1GHz Processor	Yes	No	No
Intel Xeon Platinum 8160T 24C 150W 2.1GHz Processor	Yes	No	No
Intel Xeon Platinum 8164 26C 150W 2.0GHz Processor	Yes	No	No
Intel Xeon Platinum 8168 24C 205W 2.7GHz Processor	Yes ^a	No	No
Intel Xeon Platinum 8170 26C 165W 2.1GHz Processor	Yes	No	No
Intel Xeon Platinum 8170M 26C 165W 2.1GHz Processor	Yes	No	No
Intel Xeon Platinum 8176 28C 165W 2.1GHz Processor	Yes	No	No
Intel Xeon Platinum 8176M 28C 165W 2.1GHz Processor	Yes	No	No
Intel Xeon Platinum 8180 28C 205W 2.5GHz Processor	Yes ^a	No	No
Intel Xeon Platinum 8180M 28C 205W 2.5GHz Processor	Yes ^a	No	No

a. These processors when used in an eight-socket configuration may experience a slight drop in performance under extreme workloads when the ambient temperature is above 30°C.

The server is supported operating in the following environments:

► ASHRAE A2:

- Temperature: 10°C to 35°C
- Humidity: 20% to 80%
- Altitude 0-3000 m (10,000 ft), derated 1°C per 300 m above 950 m

► ASHRAE A3:

- Temperature: 5°C to 40°C
- Humidity: 8% to 85%
- Altitude 0-3000 m (10,000 ft), derated 1°C per 175 m above 950 m

► ASHRAE A4:

- Temperature: 5°C to 45°C
- Humidity: 8% to 90%
- Altitude 0-3000 m (10,000 ft), derated 1°C per 125 m above 950 m

Non-operating environment support:

- ▶ Without packaging:
 - Temperature: -10 °C to 60 °C
 - Humidity: 8 to 90%
- ▶ Storage with packaging:
 - Temperature: -40 °C to 70 °C
 - Humidity: 8 to 90%

The server generates the following heat:

- ▶ Thermal/heat output:
 - Minimum configuration: 938 BTU/hr, 275 W
 - Typical configuration: 3412 BTU/hr, 1000 W
 - Full configuration: 11,970 BTU/hr, 3508 W

4.3 Rack selection and rack options

The racks supported by the SR950 server are listed in Table 4-2.

Table 4-2 Rack cabinets

Part number	Description	Maximum supported
Rack cabinets currently available		
9308-4EX	42U Enterprise Expansion Rack	10
9308-4PX	42U Enterprise Rack	10
9360-4PX	42U 1200mm Deep Dynamic Rack	10
9361-4PX	42U 1200mm Deep Static Rack	10
9363-4PX	42U 1100mm Dynamic Rack	10
9363-4EX	42U 1100mm Dynamic Expansion Rack	10
Withdrawn rack cabinets		
9360-4EX	42U 1200mm Deep Dynamic Expansion Rack	10
9361-4EX	42U 1200mm Deep Static Expansion Rack	10
9362-4PX	47U 1200mm Deep Static Rack	11
9362-4EX	47U 1200mm Deep Static Expansion Rack	11
9363-4CX	PureFlex System 42U Rack	10
9363-4DX	PureFlex System 42U Expansion Rack	10

For information about these racks, see the Lenovo Rack Cabinet Reference:

<http://lenovopress.com/lp0658-lenovo-rack-cabinet-reference>

Table 4-3 on page 102 lists the supported KVM consoles, keyboards, and KVM switches.

Table 4-3 Console keyboards

Part number	Description
Consoles	
17238BX	1U 18.5" Standard Console (without keyboard)
Console keyboards	
7ZB7A05469	ThinkSystem Keyboard w/ Int. Pointing Device USB - Arabic 253 RoHS v2
7ZB7A05468	ThinkSystem Keyboard w/ Int. Pointing Device USB - Belg/UK 120 RoHS v2
7ZB7A05206	ThinkSystem Keyboard w/ Int. Pointing Device USB - Czech 489 RoHS v2
7ZB7A05207	ThinkSystem Keyboard w/ Int. Pointing Device USB - Danish 159 RoHS v2
7ZB7A05208	ThinkSystem Keyboard w/ Int. Pointing Device USB - Dutch 143 RoHS v2
7ZB7A05210	ThinkSystem Keyboard w/ Int. Pointing Device USB - Fr/Canada 445 RoHS v2
7ZB7A05209	ThinkSystem Keyboard w/ Int. Pointing Device USB - French 189 RoHS v2
7ZB7A05211	ThinkSystem Keyboard w/ Int. Pointing Device USB - German 129 RoHS v2
7ZB7A05212	ThinkSystem Keyboard w/ Int. Pointing Device USB - Greek 219 RoHS v2
7ZB7A05213	ThinkSystem Keyboard w/ Int. Pointing Device USB - Hebrew 212 RoHS v2
7ZB7A05214	ThinkSystem Keyboard w/ Int. Pointing Device USB - Hungarian 208 RoHS v2
7ZB7A05215	ThinkSystem Keyboard w/ Int. Pointing Device USB - Italian 141 RoHS v2
7ZB7A05216	ThinkSystem Keyboard w/ Int. Pointing Device USB - Japanese 194 RoHS v2
7ZB7A05217	ThinkSystem Keyboard w/ Int. Pointing Device USB - Korean 413 RoHS v2
7ZB7A05218	ThinkSystem Keyboard w/ Int. Pointing Device USB - LA Span 171 RoHS v2
7ZB7A05219	ThinkSystem Keyboard w/ Int. Pointing Device USB - Norwegian 155 RoHS v2
7ZB7A05220	ThinkSystem Keyboard w/ Int. Pointing Device USB - Polish 214 RoHS v2
7ZB7A05221	ThinkSystem Keyboard w/ Int. Pointing Device USB - Portugese 163 RoHS v2
7ZB7A05222	ThinkSystem Keyboard w/ Int. Pointing Device USB - Russian 441 RoHS v2
7ZB7A05223	ThinkSystem Keyboard w/ Int. Pointing Device USB - Slovak 245 RoHS v2
7ZB7A05231	ThinkSystem Keyboard w/ Int. Pointing Device USB - Slovenian 234 RoHS v2
7ZB7A05224	ThinkSystem Keyboard w/ Int. Pointing Device USB - Spanish 172 RoHS v2
7ZB7A05225	ThinkSystem Keyboard w/ Int. Pointing Device USB - Swed/Finn 153 RoHS v2
7ZB7A05226	ThinkSystem Keyboard w/ Int. Pointing Device USB - Swiss F/G 150 RoHS v2
7ZB7A05227	ThinkSystem Keyboard w/ Int. Pointing Device USB - Thai 191 RoHS v2
7ZB7A05467	ThinkSystem Keyboard w/ Int. Pointing Device USB - Trad Chinese/US 467 RoHS v2
7ZB7A05228	ThinkSystem Keyboard w/ Int. Pointing Device USB - Turkish 179 RoHS v2
7ZB7A05229	ThinkSystem Keyboard w/ Int. Pointing Device USB - UK Eng 166 RoHS v2
7ZB7A05470	ThinkSystem Keyboard w/ Int. Pointing Device USB - US Eng 103P RoHS v2
7ZB7A05230	ThinkSystem Keyboard w/ Int. Pointing Device USB - US Euro 103P RoHS v2

Part number	Description
Console switches	
1754D2X	Global 4x2x32 Console Manager (GCM32)
1754D1X	Global 2x2x16 Console Manager (GCM16)
1754A2X	Local 2x16 Console Manager (LCM16)
1754A1X	Local 1x8 Console Manager (LCM8)
Console switch cables	
43V6147	Single Cable USB Conversion Option (UCO)
39M2895	USB Conversion Option (4 Pack UCO)
46M5383	Virtual Media Conversion Option Gen2 (VCO2)
46M5382	Serial Conversion Option (SCO)

For more information, see the list of Product Guides in the KVM Switches and Consoles category:

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

4.4 Rack installation

All models of the SR950 include a rail kit for installation in a 19-inch rack cabinet. The contents of the kits are shown in Figure 4-1.




Figure 4-1 SR950 rail kit

Optionally available is the SR950 shipping bracket kit (as shown in Figure 4-2 on page 104), which is used when the server is installed in a rack and the complete rack is then shipped to the client.




Figure 4-2 SR950 shipping bracket kit

Table 4-4 Shipping bracket kit

Part number	Feature code	Description
7XF7A03967	AUL5	ThinkSystem SR950 Shipping Bracket

4.4.1 Installing in a non-Lenovo rack

The SR950 server can be installed in non Lenovo racks. The 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).

The rack that the server is installed in should meet the following requirements:

- ▶ Minimum depth of 70 mm (2.76 in.) between the front mounting flange and inside of the front door.
- ▶ Minimum depth of 150 mm (5.9 in.) between the rear mounting flange and inside of the rear door.
- ▶ Minimum depth of 711.2 mm (28 in.) and maximum depth of 914.4 mm (36 in.) between the front and rear mounting flanges.
- ▶ Install the server only in a rack that has perforated front and rear doors or in a rack that is equipped with a rear door heat exchanger (RDHX) for cooling the rack with water. See 4.6, “Using Rear Door Heat eXchanger (RDHX)” on page 105 for additional information on the RDHX.

4.5 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 *ThinkSystem SR950 Rack Installation Instructions*.

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 server 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 to 8S by adding Compute boards and Compute trays without the need to upgrade the chassis itself.
- ▶ The system is easier to add I/O by hot-swapping and hot-adding riser cards to the I/O tray.
- ▶ The system is easier to add memory by removing the appropriate Compute tray and installing memory on the Compute board(s).

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 tray, I/O tray, fans or power supplies.

4.6 Using Rear Door Heat eXchanger (RDHX)

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 4-3 shows the capability of the RDHX.




Figure 4-3 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 4-5.

Table 4-5 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 web page:

<https://support.lenovo.com/docs/UM103398>

4.7 Power guidelines

In this section, we provide power planning considerations, describe the power subsystem in the SR950, 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:

- ▶ 4.7.1, “Considerations” on page 107
- ▶ 4.7.2, “Power supply redundancy” on page 108
- ▶ 4.7.3, “Power supply installation order” on page 109

4.7.1 Considerations

When you are planning your power source for an SR950 system rack installation, consider the following variables:

- ▶ Power supplies: The server supports up to four hot-swap power supplies: 1100W or 1600W
- ▶ Mixing power supplies: All power supplies installed in the server must be identical. Mixing of the 1100W PSU, 1600W PSU and voltage is not supported
- ▶ Input voltage range:
 - 100 - 120 VAC is supported with 1100W and 1600W power supplies
 - 200 - 240 VAC is supported with the 1600W power supplies only
- ▶ Power Distribution Unit (PDU) input: Single-phase or three-phase
- ▶ Power redundancies: AC source feed (power feed) N+N 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 Lenovo Capacity Planner. 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:

<https://datacentersupport.lenovo.com/us/en/products/solutions-and-software/software/lenovo-capacity-planner/solutions/HT504651>

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 UPS units for this system, email the Lenovo power team at power@lenovo.com.

4.7.2 Power supply redundancy

The four SR950 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 bottom to top when viewed from the rear of the chassis, as shown in Figure 4-4.




Figure 4-4 Grouping of the power supplies in the x3850 X6

The following types of power supplies are supported:

- ▶ 1100 W AC (100Vac - 120Vac or 200Vac - 240Vac)
- ▶ 1600 W AC (200Vac - 240Vac only)

The SR950 supports the following modes of redundancy. The Power redundancy is controlled by the XClarity Controller (XCC):

- ▶ Redundant (N+N)

The server is guaranteed to continue to remain operational with the loss of one power supply feed

- ▶ Non-redundant (N):

Non-Redundant: The server is not guaranteed to remain operational with the loss of a power supply feed. The system may throttle if a power supply fails in an attempt to stay powered up and continue running.

You can set and change the Power redundancy by using the XCC 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 XCC and updating the settings, see 5.2.1, “Accessing XCC” on page 142.

For information on how to configure the power redundancy for the SR950 server, refer to “Configuring power redundancy” on page 149.

4.7.3 Power supply installation order

The system can operate with 1, 2 or up to 4 power supplies, depending on the hardware installed in the system. The power supply installation order should be from bottom to top as indicated in Figure 4-4 on page 108. The following rules apply:

- Servers with two processors can operate with one power supply, installed in power-supply bay 1. The system will be non-redundant and a power feed or power supply failure will cause the system to crash.
- Servers with four or eight processors require a minimum of two power supplies, one of which must be installed in power-supply bay 1, then 2, then 3, then 4.

4.8 Power monitoring

The Lenovo XClarity Controller (XCC) and the Lenovo XClarity Energy Manager (LXEM) provide monitoring of the power consumption and cooling of the SR950 server. These features help with infrastructure planning, particularly for adding on or expanding the data center. The following monitoring capabilities are discussed:

- 4.8.1, “XClarity Controller power and cooling” on page 109
- 4.8.2, “Lenovo XClarity Administrator power and cooling” on page 112

4.8.1 XClarity Controller power and cooling

The XClarity Controller (XCC) provides an overview of the SR950’s power and cooling consumption. To view the power and cooling information from the XCC, follow these steps:

1. Log into XCC as explained in 5.2.1, “Accessing XCC” on page 142.
2. You can view a quick overview of the system utilization for the CPU, memory, I/O and system, and power utilization for each CPU and all DIMMs from the XCC home page as seen in Figure 4-5.




Figure 4-5 Power utilization

3. To access the utilization page, click **Utilization** from the left hand menu in XCC. The utilization page is broken up into four categories:
- “Temperature” on page 110

- ▶ “Power utilization” on page 110
- ▶ “System utilization” on page 111
- ▶ “Fan speed” on page 111

Temperature

The utilization page displays the real time ambient temperature of the server and component temperature for the CPU and memory DIMMs, as seen in Figure 4-6.

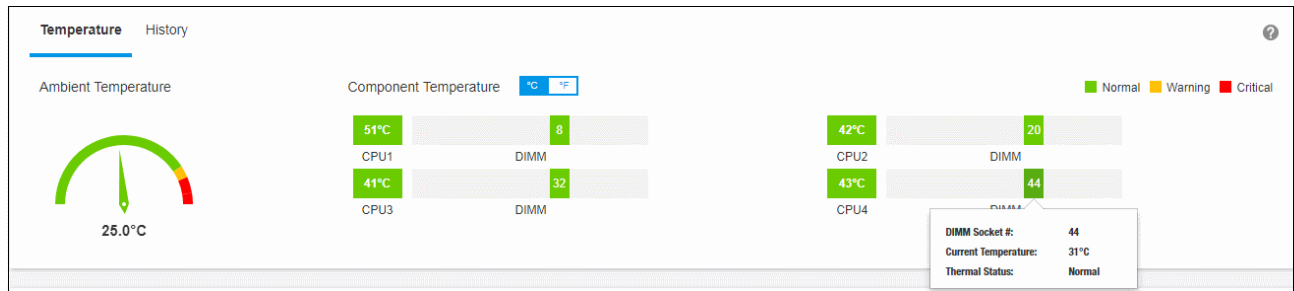


Figure 4-6 Real time ambient temperatures

The history tab at the top displays the historical temperature charts for the past 24 hours as seen in Figure 4-7.




Figure 4-7 Historic ambient temperatures up to 24 hours

Power utilization

The Utilization page also displays the power utilization. The power utilization section displays the current power consumption pie chart, as well as the historical power consumption charts for the past 24 hours.

The current consumption pie chart consists of four categories: CPU, Memory, Other and Spares. The Other means the total system power consumption minus the CPU and memory power consumption. Spare means the total available allocated power minus the total system power consumption.

The Voltage tab displays the current voltage readings and status. The power utilization section is shown in Figure 4-8 on page 111.




Figure 4-8 Power utilization in XCC

System utilization

The system utilization function represents the real-time utilization of the system, CPU, memory and I/O subsystems:

- ▶ The CPU subsystem-level utilization represents the percentage of the total CPU bandwidth currently in use, as measured by the performance counters built in to the CPU.
- ▶ The memory subsystem-level utilization represents the percentage of total memory channel controller bandwidth currently in use.
- ▶ I/O subsystem-level utilization represents the percentage of total PCIe traffic bandwidth currently in use.

The system utilization section is shown in Figure 4-9.




Figure 4-9 System utilization

Fan speed

The fan speed is also recorded. The fan speed section shows the fan speeds as a percentage of the maximum speed. Note that hovering the mouse cursor over the icon will show the fan speed percentage, as seen in Figure 4-10.

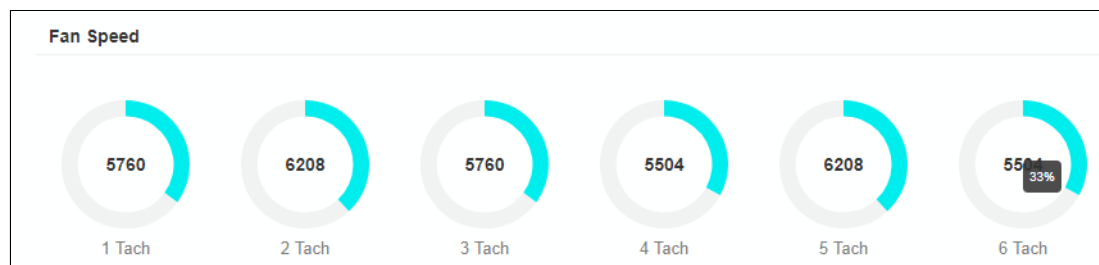


Figure 4-10 Fan speed percentage

Note: To set the power redundancy policy, power capping and power restore settings for the SR950 server, refer to 5.2.4, “Configure power polices with XCC” on page 149.

4.8.2 Lenovo XClarity Administrator power and cooling

Lenovo XClarity Administrator (LXCA) allows control of the power usage of the SR950 server through improved data analysis to lower the overall total cost of the ownership (TCO) of the server. LXCA helps to optimize data center efficiency by enabling the following:

- ▶ Manage power consumption and temperature in your data center through a simple management console
- ▶ Monitor room, row, rack, and device levels at the data center
- ▶ Track platform power consumption, inlet temperature, CPU, and memory power consumption
- ▶ Set power limits and policies, CPU and memory power limits, as well as do-not-exceed power limits (cap)
- ▶ Optimize power use through cooling analysis, low-utilization analysis, server power characteristics, and workload placement analysis
- ▶ Reduce platform power consumption during power supply event
- ▶ Use with Red Hat Linux, SUSE Linux, and Microsoft Windows to monitor rack, blade and tower servers

For additional information and an example of using Lenovo XClarity Administrator, refer to 7.2, “Lenovo XClarity Administrator” on page 225.

4.9 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 SR950 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 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.

4.10 Uninterruptible power supply units

There are several rack-mounted uninterruptible power supply (UPS) units that can be used with the SR950 server. 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://datacentersupport.lenovo.com/us/en/solutions/lno-pwrconf>

Table 4-6 lists the uninterruptible power supply (UPS) units that are offered by Lenovo.

Table 4-6 Uninterruptible power supply units

Part number	Description
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)
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 Product Guides in the UPS category:

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

For additional information, refer to the UPS Technical Reference guide, which is located here:

<http://datacentersupport.lenovo.com/no/en/documents/lno-powinf>

4.11 PDU and line cord selection

There are several PDUs that can be used to power the SR950 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 Lenovo Capacity Planner tool to help determine your total power load of the server, which is available at this website:

<https://datacentersupport.lenovo.com/us/en/products/solutions-and-software/software/lenovo-capacity-planner/solutions/HT504651>

This section describes the power supply unit to PDU line cord options that are available to connect your SR950 power supplies to your PDU source.

4.11.1 Server-to-PDU power cord options

The power supplies that are installed in the SR950 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 power supplies do not include a power cord. For models of the SR950, the inclusion of a power cord is model dependent. Configure-to-order models can be configured without a power cord if desired.

Table 4-7 lists power cables options for the 1100 W and 1600 W power supplies.

Table 4-7 Power cords

Part number	Feature code	Description
Rack cables		
00Y3043	A4VP	1.0m, 10A/100-250V, C13 to C14 Jumper Cord
39Y7937	6201	1.5m, 10A/100-250V, C13 to C14 Jumper Cord
4L67A08369	6570	2.0m, 13A/100-250V, C13 to C14 Jumper Cord
4L67A08366	6311	2.8m, 10A/100-250V, C13 to C14 Jumper Cord
4L67A08370	6400	2.8m, 13A/100-250V, C13 to C14 Jumper Cord
39Y7932	6263	4.3m, 10A/100-250V, C13 to C14 Jumper Cord
4L67A08371	6583	4.3m, 13A/100-250V, C13 to C14 Jumper Cord
Country-specific cables		
39Y7930	6222	2.8m, 10A/250V, C13 to IEC 60320 (Argentina) Line Cord
81Y2384	6492	4.3m, 10A/250V, C13 to IEC 60320 (Argentina) Line Cord
39Y7924	6211	2.8m, 10A/250V, C13 to AS/NZS 3112 (Australia/NZ) Line Cord
81Y2383	6574	4.3m, 10A/250V, C13 to AS/NZS 3112 (Australia/NZ) Line Cord
69Y1988	6532	2.8m, 10A/250V, C13 to NBR 14136 (Brazil) Line Cord
81Y2387	6404	4.3m, 10A/250V, C13 to NBR 14136 (Brazil) Line Cord
39Y7928	6210	2.8m, 10A/220V, C13 to GB 2099.1 (China) Line Cord

Part number	Feature code	Description
81Y2378	6580	4.3m, 10A/250V, C13 to GB 2099.1 (China) Line Cord
39Y7918	6213	2.8m, 10A/250V, C13 to DK2-5a (Denmark) Line Cord
81Y2382	6575	4.3m, 10A/250V, C13 to DK2-5a (Denmark) Line Cord
39Y7917	6212	2.8m, 10A/250V, C13 to CEE 7/7 (Europe) Line Cord
81Y2376	6572	4.3m, 10A/250V, C13 to CEE 7/7 (Europe) Line Cord
39Y7927	6269	2.8m, 10A/250V, C13 to IS 6538 (India) Line Cord
81Y2386	6567	4.3m, 10A/250V, C13 to IS 6538 (India) Line Cord
39Y7920	6218	2.8m, 10A/250V, C13 to SI 32 (Israel) Line Cord
81Y2381	6579	4.3m, 10A/250V, C13 to SI 32 (Israel) Line Cord
39Y7921	6217	2.8m, 10A/250V, C13 to CEI 23-16 (Italy) Line Cord
81Y2380	6493	4.3m, 10A/250V, C13 to CEI 23-16 (Italy) Line Cord
4L67A08362	6495	4.3m, 12A/200V, C13 to JIS C-8303 (Japan) Line Cord
39Y7922	6214	2.8m, 10A/250V, C13 to SABS 164-1 (South Africa) Line Cord
81Y2379	6576	4.3m, 10A/250V, C13 to SANS 164-1 (South Africa) Line Cord
39Y7926	6335	4.3m, 12A/100V, C13 to JIS C-8303 (Japan) Line Cord
39Y7925	6219	2.8m, 12A/220V, C13 to KSC 8305 (S. Korea) Line Cord
81Y2385	6494	4.3m, 12A/250V, C13 to KSC 8305 (S. Korea) Line Cord
39Y7919	6216	2.8m, 10A/250V, C13 to SEV 1011-S24507 (Swiss) Line Cord
81Y2390	6578	4.3m, 10A/250V, C13 to SEV 1011-S24507 (Swiss) Line Cord
23R7158	6386	2.8m, 10A/125V, C13 to CNS 10917 (Taiwan) Line Cord
81Y2375	6317	2.8m, 10A/250V, C13 to CNS 10917 (Taiwan) Line Cord
81Y2374	6402	2.8m, 13A/125V, C13 to CNS 10917 (Taiwan) Line Cord
4L67A08363	AX8B	4.3m, 10A/125V, C13 to CNS 10917 (Taiwan) Line Cord
81Y2389	6531	4.3m, 10A/250V, C13 to CNS 10917 (Taiwan) Line Cord
81Y2388	6530	4.3m, 13A/125V, C13 to CNS 10917 (Taiwan) Line Cord
39Y7923	6215	2.8m, 10A/250V, C13 to BS 1363/A (UK) Line Cord
81Y2377	6577	4.3m, 10A/250V, C13 to BS 1363/A (UK) Line Cord
90Y3016	6313	2.8M, 10A/125V, C13 to NEMA 5-15P (US) Line Cord
46M2592	A1RF	2.8m, 10A/250V, C13 to NEMA 6-15P (US) Line Cord
00WH545	6401	2.8M, 13A/125V, C13 to NEMA 5-15P (US) Line Cord
4L67A08359	6370	4.3m, 10A/125V, C13 to NEMA 5-15P (US) Line Cord
4L67A08361	6373	4.3m, 10A/250V, C13 to NEMA 6-15P (US) Line Cord
4L67A08360	AX8A	4.3m, 13A/125V, C13 to NEMA 5-15P (US) Line Cord

110V customers: If you plan to use the ThinkSystem 1100W power supply with a 110V power source, select a power cable that is rated above 10A. Power cables that are rated at 10A or below are not supported with 110V power.

4.11.2 PDU selection

Table 4-8 lists the power distribution units (PDUs) that are offered by Lenovo.

Table 4-8 Power distribution units

Part number	Description
0U Basic PDUs	
00YJ776	0U 36 C13/6 C19 24A/200-240V 1 Phase PDU with NEMA L6-30P line cord
00YJ777	0U 36 C13/6 C19 32A/200-240V 1 Phase PDU with IEC60309 332P6 line cord
00YJ778	0U 21 C13/12 C19 32A/200-240V/346-415V 3 Phase PDU with IEC60309 532P6 line cord
00YJ779	0U 21 C13/12 C19 48A/200-240V 3 Phase PDU with IEC60309 460P9 line cord
Switched and Monitored PDUs	
00YJ780	0U 20 C13/4 C19 Switched and Monitored 32A/200-240V/1Ph PDU w/ IEC60309 332P6 line cord
00YJ781	0U 20 C13/4 C19 Switched and Monitored 24A/200-240V/1Ph PDU w/ NEMA L6-30P line cord
00YJ782	0U 18 C13/6 C19 Switched / Monitored 32A/200-240V/346-415V/3Ph PDU w/ IEC60309 532P6 cord
00YJ783	0U 12 C13/12 C19 Switched and Monitored 48A/200-240V/3Ph PDU w/ IEC60309 460P9 line cord
46M4002	1U 9 C19/3 C13 Switched and Monitored DPI PDU (without line cord)
46M4003	1U 9 C19/3 C13 Switched and Monitored 60A 3 Phase PDU with IEC 309 3P+Gnd line cord
46M4004	1U 12 C13 Switched and Monitored DPI PDU (without line cord)
46M4005	1U 12 C13 Switched and Monitored 60A 3 Phase PDU with IEC 309 3P+Gnd line cord
Ultra Density Enterprise PDUs (9x IEC 320 C13 + 3x IEC 320 C19 outlets)	
71762NX	Ultra Density Enterprise C19/C13 PDU Module (without line cord)
71763NU	Ultra Density Enterprise C19/C13 PDU 60A/208V/3ph with IEC 309 3P+Gnd line cord
C13 Enterprise PDUs (12x IEC 320 C13 outlets)	
39M2816	DPI C13 Enterprise PDU+ (without line cord)
39Y8941	DPI Single Phase C13 Enterprise PDU (without line cord)
C19 Enterprise PDUs (6x IEC 320 C19 outlets)	
39Y8948	DPI Single Phase C19 Enterprise PDU (without line cord)
39Y8923	DPI 60A 3 Phase C19 Enterprise PDU with IEC 309 3P+G (208 V) fixed line cord
Front-end PDUs (3x IEC 320 C19 outlets)	
39Y8938	DPI 30amp/125V Front-end PDU with NEMA L5-30P line cord
39Y8939	DPI 30amp/250V Front-end PDU with NEMA L6-30P line cord
39Y8934	DPI 32amp/250V Front-end PDU with IEC 309 2P+Gnd line cord
39Y8940	DPI 60amp/250V Front-end PDU with IEC 309 2P+Gnd line cord

Part number	Description
39Y8935	DPI 63amp/250V Front-end PDU with IEC 309 2P+Gnd line cord
NEMA PDUs (6x NEMA 5-15R outlets)	
39Y8905	DPI 100-127V PDU with Fixed NEMA L5-15P line cord
Line cords for PDUs that ship without a line cord ^a	
40K9611	DPI 32a Line Cord (IEC 309 3P+N+G)
40K9612	DPI 32a Line Cord (IEC 309 P+N+G)
40K9613	DPI 63a Cord (IEC 309 P+N+G)
40K9614	DPI 30a Line Cord (NEMA L6-30P)
40K9615	DPI 60a Cord (IEC 309 2P+G)
40K9617	DPI Australian/NZ 3112 Line Cord
40K9618	DPI Korean 8305 Line Cord

a. 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.

For more information, see the Lenovo Press documents in the PDU category:

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

For additional information refer to the International and North American PDU Technical Reference Guides, located here:

<http://datacentersupport.lenovo.com/no/en/documents/lnvo-powinf>

Preparing the hardware

This chapter provides assistance for making configuration, monitoring, and maintenance decisions when implementing the SR950 server. 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.

The two main methods of preparing the hardware are discussed here: with XClarity Provisioning Manager and with XClarity Controller.

Topics in this chapter are:

- ▶ 5.1, “Preparing the hardware with XClarity Provisioning Manager (UEFI)” on page 120
- ▶ 5.2, “Preparing the hardware with XClarity Controller” on page 142
- ▶ 5.3, “Other Lenovo tools for updating firmware” on page 159
- ▶ 5.4, “Overview of firmware deployment tools: Using the right tool” on page 176
- ▶ 5.5, “Backing up the server configuration” on page 177
- ▶ 5.6, “Tethering with XClarity Mobile” on page 178

5.1 Preparing the hardware with XClarity Provisioning Manager (UEFI)

The Lenovo XClarity Provisioning Manager (LXPM) is a UEFI application embedded tool. By default, the Lenovo XClarity Provisioning Manager Graphical User Interface (GUI) is displayed when you press F1 when prompted during boot.

The following topics are discussed in this section:

- ▶ 5.1.1, “Accessing Lenovo XClarity Provisioning Manager” on page 120
- ▶ 5.1.2, “Viewing system inventory” on page 124
- ▶ 5.1.3, “UEFI settings” on page 125
- ▶ 5.1.4, “Operating modes” on page 127
- ▶ 5.1.5, “Configuring common UEFI settings” on page 128
- ▶ 5.1.6, “Updating firmware with LXPM” on page 134
- ▶ 5.1.7, “Configuring RAID with LXPM” on page 135
- ▶ 5.1.8, “Configure Networking with LXPM” on page 139
- ▶ 5.1.9, “Installing an Operating System with LXPM” on page 141

5.1.1 Accessing Lenovo XClarity Provisioning Manager

LXPM provides an easy-to-use interface for setting up your server. You can perform the following tasks:

- ▶ View system inventory to see information about the devices that are installed in your server.
- ▶ Configure traditional UEFI settings
- ▶ Update the key system components firmware to the latest level.
- ▶ Set up RAID volumes for your server storage.
- ▶ Configure the servers IP address to access the XCC.
- ▶ Install an operating system on your server (unattended mode).
- ▶ Run diagnostic test on the hard disk drives and memory that are installed in the system.

There are two ways in which to enter F1 set-up to access LXPM.

- ▶ At boot-up:

When the server is booting, select F1 when prompted. The system will then enter into F1 set-up.

- ▶ Set to boot into F1 via XCC:
 - a. Connect to the XCC (refer to 5.2, “Preparing the hardware with XClarity Controller” on page 142 for connecting to the XCC).
 - b. Select **System Configuration** → **Boot Options** → **From the ‘One Time Boot Device’ drop down select ‘F1 system setup’**
 - c. The XCC One Time Boot Device options are shown in Figure 5-1 on page 121.




Figure 5-1 Boot option via XCC

- d. Select how the server should be rebooted as seen in Figure 5-2.




Figure 5-2 Reboot options from XCC

On the next reboot, the server will now automatically enter into F1.

Figure 5-3 on page 122 shows the boot screen that will be displayed once F1 has been initiated remotely as described above.

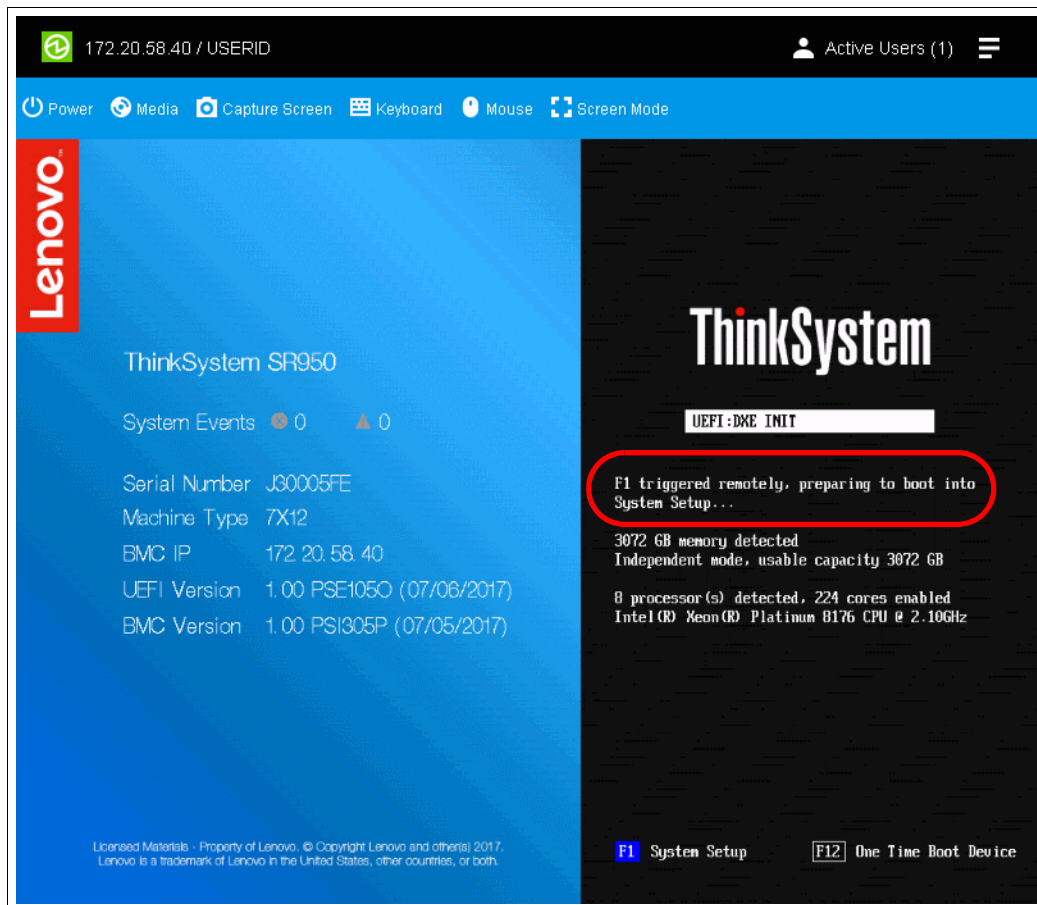


Figure 5-3 F1 booting screen of the XCC

The XClarity Provisioning Manager welcome screen is shown in Figure 5-4.

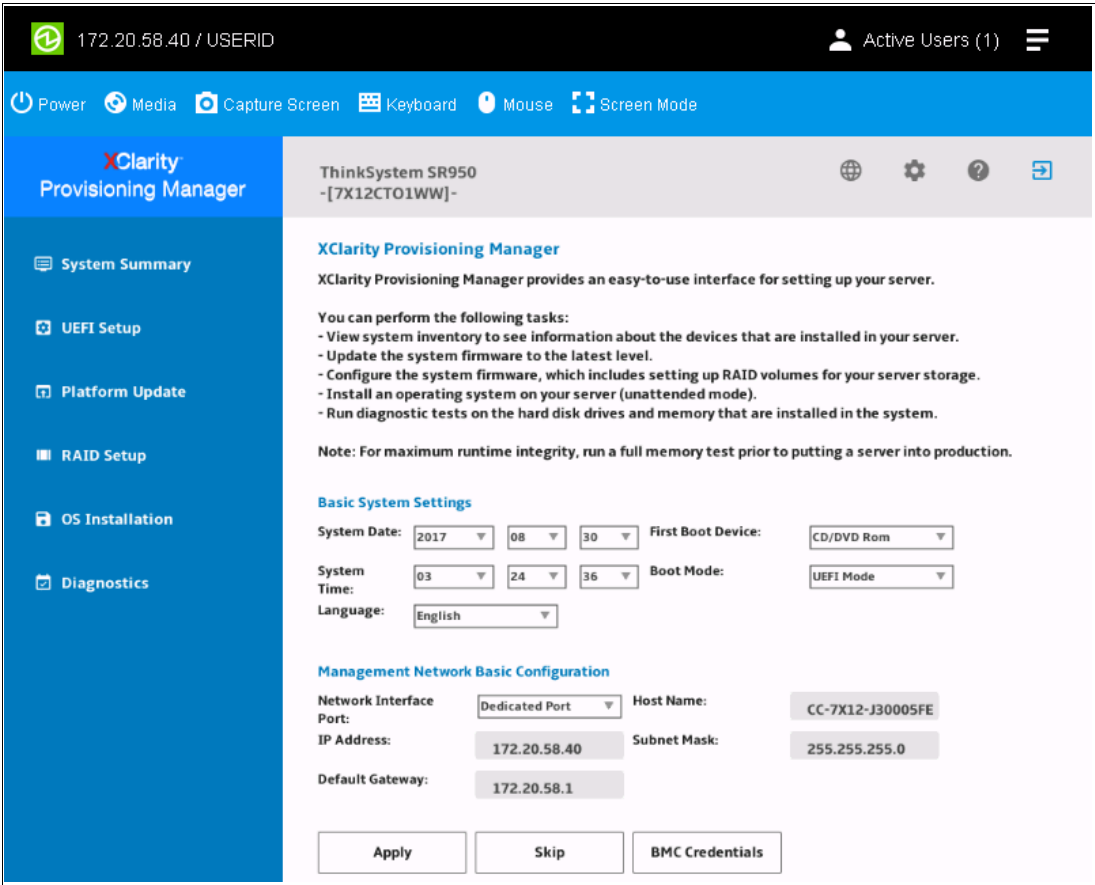


Figure 5-4 XClarity Provisioning Manager welcome screen

5.1.2 Viewing system inventory

The XClarity Provisioning Manager collects the systems inventory. You can access the inventory via the System Summary menu as shown in Figure 5-5.

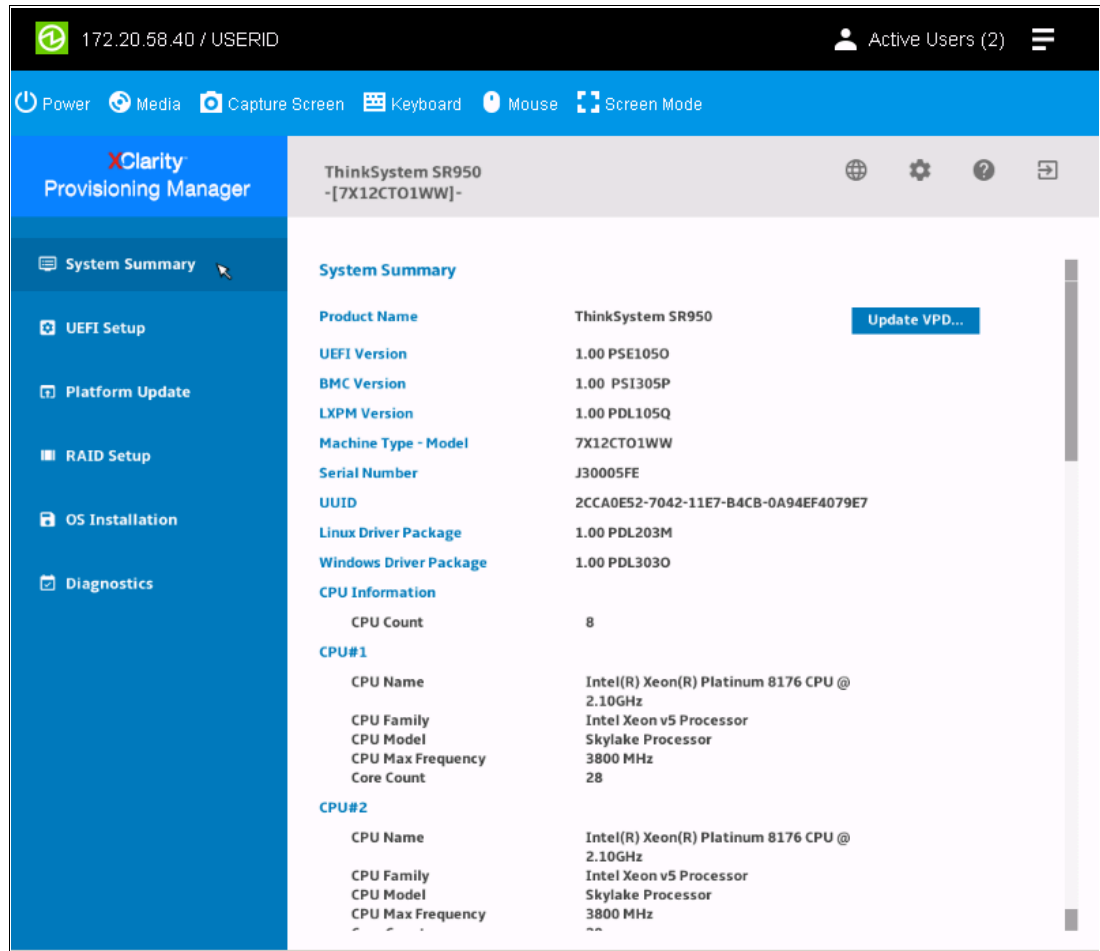


Figure 5-5 System summary screen in LXPM

The following inventory is collected with LXPM:

- ▶ Product name
- ▶ UEFI version
- ▶ BMC version
- ▶ LXPM version
- ▶ Machine type - model
- ▶ Serial number
- ▶ UUID
- ▶ Linux and Windows driver packages
- ▶ CPU information: Number of CPUs, name, family, model, max frequency, core count
- ▶ DIMMs: Number of DIMMs, capacity, type, speed, vendor
- ▶ PCI slots: Number of PCI cards, name of PCI cards
- ▶ Disks: Number of disks, model of disks

5.1.3 UEFI settings

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.

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.

To access the UEFI settings in LXPM, click **UEFI Setup** → **System Settings**. Figure 5-6 shows the UEFI system setting options.

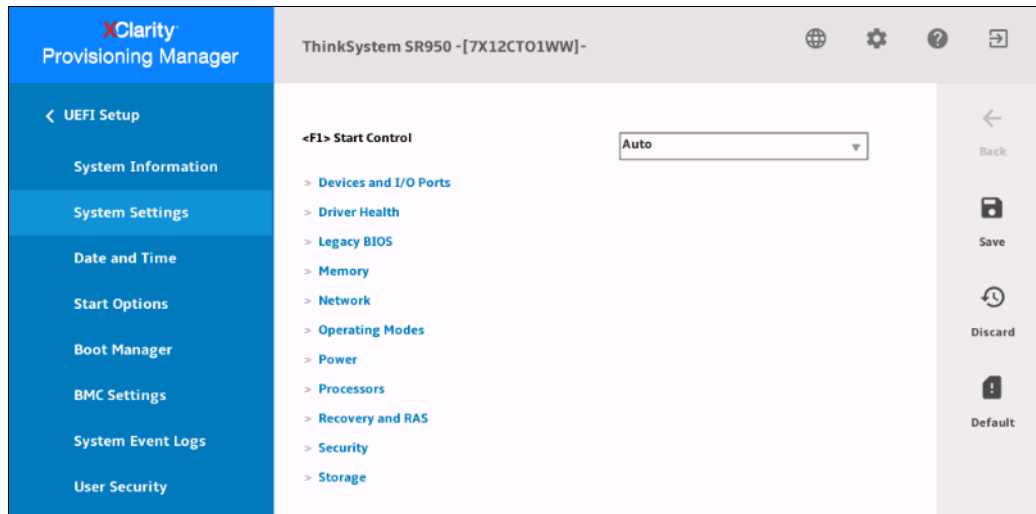


Figure 5-6 UEFI system settings menu

Table 5-1 lists the most commonly used UEFI settings and their default values.

Table 5-1 UEFI default settings

UEFI value	Default setting
Processor settings	
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
Trusted Execution Technology	Disabled
Intel Virtualization Technology	Enabled
Hardware Prefetcher	Enabled
Adjacent Cache Prefetch	Enabled
DCU Streamer Prefetcher	Enabled

UEFI value	Default setting
DCU IP Prefetcher	Enabled
Direct Cache Access (DCA)	Enabled
Energy Efficient Turbo	Enabled
Uncore Frequency Scaling	Enabled
MONITOR/MWATT	Enabled
SNC	Disabled
Snoop Preference	HS w. Directory + OSB +HitMe cache
Per Core P-state	Enabled
XPT Prefetcher	Enabled
UPI Prefetcher	Enabled
Cores in CPU Package	All
UPI Link Frequency	Max Performance
CPU Frequency Limits	Full turbo uplift
Memory Settings	
Memory Mode	Disabled
Memory Speed	Max Performance
Memory Power Management	Disabled
Socket Interleave	NUMA
Patrol Scrub	Enabled
Memory Data Scrambling	Enabled
Sparing	Disabled
Page Policy	Closed
Cold Boot Fast	Enabled
Memory Test	Automatic
Power	
Power/Performance Bias	Platform Controlled
Platform Controlled Type	Efficiency - Favor Performance
Workload Configuration	Balanced
Advanced RAS	
Machine Check Recovery	Enabled
PCI Error Recovery	Disabled

5.1.4 Operating modes

The SR950 server provides 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 setup settings provide several predefined setups for commonly wanted operation conditions. These predefined values are referred to as *operating modes*. Access the menu in UEFI by selecting **UEFI Setup** → **System Settings** → **Operating Modes** → **Choose Operating Mode**. You can see the five operating modes from which to choose, as shown in Figure 5-7. When a mode is chosen, the affected settings change to the appropriate predetermined values.




Figure 5-7 Operating modes

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.

- ▶ Custom Mode

Select this mode only if you understand the function of the low-level XCC settings. This mode is the only choice with which you can change the low-level XCC settings that affect the performance and power consumption of the server.

- ▶ 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.

5.1.5 Configuring common UEFI settings

The UEFI default settings are configured to provide optimal performance with reasonable power consumption. Other operating modes, as described in 5.1.4, “Operating modes” on page 127, are also available to meet various power and performance requirements. However, individual system settings enable users to fine-tune the characteristics of the SR950 server.

This section describes the following individual LXPM UEFI system settings:

- ▶ “System power settings”
- ▶ “Processor settings” on page 129
- ▶ “Memory settings” on page 132
- ▶ “UEFI settings for different workload types” on page 133

System power settings

To access the LXPM UEFI power settings of the SR950 server, from the left hand menu, click **UEFI Setup** → **System Settings** → **Power**

Figure 5-8 shows the available options in the UEFI system Power settings.

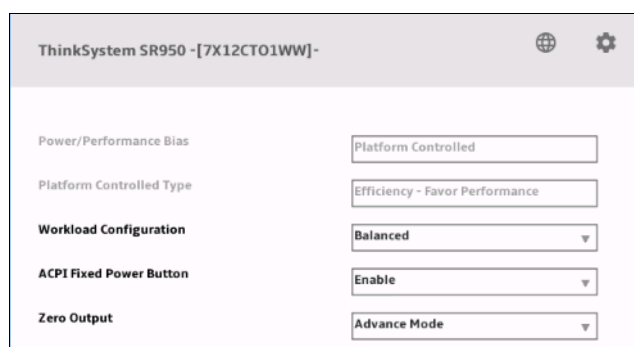


Figure 5-8 LXPM UEFI System settings - power

Power settings include the following options:

- ▶ 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, refer to 5.1.4, “Operating modes” on page 127.

- ▶ Platform Controlled Type

This option refers to the operating mode that your system is set to run in. For more information about operating modes, refer to 5.1.4, “Operating modes” on page 127.

- ▶ 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.

- ▶ **ACPI Fixed Power Button**

APCI stands for Advanced Configuration and Power Interface. When this option is enabled in UEFI, the server will allow power management of unused components in the server by the operating system, putting unused components to sleep.

- ▶ **Zero output**

The Zero Output setting can be in one of two states: Disabled, or Advanced Mode.

- Disabled (default):

When Disabled, the power supplies (PSUs) will operate in an Active/Active state.

When the PSU are operating in Active/Active mode, both PSUs share the power load of the system. If one of the PSUs fail, the power load is handled by the remaining PSU.

- Advanced Mode:

When Advanced Mode is activated, the power supplies will operate in an Active/Standby state when the load allows. When in an Active/Standby state, only one of the PSUs are operating and handling the power load of the system. The other PSU remains idle and will become active in the event of a PSU failure.

Processor settings

Processor settings control the various performance and power features that are available on the installed processor.

Figure 5-9 on page 130 shows the UEFI Processor system settings window with the default values.

» Processor Details	
Turbo Mode	Enable
CPU P-state Control	Autonomous
C-States	Autonomous
Hyper-Threading	Enable ▼
Execute Disable Bit	Enable ▼
Trusted Execution Technology	Disable ▼
Intel Virtualization Technology	Enable ▼
Hardware Prefetcher	Enable ▼
Adjacent Cache Prefetch	Enable ▼
DCU Streamer Prefetcher	Enable ▼
DCU IP Prefetcher	Enable ▼
DCA	Disable ▼
Energy Efficient Turbo	Enable
Uncore Frequency Scaling	Enable ▼
MONITOR/MWAIT	Enable
SNC	Disable ▼
Snoop Preference	HS w. Directory + OSB +HitME cache ▼
Per Core P-state	Enable ▼
XPT Prefetcher	Enable ▼
UPI Prefetcher	Enable ▼
Cores in CPU Package	All ▼
UPI Link Frequency	Max Performance
CPU Frequency Limits	Full turbo uplift ▼

Figure 5-9 LXPM UEFI System settings - processor

The most common processor feature options are listed below:

► 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 32.

► 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.
- ▶ 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 31.
- ▶ 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 31.
- ▶ 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.
- ▶ UPI 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

Memory settings

The Memory settings panel provides the available memory operation options, as shown in Figure 5-10.

> System Memory Details	
Total Memory Size	3072 GB
Mirror Mode	Disable
Sparing	Disable
Memory Speed	Max Performance
Memory Power Management	Disable
Socket Interleave	NUMA
Patrol Scrub	Enable
Memory Data Scrambling	Enable
Page Policy	Closed
Cold Boot Fast	Enable
Memory Test	Automatic

Figure 5-10 UEFI LXPM System settings - memory

The memory settings include the following options:

- ▶ **Mirror Mode (Default: Disabled)**
This option enables or disables memory mirroring. Memory mirroring is discussed in more detail in “Mirrored channel mode” on page 38. There are three options for memory mirroring mode:
 - Disabled: This option is the default.

- Full: This enables full memory mirroring across all DIMMs.
- Partial: This allows memory mirroring to be enabled/disabled when the system detects below 4GB of available memory.

Memory mirror mode: Memory mirror mode cannot be used with Memory sparing.

- ▶ Sparing (Default: Disabled)
This option enables or disables memory rank sparing, as described in “Rank sparing mode” on page 38.
- ▶ 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.
 - 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.
- ▶ 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.
- ▶ Cold Boot Fast (Default: Enabled)
This option enables or disables Cold Boot Fast.
- ▶ Memory Test (Default: Enabled)
This option enables, disabled or automatically allows for memory testing of the memory in the system.

Note: For any changes made in the UEFI Setup pages, ensure you click the **Save** button to commit the changes before you exit the setup.

UEFI settings for different workload types

Tuning the SR950 server 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 SR950 server 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 ^a	Low latency	Database	HPC
Turbo Mode ^b	Enabled	Enabled	Enabled	Enabled	Enabled
CPU P-States	Enabled	Enabled	Disabled	Enabled	Enabled
C-states	Enabled	Enabled	Disabled	Enabled	Disabled
Prefetcher	Enabled	Enabled	Enabled	Enabled	Enabled
Hyper-Threading	Enabled	Enabled	Disabled	Enabled	Disabled
UPI 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 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 Turbo Mode might also increase power consumption. The actual processing performance boost that you get from Turbo Mode 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.

5.1.6 Updating firmware with LXPM

You can update key firmware components of the system from the LXPM interface. LXPM allows you to update the following components:

- ▶ UEFI
- ▶ XCC (XClarity Controller)
- ▶ XClarity Provisioning Manager (LXPM)
- ▶ Windows Driver Bundle
- ▶ Linux Driver Bundle

You can access the LXPM firmware updates (either by browsing to the update package on a network share or over USB) by clicking **Platform Update**, selecting the appropriate checkboxes, then clicking **Update**.

The firmware update screen for LXPM is shown in Figure 5-11 on page 135.

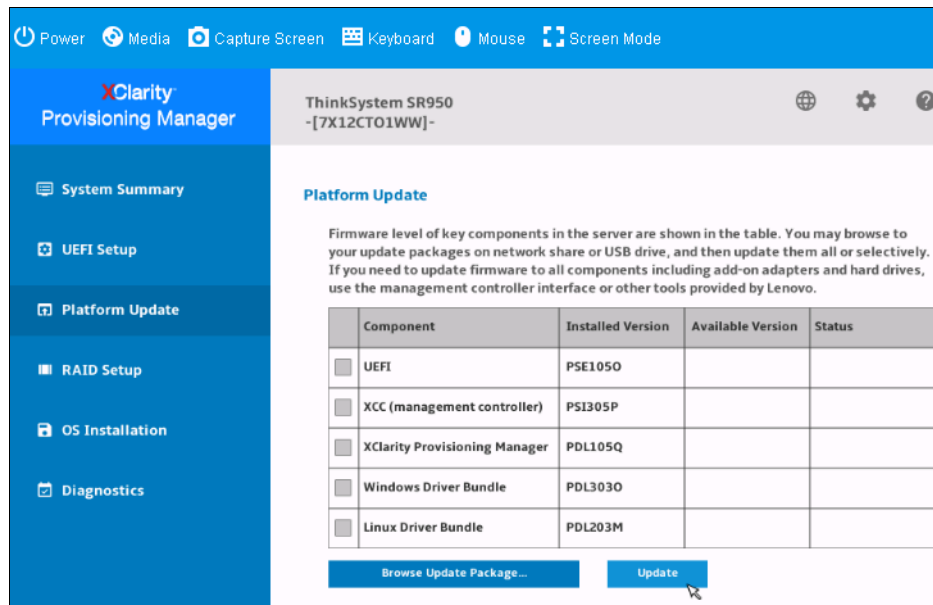


Figure 5-11 LXPM firmware update options

If other component firmware updates are required (for example, add-on adapters and drives) then use the XCC user interface (discussed in 5.2, “Preparing the hardware with XClarity Controller” on page 142) or other management tools (as discussed in 5.3, “Other Lenovo tools for updating firmware” on page 159).

5.1.7 Configuring RAID with LXPM

LXPM allows you to configure RAID on the server. Click **RAID Setup** via the navigation menu as shown in Figure 5-12.

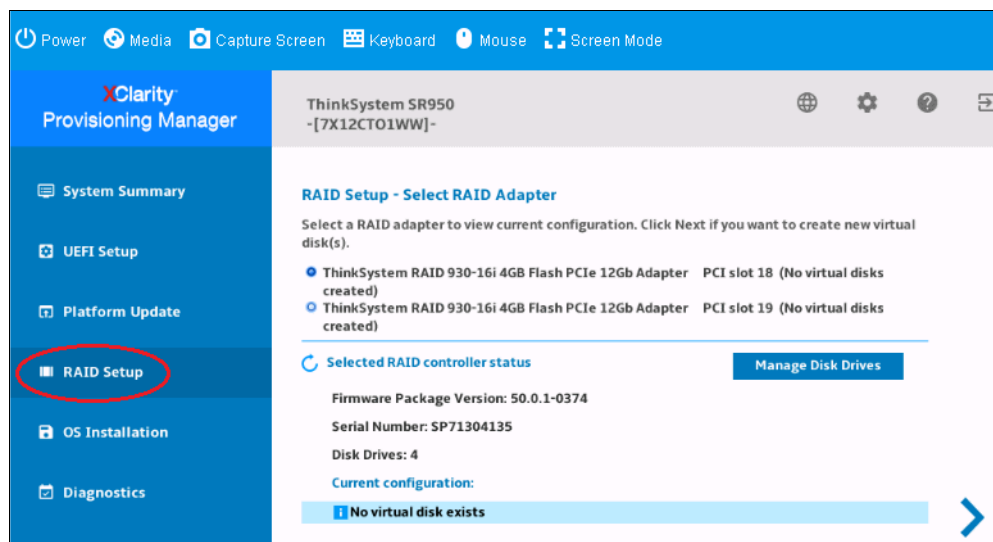


Figure 5-12 LXPM RAID setup screen

Tip: RAID setup can also be performed via XClarity Controller (XCC), as described in 5.2.5, “Configuring RAID with XCC” on page 153.

To manage the disk drives in the system and create a RAID array, ensure you have a compatible RAID controller in the system. See 3.9, “Controllers for internal storage” on page 72 for compatible RAID controllers and their functions.

The steps to set up RAID are as follows:

1. Select the RAID adapter from the Select RAID Adapter section and click **Manage Disk Drives**, as shown in Figure 5-12 on page 135. The Manage Disk Drives page appears, as shown in Figure 5-13.

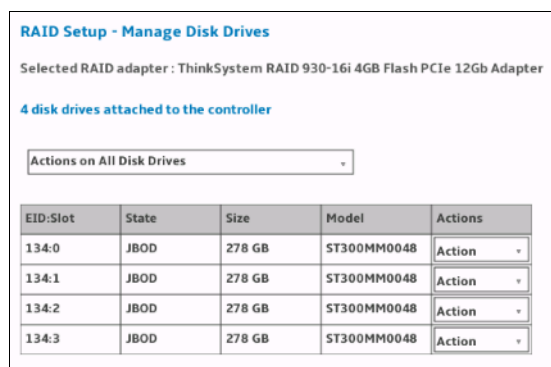


Figure 5-13 RAID Setup - Manage Disk Drives

2. Click the **Actions on All Disk Drives** drop-down menu and select **Change all disk drives from JBOD to UGood**. This will change all drives with a current state of JBOD (just a bunch of disks). See Figure 5-14.



Figure 5-14 Action on all disk drives

Alternatively, click the **Actions** drop-down menu next to each disk to change the state of each disk to Good, one at a time.

3. Click **Next** to Continue.
4. A warning will appear as shown in Figure 5-15. Click **Yes** to continue.




Figure 5-15 Warning on disk drive states

Warning: If you proceed at this point any existing data on the disk drives will be lost.

5. The system will take a moment to update the disk drives states from JBOD to a RAID capable state.

6. LXPM will take you back to the RAID Setup screen. Click **Next** as shown in Figure 5-16 to create an array(s) from the disks.

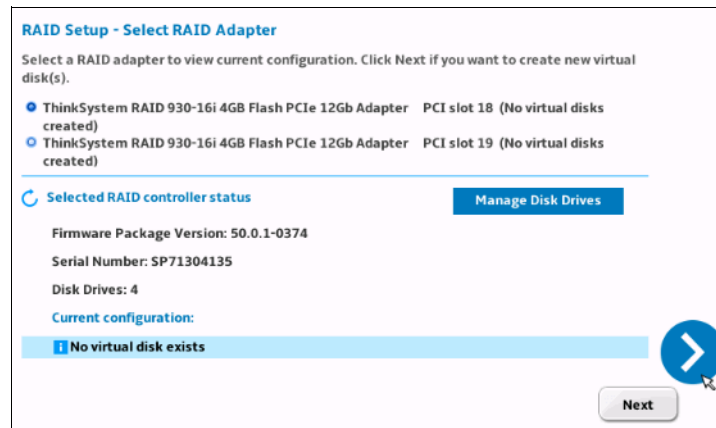


Figure 5-16 Create virtual disks

7. You are now presented with the Select RAID configuration type panel, Figure 5-17.




Figure 5-17

There are two options:

- **Simple configuration:** This will use all the disks drives available to the system to create a RAID array and hot-spare. See Figure 5-18.




Figure 5-18 Simple RAID configuration

- **Advanced configuration:** Create new disk array and virtual disks or use existing free capacity on existing disk array. Figure 5-19 shows the new disk array option.

RAID Setup - Select RAID Level and Select Drives

Create disk array by specifying the RAID level and disk drive.

Set RAID Level:

RAID 5 1 to 32 disk drives are needed for RAID 0.

Select Drives:

<input checked="" type="checkbox"/>	EID:Slot	Type	State	Capacity	Role
<input checked="" type="checkbox"/>	134:0	SAS/HDD	UGood	278 GB	Member
<input checked="" type="checkbox"/>	134:1	SAS/HDD	UGood	278 GB	Member
<input checked="" type="checkbox"/>	134:2	SAS/HDD	UGood	278 GB	Member
<input checked="" type="checkbox"/>	134:3	SAS/HDD	UGood	278 GB	Member

Figure 5-19 Advanced RAID configuration

8. Select the disks and RAID level required and click **Next**. LXPm will display a verification screen to confirm the details, Figure 5-20.

RAID Setup - Verify settings

Verify disk array and virtual disk settings. You can click back to change settings.

Disk Array:

RAID Level: 5
Disk Drives: 3
Dedicated Hot Spares: 1

New Virtual Disk:

Virtual Disk Name	Volume1
Capacity	556 GB
Stripe Size	64K
Read Policy	No read ahead
Write Policy	Write through
IO Policy	Direct
Disk Cache Policy	Unchanged

Next

Figure 5-20 RAID Setup - verification

9. Click **Next** to continue with the RAID creation. A warning message will appear; click **Yes** to continue.
10. The system may take some time to set up the RAID array depending on the number of disks installed. A status window will be displayed while the array is created, Figure 5-21.

RAID Setup - Creating Virtual Disk(s)

Creating virtual disk Volume1

● ● ● ● ●

Figure 5-21 Creating RAID array

11. Once the array has been created, a confirmation message will be displayed as shown in Figure 5-22.




Figure 5-22 RAID array created

12. The array will now be visible from the RAID Setup screen, as seen in Figure 5-23.




Figure 5-23 RAID Setup screen with newly created array

5.1.8 Configure Networking with LXPM

The networking for the XCC can be configured in LXPM from the settings menu as shown in Figure 5-24 on page 140.

XClarity Provisioning Manager

ThinkSystem SR950
-[7X12CTO1WW]-

Network Settings

Host Ethernet Adapter

- ☒ Obtain an IP address automatically
- ☐ Use the following IP address

IP Address: []

Subnet Mask: []

Default Gateway: []

Network Share via host Ethernet adapter

Provide the remote CIFS shared folder information to use for retrieving and saving files.

IP Address: []

User Name: []

Format: Domain@Username or User name

Password: []

Path: []

Example: My Share\My file

Buttons: Save and Connect, Save without Connect, Cancel

Figure 5-24 Network configuration settings in LXPM

Other network settings can be configured via LXPM from the left hand menu under **UEFI Setup** → **BMC Settings** → **Network Settings** as shown in Figure 5-25.

XClarity Provisioning Manager

ThinkSystem SR950 -[7X12CTO1WW]-

UEFI Setup

- System Information
- System Settings
- Date and Time
- Start Options
- Boot Manager
- BMC Settings**
- System Event Logs
- User Security

Network Settings

Network Interface Port: Dedicated

Fail-Over Rule: None

Burned-in MAC Address: 08-94-EF-40-79-E5

Hostname: XCC-7X12-J30005FE

DHCP Control: DHCP with fallback

IP Address: 172.20.58.40

Subnet Mask: 255.255.255.0

Default Gateway: 172.20.58.1

IPv6: Enable

Local Link Address: FE80:0000:0000:0000:0A94:EFFE:FE40:79E5/64

VLAN Support: Disable

> Advanced Setting for BMC Ethernet

Save Network Settings

Buttons: Back, Save, Discard, Default

Figure 5-25 Networking settings

Note: All of the original UEFI settings including network settings can be configured under the **UEFI Setup** pages. Refer to the 5.1.5, “Configuring common UEFI settings” on page 128 section for additional information on the UEFI settings for the server.

5.1.9 Installing an Operating System with LXPM

From the LXPM GUI you can install an operating system.

For compatible operating systems, refer to 3.23, “Operating system support” on page 94. To install an Operating System click the **OS Installation** option from the left hand menu, as shown in Figure 5-26.

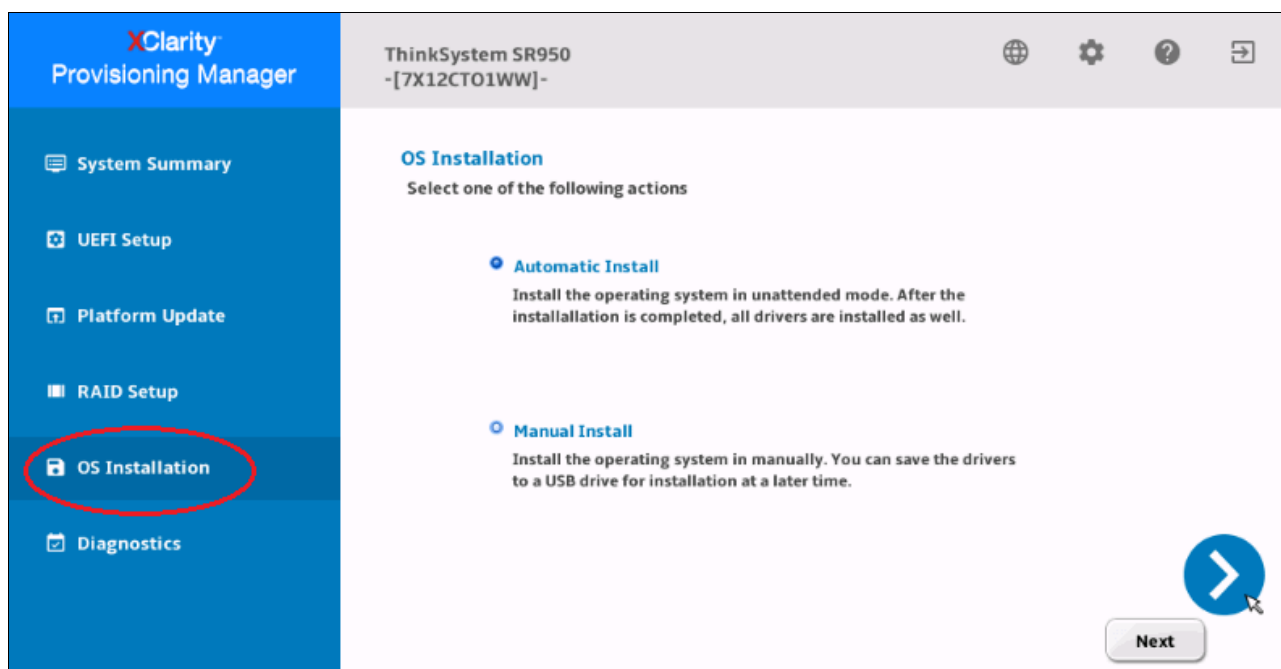


Figure 5-26 LXPM Operating System installation screen

There two methods for installing the operating system with LXPM:

- ▶ **Automatic Install:** This is an unattended mode where the operating system is installed and all necessary drivers are installed without user intervention.
- ▶ **Manual Install:** This is a manual installation where user intervention for driver installation is required.

Select a method of installation and click **Next**. There are four steps to installing the OS with LXPM:

1. Drive selection
2. OS selection
3. Installation settings
4. Partition options

For a detailed explanation of installing operating systems, refer to Chapter 6, “Operating system installation” on page 183.

5.2 Preparing the hardware with XClarity Controller

The SR950 server includes the integrated XClarity Controller (XCC). The XCC is a separate, independent service processor that is running whenever the server is plugged into a power source, even if the system is turned off. The XCC is based on the Pilot4 XE401 baseboard management controller (BMC) using a dual-core ARM Cortex A9 service processor.

The XCC monitors the hardware components of the server for potential hardware faults, allows you to update firmware, remotely deploy an OS and remotely manage it.

This section includes the following topics:

- ▶ 5.2.1, “Accessing XCC” on page 142
- ▶ 5.2.2, “Trouble shooting the XCC communications” on page 147
- ▶ 5.2.3, “Updating the firmware with XCC” on page 147
- ▶ 5.2.4, “Configure power policies with XCC” on page 149
- ▶ 5.2.5, “Configuring RAID with XCC” on page 153
- ▶ 5.2.6, “Configure networking with XCC” on page 157
- ▶ 5.2.7, “Installing an operating system with XCC” on page 158

This chapter discusses how to prepare the server with XCC. Management functions with XCC are discussed in 7.5, “Lenovo XClarity Controller management functions” on page 263.

5.2.1 Accessing XCC

XCC 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)

Remote presence, which is a standard feature of XCC as provided with the SR950, 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 XCC memory and mapping it to the server as a virtual drive.

The greatest strength of the XCC is the ability to completely monitor and manage the server from over the network. The following describes how to configure and connect to the XCC with the following topics:

- ▶ “Supported web browsers with XCC”
- ▶ “Default local network access to XCC” on page 143
- ▶ “Configuring a dedicated or shared network port” on page 144
- ▶ “Configuring static or DHCP network access” on page 145

Supported web browsers with XCC

The supported web browsers with XCC are listed below. The Firefox browser is recommended for JAWS users.

- ▶ Chrome 48.0 or above(55.0 or above for Remote Console)
- ▶ Firefox ESR 38.6.0 or above
- ▶ Microsoft Internet Explorer 11 (11.0.9600.18499 or above for Remote Console)
- ▶ Microsoft Edge
- ▶ Safari 9.0.2 or above (IOS7 or later and OS X)

If a browser version is not listed as supported, it may still work, especially on browsers with rapid release cycles (e.g. Chrome, Firefox). Note that support for the remote control feature is not available through the browser on mobile device operating systems.

For increased security, only high strength ciphers are supported when using https. When using https, the combination of your client OS and browser must support one of the following cipher suites:

- ▶ ECDHE-RSA-AES256-GCM-SHA384
- ▶ ECDHE-ECDSA-AES256-GCM-SHA384
- ▶ ECDHE-RSA-AES256-SHA384
- ▶ ECDHE-ECDSA-AES256-SHA384
- ▶ DHE-DSS-AES256-GCM-SHA384
- ▶ DHE-RSA-AES256-GCM-SHA384
- ▶ DHE-RSA-AES256-SHA256
- ▶ DHE-DSS-AES256-SHA256
- ▶ ECDH-RSA-AES256-GCM-SHA384
- ▶ ECDH-ECDSA-AES256-GCM-SHA384
- ▶ ECDH-RSA-AES256-SHA384
- ▶ ECDH-ECDSA-AES256-SHA384
- ▶ AES256-GCM-SHA384

Default local network access to XCC

The default network connection for the XCC on the SR950 is through the System Management port on the back of the server as shown in Figure 5-27.




Figure 5-27 XClarity Controller (XCC) port on the rear of the SR950

The XCC network access label is on the front of the server accessible via a pull-out tab. The label provides the default IPv4 host name and default IPv6 link local address of the XCC. A sample label is shown in Figure 5-28 on page 144.

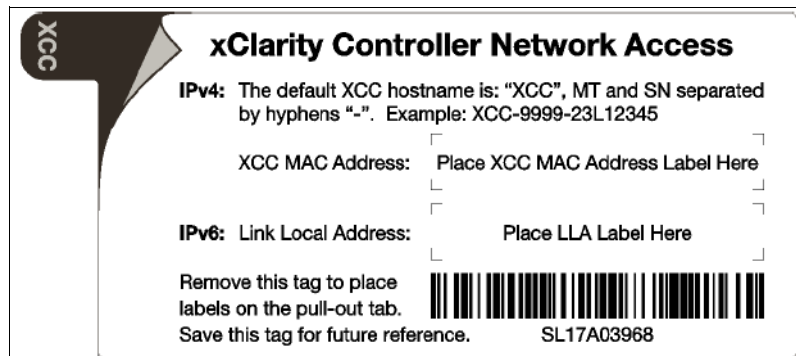


Figure 5-28 Default network details for the SR950

If you are local to the server you can connect a LAN cable to the dedicated XCC Ethernet port at the rear of the server to your laptop. Modify the IP settings of your laptop so that it is on the same network as the XCC IP address. Once connected, you can open a supported web browser and enter in the IP Address of the XCC. The XCC homepage will then be displayed.

Configuring a dedicated or shared network port

The SR950 offers two RJ45 ports that you can use to access XCC remotely, either the dedicated management port shown in Figure 5-27 on page 143, or via a port on an installed Ethernet LOM adapter which will be shared with the installed operating system. You select the port by setting Dedicated or Shared for the Network Interface Port in the Network Configuration panel of F1 setup.

Refer to 5.1.1, “Accessing Lenovo XClarity Provisioning Manager” on page 120 for information on how to access the F1 setup options (LXPM). Once in LXPM, you can access this panel by selecting **UEFI Setup** → **BMC Settings** → **Network Settings**, as shown in Figure 5-29.

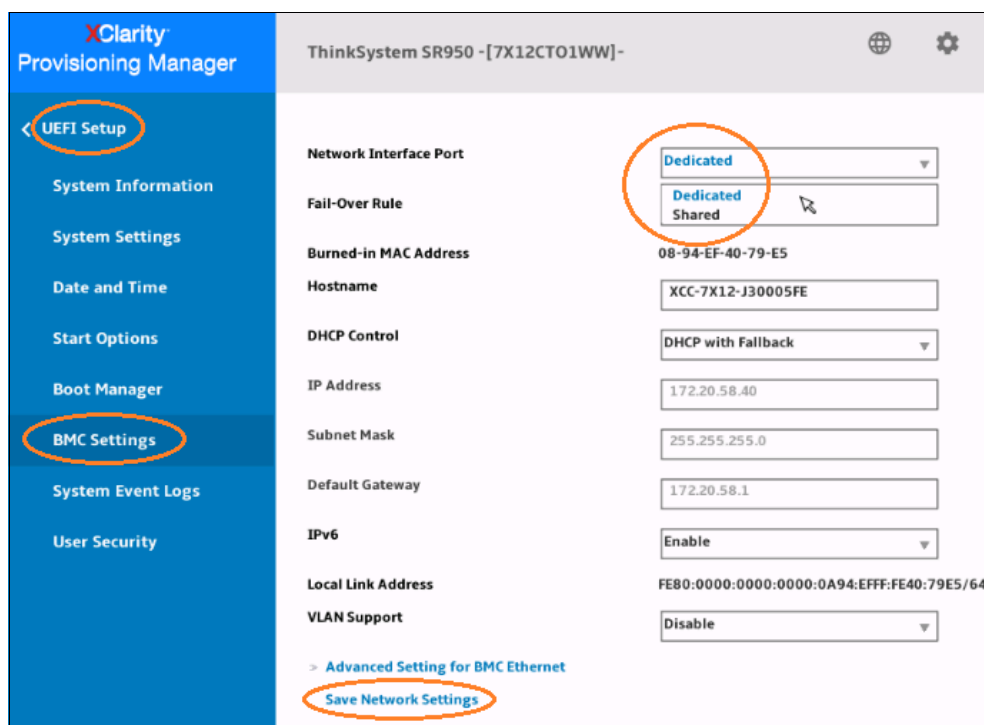


Figure 5-29 Configuring dedicated or shared XCC port

XCC dedicated versus shared Ethernet port

When configured as **Dedicated**, you are connecting to the network via the system management port as shown in Figure 5-27 on page 143.

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 between the management port and on an Ethernet adapter.

Although the Shared 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 XCC 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 other adapter's Ethernet port in a network team by using 802.3ad load balancing.

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 XCC MAC and dropped by the second Ethernet port's MAC.

Although the XCC uses a dedicated ARM Cortex A9 service processor, there are limitations as to the amount of network traffic that the XCC 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) embedded into LXPM is not as tolerant. For this reason (maintaining secured access), we recommend you keep your XCC management traffic on a separate management network.

Configuring static or DHCP network access

You can set the XCC with a static IP or for DHCP or DHCP with fallback (to static).

This can be set in one of two ways:

- ▶ Through the XClarity Controller (XCC) itself
- ▶ Through the F1 option via Lenovo XClarity Provisioning Manager (LXPM)

Both of these methods are discussed below.

Configure access via XClarity Controller

Connect to the XCC using the default IP address as discussed in "Default local network access to XCC" on page 143. In XCC, you can set the networking under **BMC Configuration** → **Network** as shown in Figure 5-30 on page 146.

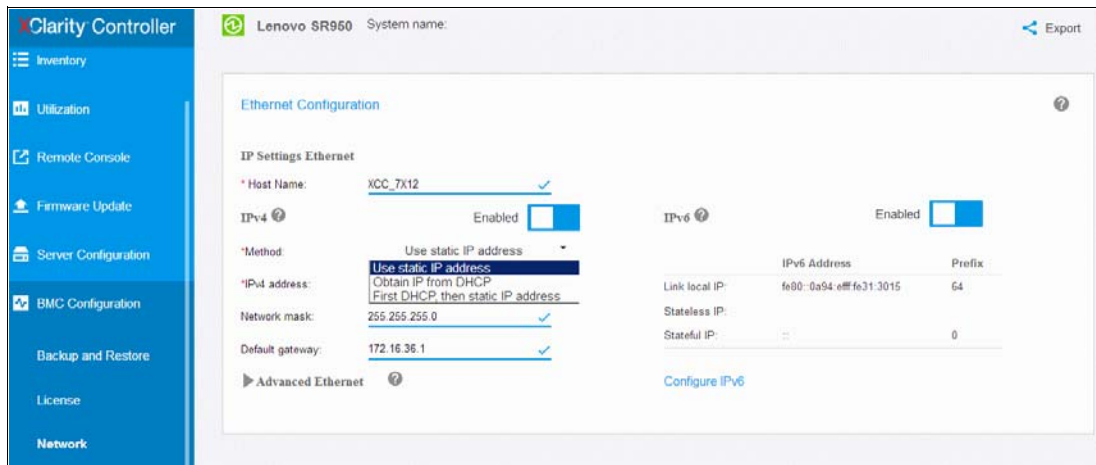


Figure 5-30 Set the XCC IP via XCC BMC Configuration

If you choose a static IP connection, make sure that you specify an IPv4 or IPv6 address that is available on the network. If you choose a DHCP connection, make sure that the MAC address for the server has been configured in the DHCP server.

Configuring access via Lenovo XClarity Provisioning Manager

To connect to LXPM, refer to 5.1.1, “Accessing Lenovo XClarity Provisioning Manager” on page 120 (at bootup). In LXPM, you can set the networking under **UEFI Setup** → **BMC Settings** → **Network Settings**, as shown in Figure 5-31.

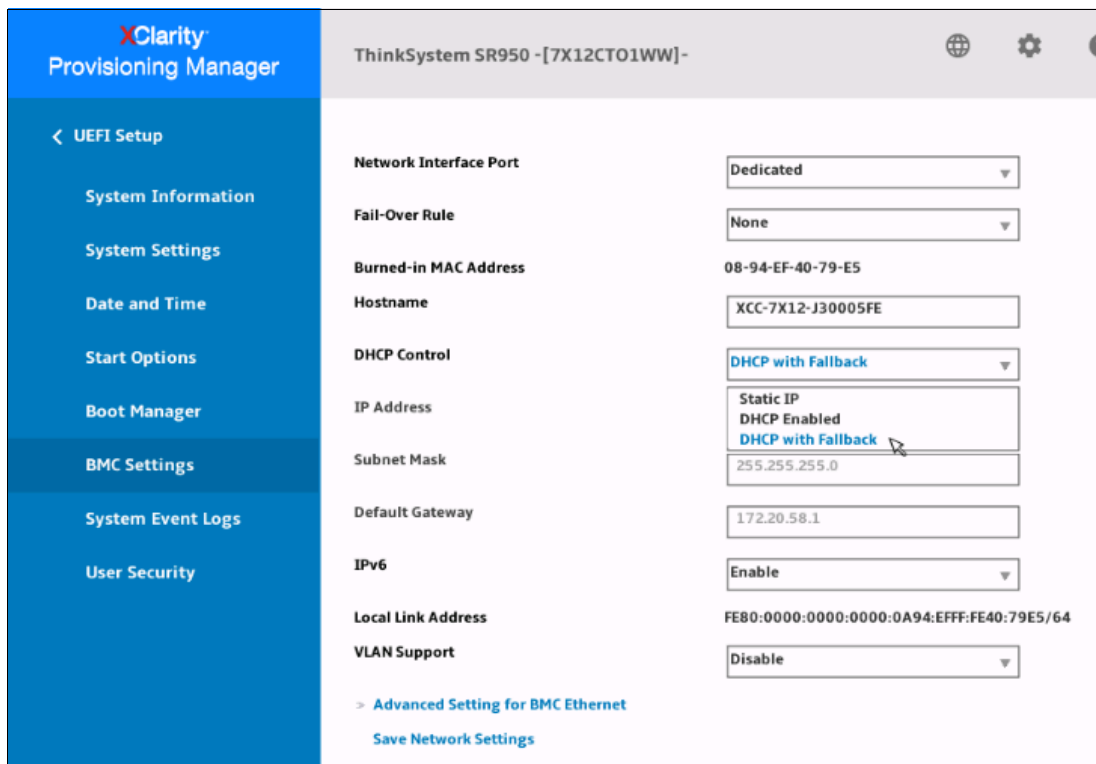


Figure 5-31 XCC Static or DHCP setup screen

If you choose a static IP connection, make sure that you specify an IPv4 or IPv6 address that is available on the network.

If you choose a DHCP connection, make sure that the MAC address for the server has been configured in the DHCP server.

5.2.2 Trouble shooting the XCC communications

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 XCC port with a notebook and Ethernet patch cable, pinging, and then starting a web session.

If you can ping the XCC, you have a good direct network link. If the web session fails, complete the following steps:

1. Try another web browser (For supported web browsers, refer to “Supported web browsers with XCC” on page 143.).
2. Access the F1 setup and restart the BMC. Select **UEFI Setup** → **BMC Settings** → **Restart BMC**. You must wait approximately 5 minutes before you may be able to ping it. This 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 through F1 setup by selecting **UEFI Setup** → **BMC Settings** → **Reset Factory Default Settings**.

If none of the above, fix the problem, contact Lenovo support for assistance.

5.2.3 Updating the firmware with XCC

The XCC has the ability to update the system and adapter firmware levels. For additional ways to update the firmware (and for updating drivers for an OS), refer to 5.3, “Other Lenovo tools for updating firmware” on page 159.

From the **Firmware Update** screen on the left hand menu, the current status and versions of firmware for the BMC, UEFI, LXPM, LXPM drivers, and adapters are displayed, including the BMC primary and backup versions.

There are four categories for the firmware status:

- ▶ Active: The firmware is active.
- ▶ Inactive: The firmware is not active.
- ▶ Pending: The firmware is waiting to become active.
- ▶ N/A: No firmware has been installed for this component.

Figure 5-32 on page 148 shows the firmware update screen with firmware information on the system and its adapters.

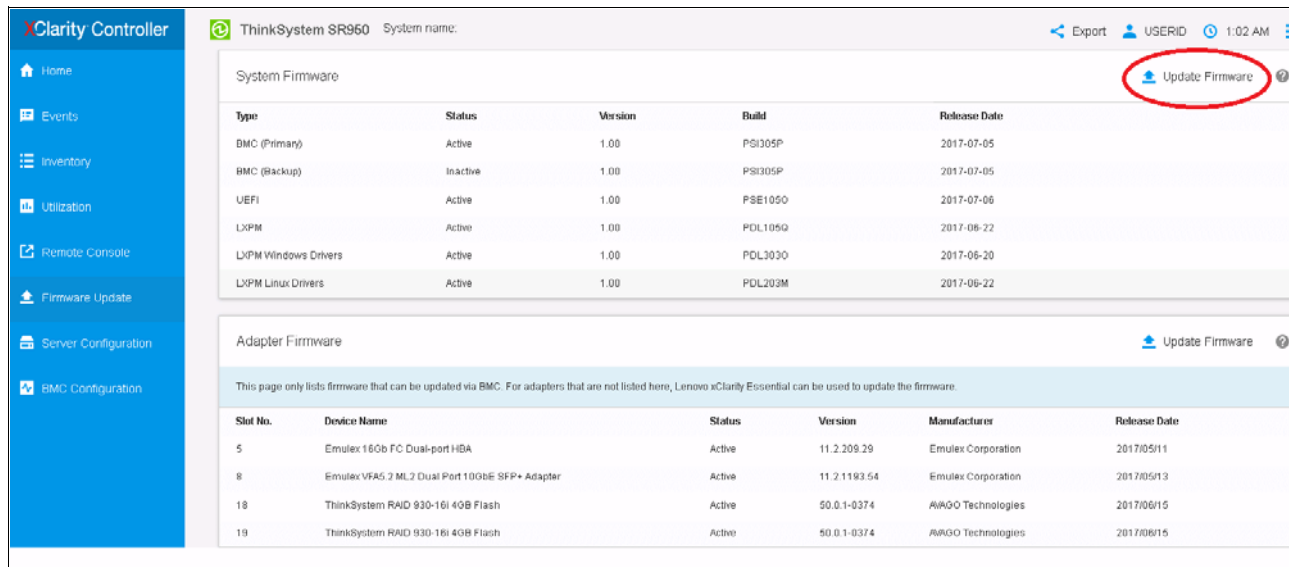


Figure 5-32 XCC firmware update screen

To update the server firmware, complete the following steps:

1. Click **Update Firmware** (as highlighted in Figure 5-32). The Update Server Firmware window opens as shown in Figure 5-33.

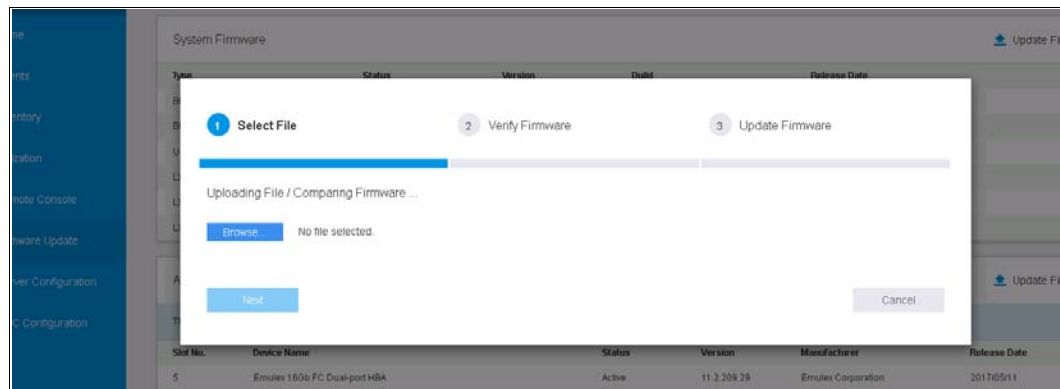


Figure 5-33 Firmware update selection

2. Click **Browse** to select the firmware update file that you want to use.
3. Navigate to the file you want to select and click **Open**. You are returned to the Update Server Firmware window with the selected file displayed.
4. Click **Next** to begin the upload and verify process on the selected file. A progress meter will be displayed as the file is being uploaded and verified. You can view this status window to verify that the file you selected to update is the correct file. The status window will have information regarding the type of firmware file that is to be updated such as BMC, UEFI, or LXPM.
5. After the firmware file is uploaded and verified successfully, click **Next** to select the device you want to update.
6. Click **Update** to begin the firmware update. A progress meter shows the progress of the update.

- When the firmware update is completed successfully, click **Finish**. If the update requires the XClarity Controller to be restarted in order to take effect, a warning message will be displayed. For details on how to restart the XClarity Controller, refer to “Power actions” on page 152.

5.2.4 Configure power polices with XCC

You can configure the power redundancy, power capping and power restore policies via XCC. This section discusses the following topics:

- ▶ “Configuring power redundancy”
- ▶ “Power capping” on page 150
- ▶ “Power restore policy” on page 151
- ▶ “Power actions” on page 152

Configuring power redundancy

The available power redundancy modes for the SR950 include:

- ▶ Redundant (N+N): In this mode, the server will remain operational in the event of a loss of a power supply.
- ▶ Non-redundant mode: In this mode, the server is not guaranteed to remain operational if a power supply is lost. The server will throttle if a power supply fails in an attempt to remain running.

Complete the following steps to configure the power redundancy by using the XCC web interface:

- Connect to the XCC interface as described in 5.2.1, “Accessing XCC” on page 142. Access the Server Power Policy page by selecting **Server Configuration** → **Power Policy** as shown in Figure 5-34.




Figure 5-34 XCC power policy page

2. Select the power redundancy that you want to implement on the system and click **OK** to implement that policy as shown in Figure 5-35.

	Estimated Usage	Total Power Capacity	Power Supply Failure Limit
<input checked="" type="radio"/> Redundant (N+N)	62%	3840W	2
<input type="radio"/> Non-redundant	37%	6400W	0

Apply Reset

Figure 5-35 Applying the power supply redundancy

The estimated usage, total power capacity, and power supply failure limit are shown on the power policy page. These have the following meanings:

- ▶ **Estimated usage:** The ratio of the maximum achievable power consumption of all components currently installed in the chassis to the Total Power Capacity.
- ▶ **Total Power Capacity:** The maximum DC power that is available to the server. It is calculated based on the number of power supplies installed, the wattage of the power supplies, and the selected power supply policy.
- ▶ **Power Supply Failure Limit:** This is the maximum number of power supplies that can fail while still guaranteeing the operation of the selected policy.

In this scenario, the system has four 1600W power supplies installed, configured as redundant (N+N). This means that two of the four power supplies will not be in use and are provisioned as backup power supplies in the event of a feed failure or power supply failure.

With two functioning power supplies, the total system capacity is approximately 3800W (or slightly higher as the system is allowed to throttle if needed).

The current usage is 62%. This means that the system has approximately 38% of its total power available (or approximately 1460W). The system can afford to lose two power supplies and continue operation.

Power capping

From the Power Policy page, you can also set power capping. To set the overall power limit on your server, click **Enable** under the Power Capping Policy section. Figure 5-36 on page 151 shows the power capping window in the XCC.

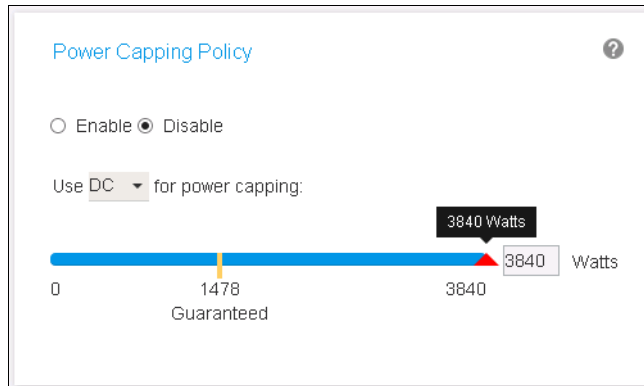


Figure 5-36 Power capping policy in the XCC

There are two options for power capping: Enabled and Disabled.

- ▶ **Power Capping Enabled:** Sets the overall system power limit in a situation where power on a component would cause the limit to be exceeded, the component would not be permitted to power on.
- ▶ **Power Capping Disabled:** The maximum power limit will be determined by the active power Redundancy policy.

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 restore policy

The power policy page also allows you to configure the restore policies. The policy that is set here will determine the action the system takes in the event of a power outage once power is restored. There are three options: Always off, Restore and Always on.

- ▶ **Always off:** System will remain off once power is restored
- ▶ **Restore:** Restore the system to the same state before the power failed. When restore is enabled you can select to have a random delay between 1 and 15 seconds for Power On. If the system state was on before power failed, the system will delay Power On once power is restored. This ensures the power has been properly restored.
- ▶ **Always on:** System will automatically Power On once power is restored

The power restore policy options are shown in Figure 5-37 on page 152.

Power Restore Policy

☐ Always Off
System will remain Off once power is restored.

☒ Restore
Restore System to the same state before power failed.

☒ Provides a random delay between 1 and 15 seconds for Power On.
If system state was on before power failed, the system will delay Power On once power is restored.

☐ Always On
System will automatically Power On once power is restored.

Apply

Reset

Figure 5-37 Power restore policy settings

Power actions

Some actions may require powering on and off components of the server. XCC can power on and off the following components listed in Table 5-3.

Table 5-3 Power actions available in the XCC

Power Action	Description
Power on server	Select this action item to power on the server and boot the operating system.
Power off server normally	Select this action item to shut down the operating system and power off the server.
Power off server immediately	Select this action item to power off the server without first shutting down the operating system.
Restart server normally	Select this action item to shut down the operating system and power cycle the server.
Restart server immediately	Select this action item to power cycle the server immediately without first shutting down the operating system.
Boot server to system setup	Select this item to power on or reboot the server and automatically boot into system setup without needing to press F1 during boot.
Trigger non-maskable interrupt (NMI)	Select this action item to force a Non-maskable Interrupt (NMI) on a “hung” system. Selection of this action item allows the platform operating system to perform a memory dump that can be used for debug purposes of the system hang condition. The auto reboot on NMI setting from the F1 system setup menu determines whether or not the XClarity Controller will reboot the server after the NMI.
Schedule power actions	Select this action item to schedule daily and weekly power and restart actions for the server.
Restart management controller	Select this action item to restart the XClarity Controller

To access the power options, go to the XCC home page. Go to **Quick Actions** → **Power Action** as shown in Figure 5-38 on page 153.

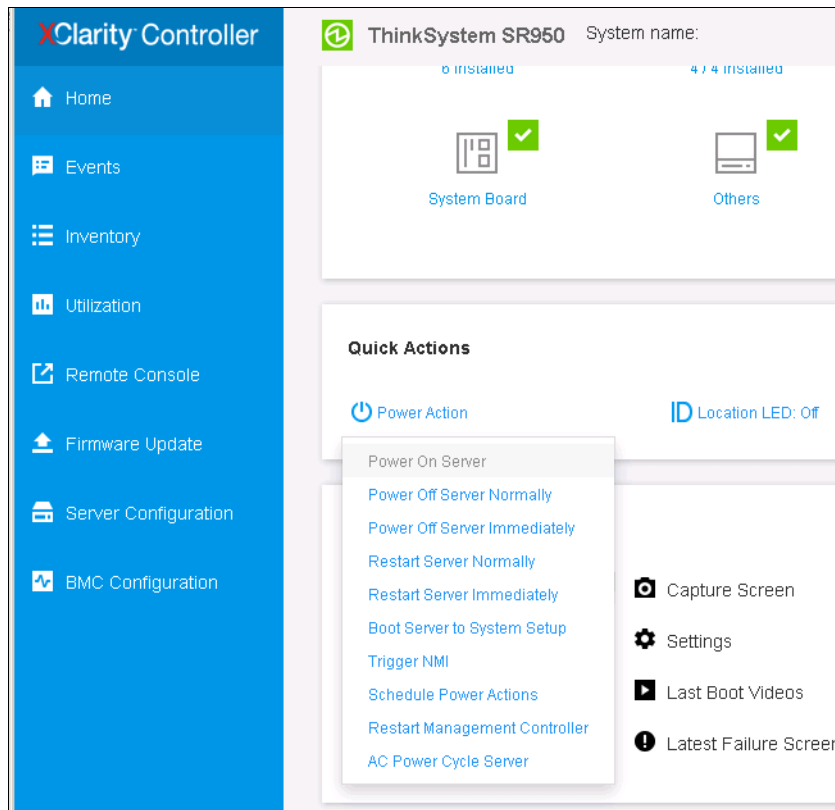


Figure 5-38 Power action menu

5.2.5 Configuring RAID with XCC

XCC allows configuration and viewing of a storage array, associated virtual disks and disk drives for the RAID adapter.

To configure RAID, follow these steps:

1. Connect to the XCC (as discussed in 5.2.1, “Accessing XCC” on page 142).
2. Click **Server Configuration** → **RAID Setup** from the left hand menu.

The initial RAID page in XCC will show the controllers and virtual disks in ‘read-only mode while the OS is running’ as seen in Figure 5-39 on page 154. Until you enable edit mode, the Controller Actions will remain greyed out.

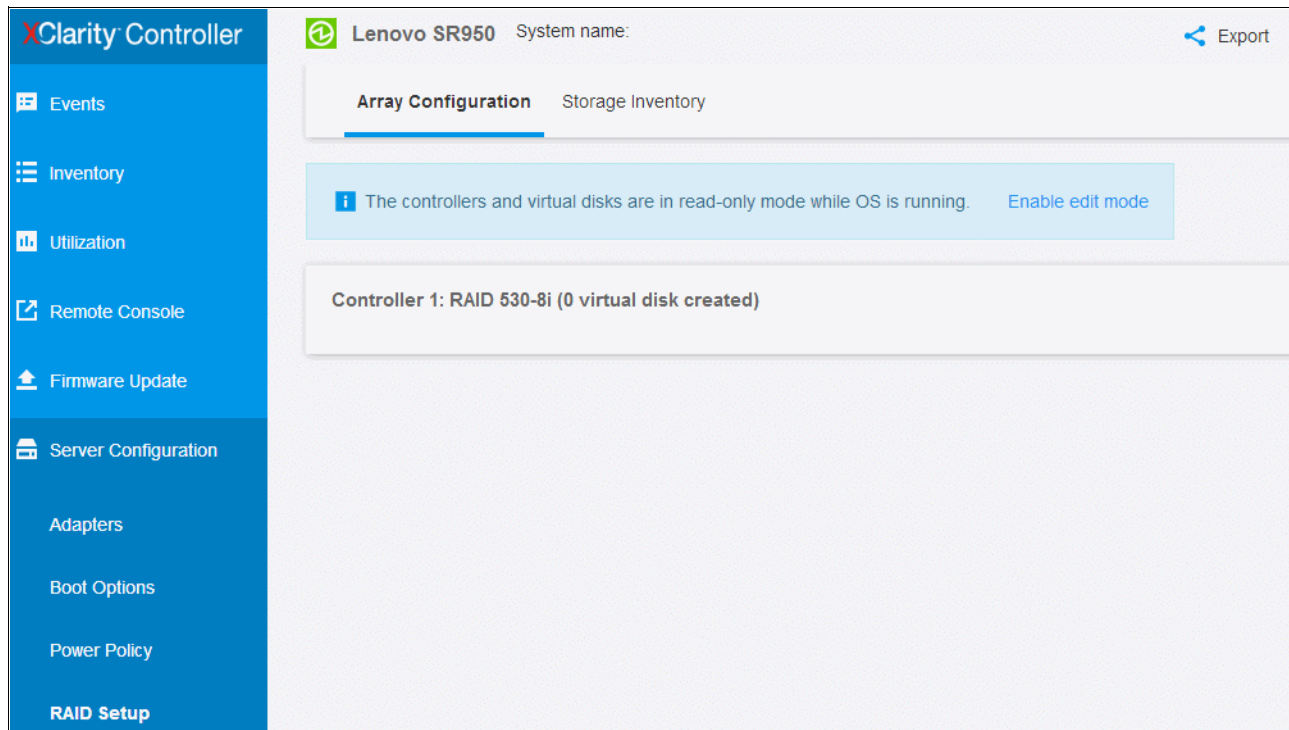


Figure 5-39 RAID Setup - in read only mode

3. Click **Enable edit mode** to allow the XCC to create and/or manage arrays, as seen in Figure 5-40. The option to create an array is now visible.

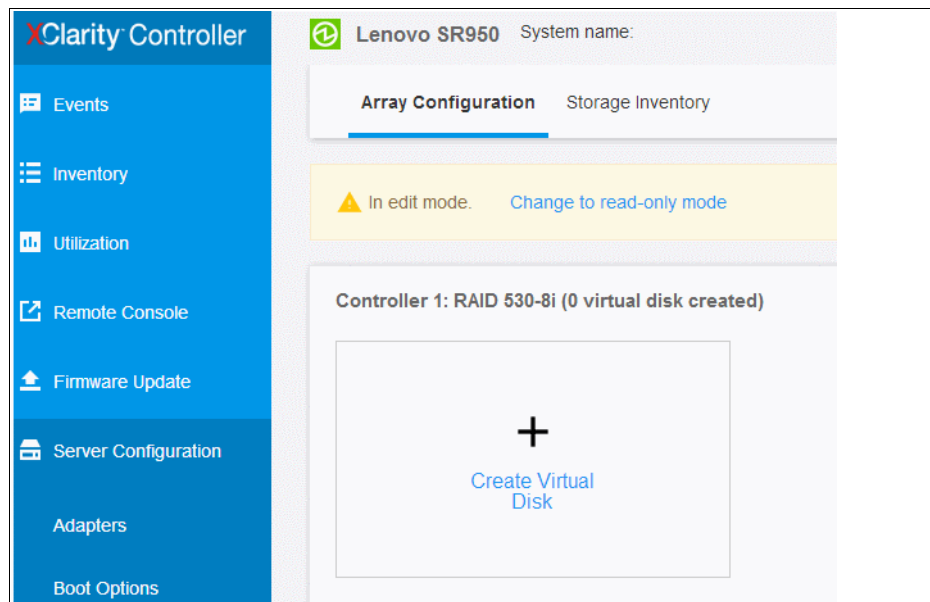


Figure 5-40 RAID Setup - enable edit mode

4. To create an array click the **Create Virtual Disk** option. Figure 5-41 on page 155 shows the RAID creation screen. Tick the disks you wish to create an array with and click **Add Member**. Next, select the RAID level from the drop down menu. Click **Next**.

1 Select Disk Drive/Disk Array

2 Create Virtual Disk

Get start:

Create new virtual disk on a new disk array

Select RAID level:

RAID 0 a minimum of 1 drives for RAID 0.

Selected disk drives:

Role	Disk Drive	Capacity
Member	Drive 4	300GB

Add member

Remove

Next >

Figure 5-41 RAID Setup - select drives and RAID level

- In the Create Virtual Disk page, Figure 5-42, click the **Actions** option and set the virtual disk name, capacity, stripe size and other settings. Click **Next**.

1 Select Disk Drive/Disk Array

2 Create Virtual Disk

By default 1 virtual disk will be created with all the available capacity. You can change the capacity if need create multiple virtual disks.

Virtual Disk Name	Capacity	Stripe Size	RAID Level	Actions
VD_2	278.46 GB	64 K	RAID 0	+

Edit Virtual Disk

Virtual Disk Name: VD_2

Capacity: 285148 MB

Stripe Size: 64 K

Read Policy: No Read Ahead

Write Policy: Write through

I/O Policy: Direct I/O

Access Policy: Read Write

Disk Cache Policy: Unchanged

Initialization Status: No Initialization

Figure 5-42 RAID Setup - options

- A confirmation of the RAID array configuration details is displayed as seen in Figure 5-43 on page 156.

1
Select Disk Drive/Disk Array

Review the summary and go back if you need to make corrections.

Disk array

RAID Level	RAID 0
Number of drives	1
Hot spare	0
Total capacity	278.46 GB
Free capacity	0 GB

New Virtual Disks

VD_2	278.46 GB
------	-----------

< Back
Start Creating

Figure 5-43 RAID Setup - create virtual disks

- Click the back button to amend any settings if needed. If the settings are correct, click the **Start Creating** button to create the virtual array. The server will take a moment to create the array.
- Once the array has been created a confirmation will be displayed and the RAID array will appear as seen in Figure 5-44.

Delete: If you wish to delete the array click the **Controller Actions** → **Clear RAID Configuration**.

Array Configuration
Storage Inventory

In edit mode.
Change to read-only mode

Controller 1: RAID 530-8i (1 virtual disk created)

Virtual Disk 1 Name: VD_1
Optimal
278.465GB
Disk Array 0, RAID 0

+
No disks available for virtual disk creation.

Controller Actions

- Clear RAID configuration
- Manage foreign configuration

Figure 5-44 RAID array

- Once the RAID array has been created, you can view the details of the array under the **'Storage Inventory'** tab. Figure 5-45 shows the Storage Inventory page with an array that has been created.

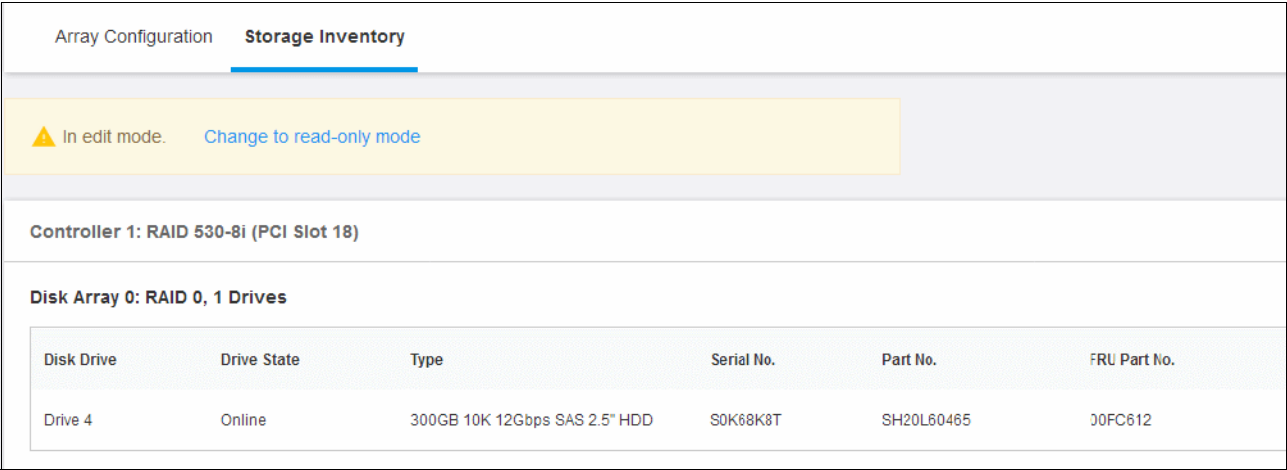
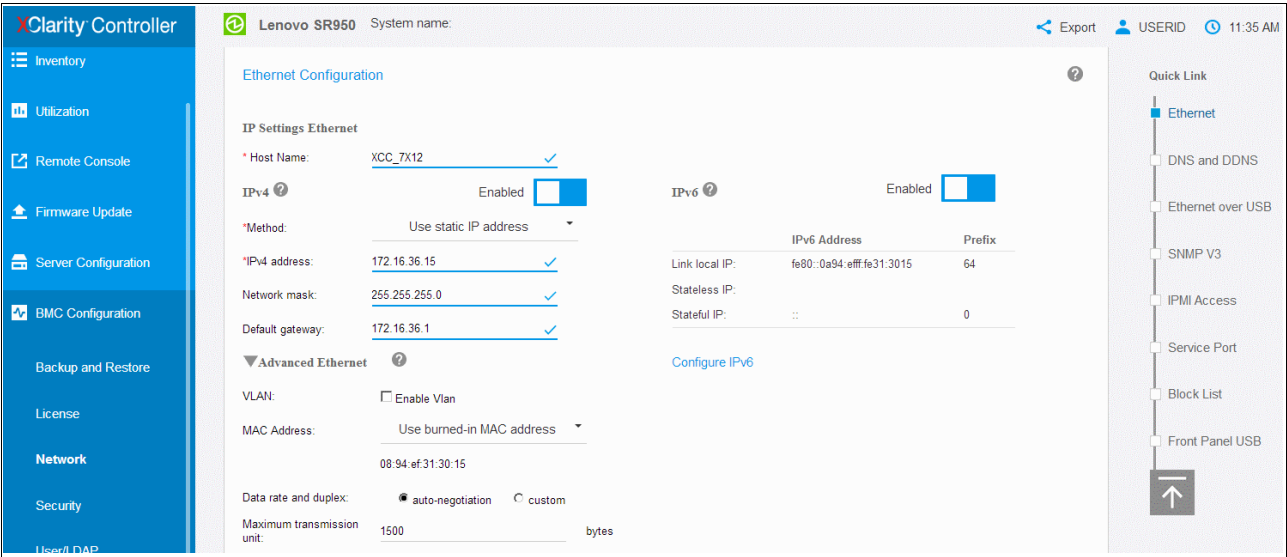


Figure 5-45 Storage inventory tab in the RAID setup page

5.2.6 Configure networking with XCC

The XCC allows the configuration of the following networking options for the server. To access the networking options in XCC, go to **BMC Configuration** → **Network** as shown in Figure 5-46.



► DNS and DDNS:

Configure if additional DNS server addresses can be included. DNS lookup is enabled by default. Note: DDNS enables or disables dynamic DNS (where the BMC notifies a domain name server to change, in real time).

► Ethernet over USB:

Enable external Ethernet to Ethernet over USB port forwarding. This provides a path between the external network interface and the server, through the BMC Ethernet over USB interface. The BMC opens a listening port and forwards the traffic on the port between the external network and the host.

► SNMPv3:

Enable and disable SNMP v3 for alert trapping.

► IPMI Access:

Enable or disable access of IPMI over LAN.

► Service Port:

Enable or disable a BMC service to change port numbers. Note: restart the BMC for new settings to take effect and that the port number selected is not in use by any other service. Also note that the following port numbers are reserved and can only be used for the corresponding services:

- IPMI port: 623
- SLP: 427
- SSDP: 1900

► Block List:

Configure IP or MAC addresses that will be denied access to the BMC.

► Front Panel USB:

The front panel USB connector can be switched so that it is attached either to the server host or to the BMC. This feature allows you to switch where the front panel USB port is connected to:

- In Host only Mode, the USB port is dedicated to the server.
- In BMC only Mode, the USB port is dedicated to the BMC.
- In Shared Mode, the BMC owns the USB port when the server is powered off. The BMC will return the USB port ownership back to the server when the server power is on and no longer using the USB port.
- When the port is configured for Shared Mode, pressing the Identification (ID) button for more than 3 seconds will switch ownership of the port between the server and the BMC.

5.2.7 Installing an operating system with XCC

From the XCC GUI you can install an operating system. You can remote mount an ISO image to the server using the XCC interface. For more information, see Chapter 6, “Operating system installation” on page 183.

5.3 Other Lenovo tools for updating firmware

This section discusses additional tools from Lenovo that enable you to update the servers firmware and operating system drivers. The following tools are discussed:

- ▶ 5.3.1, “XClarity Essentials UpdateXpress” on page 159
- ▶ 5.3.2, “XClarity Essentials Bootable Media Creator (BoMC)” on page 165
- ▶ 5.3.3, “XClarity Essentials OneCLI” on page 171
- ▶ 5.3.4, “Lenovo XClarity Administrator” on page 173

Warning: Installing the wrong firmware update might cause the server to malfunction. Before you install a firmware or device-driver update, read any readme and change history files that are provided with the downloaded update.

These files contain important information about the update and the procedure for installing the update, including any special procedure for updating from an early firmware or device-driver version to the latest version. The readme may also describe a required order of installation of updates.

5.3.1 XClarity Essentials UpdateXpress

Lenovo XClarity Essentials UpdateXpress is ideal for post-OS update, and periodic refresh of firmware and OS drivers.

UpdateXpress acquires updates automatically from the Lenovo Support website. it runs inside the Windows or Linux GUI environments to update the entire system or it can be used to apply the latest individual updates.

UpdateXpress uses the standard HTTP (port 80) and HTTPS (port 443) to get the updates from Lenovo. Your firewall must allow these ports. UpdateXpress is supported on Windows, Linux, and VMware operating systems. It is supported on 32-bit and 64-bit operating systems.

For more information about supported operating systems, see *UpdateXpress User's Guide*, which is available at this website:

<https://support.lenovo.com/documents/LNV0-XPRESS>

Complete the following steps to use UpdateXpress:

1. Start the UpdateXpress setup utility that you downloaded from the website.
2. Accept the license agreement.
3. The Welcome window opens, as shown in Figure 5-47 on page 160. Click **Next**.

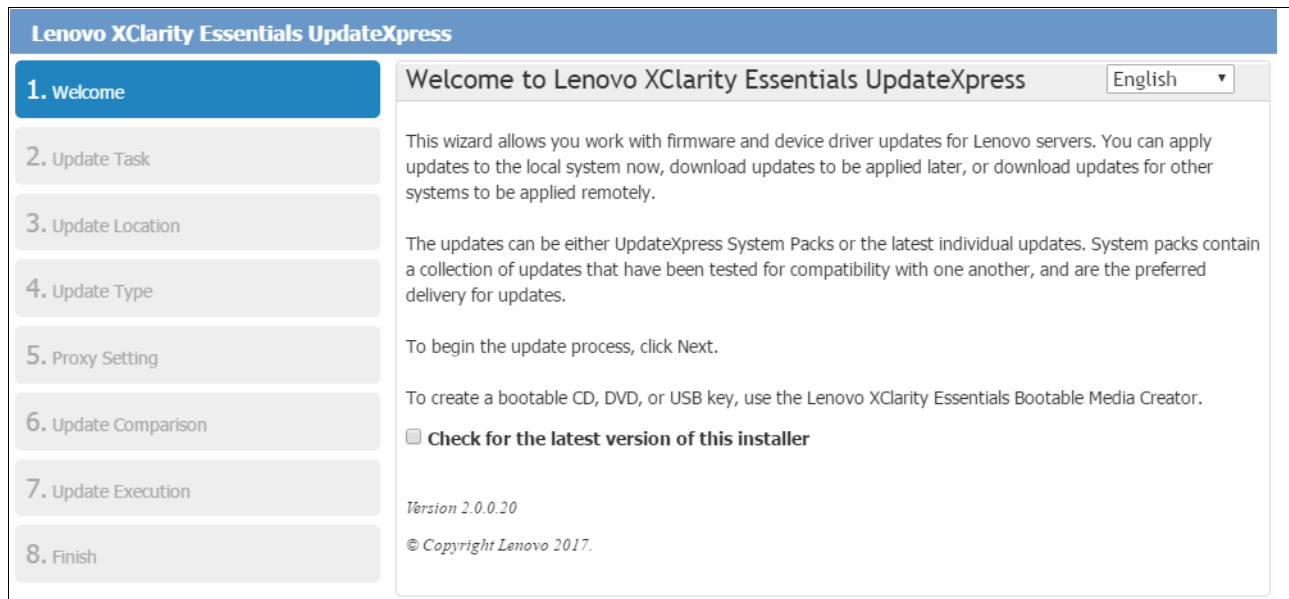


Figure 5-47 Welcome window

4. Select the way in which you will want to download/apply the updates, as shown in Figure 5-48. Click **Next**.

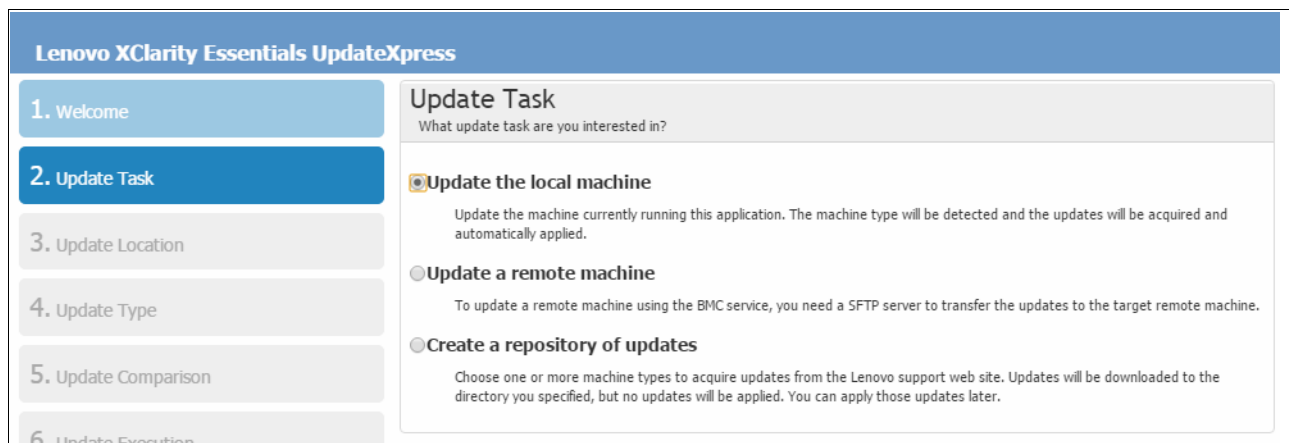


Figure 5-48 Selecting the update task

5. Select the update type, either a System Pack (a bundle of firmware and device drivers for everything) or latest available individual updates (update firmware and devices individually). In this scenario we are downloading individual updates, however the recommended method is to use the System Packs. Click **Next**.

Lenovo XClarity Essentials UpdateXpress

1. Welcome

2. Update Task

3. Update Type

4. Proxy Setting

5. Machine Types

Update Type

Select the type of updates you wish to search for

☒ **UpdateXpress System Packs (UXSP's)**

UpdateXpress System Packs contain an integration-tested bundle of online, updateable firmware and device driver updates for each Lenovo support and BladeCenter server. This is the preferred method to obtain firmware and driver updates for the server.

☐ **Latest available individual updates**

Use the latest individual version of each firmware and device driver package. This is the preferred method when you wish to install the latest updates, or when Lenovo support instructs you to install the latest updates. This option may download newer updates and hotfixes, if available, than the UXSP option above.

Figure 5-49 Update type

- If you require Internet connection you can set the HTTP Proxy settings here as shown in Figure 5-50. Click **Next**.

Lenovo XClarity Essentials UpdateXpress

1. Welcome

2. Update Task

3. Update Type

4. Proxy Setting

HTTP Proxy Setting

Enter the HTTP Proxy information here if you require to connect to the Internet. Internet connection is needed in order to download the appropriate updates from the Lenovo support web site.

☐ **I require a proxy server to connect to the Internet.**

Figure 5-50 Proxy settings

- Select the machine(s) from the list of servers that you require updates for, Figure 5-51. Click **Next**.

Lenovo XClarity Essentials UpdateXpress

1. Welcome

2. Update Task

3. Update Type

4. Proxy Setting

5. Machine Types

6. Operating Systems

7. Target Directory

8. Update Selection

9. Acquire Updates

Machine Types

Check for updates for machine types that you specify.

* UpdateXpress only supports Lenovo Server.

<input type="checkbox"/>	Lenovo ThinkSystem SR530 Server	7X08
<input type="checkbox"/>	Lenovo ThinkSystem ST550 Server	7X09
<input type="checkbox"/>	Lenovo ThinkSystem ST550 Server	7X10
<input type="checkbox"/>	Lenovo ThinkSystem SR950 Server	7X11
<input type="checkbox"/>	Lenovo ThinkSystem SR950 Server	7X12
<input checked="" type="checkbox"/>	Lenovo ThinkSystem SR950 Server	7X13
<input type="checkbox"/>	Lenovo ThinkSystem SN850 Server	7X15
<input type="checkbox"/>	Lenovo ThinkSystem SN550 Server	7X16
<input type="checkbox"/>	Lenovo ThinkSystem SR850 Server	7X18
<input type="checkbox"/>	Lenovo ThinkSystem SR850 Server	7X19
<input type="checkbox"/>	Lenovo ThinkSystem D2 Enclosure	7X20
<input type="checkbox"/>	Lenovo ThinkSystem SD530 Server / ThinkSystem D2	7X21

Add Remove Update List Reset List

Figure 5-51 Select machine(s) to be updated

8. Select the operating system that you will apply the updates on, Figure 5-52.

The screenshot shows the 'Lenovo XClarity Essentials UpdateXpress' interface. On the left is a vertical navigation pane with six items: '1. Welcome', '2. Update Task', '3. Update Type', '4. Proxy Setting', '5. Machine Types', and '6. Operating Systems' (which is highlighted in blue). The main area is titled 'Operating Systems' with the instruction 'Select the operating systems that are running on the machines you wish to apply updates to.' Below this is a table with checkboxes and operating system names:

<input type="checkbox"/>	Operating System
<input type="checkbox"/>	Windows 2008
<input checked="" type="checkbox"/>	Windows 2012
<input type="checkbox"/>	Windows 2016
<input type="checkbox"/>	Red Hat Enterprise Linux (RHEL) 7
<input type="checkbox"/>	SUSE Linux Enterprise Server (SLES) 12

Figure 5-52 Operating system selection

9. Select the target path that will host the individual driver and firmware updates, Figure 5-53.

The screenshot shows the 'Lenovo XClarity Essentials UpdateXpress' interface. The left navigation pane shows '1. Welcome', '2. Update Task', '3. Update Type', and '4. Proxy Setting'. The main area is titled 'Target Directory' with the instruction 'Choose a directory on your computer to save the updates.' Below this, there is a 'Directory:' label, a text input field containing 'C:\Firmware Updates', and a 'Browse...' button.

Figure 5-53 Target directory for drivers and firmware

10. The update selection screen will now look for and load any new updates that are available to the server, Figure 5-54.

The screenshot shows the 'Lenovo XClarity Essentials UpdateXpress' interface. The left navigation pane shows '1. Welcome', '2. Update Task', '3. Update Type', '4. Proxy Setting', and '5. Machine Types'. The main area is titled 'Update Selection' with the instruction 'Select updates to download and install.' Below this, there is a progress bar area with a loading icon and the text 'Acquiring... Obtaining update information (50% complete).'. A 'Cancel' button is visible. At the bottom, there is a list of updates with a checkbox and a green checkmark next to the text 'Lenovo ThinkSystem SR950 Server type 7X13 - Windows 2012 - 33 update(s)'.

Figure 5-54 Acquiring update options

11. In this example there are 33 updates ready to download as shown in Figure 5-55. Use the checkboxes to select and deselect the updates required for the SR950 server and click **Next**.

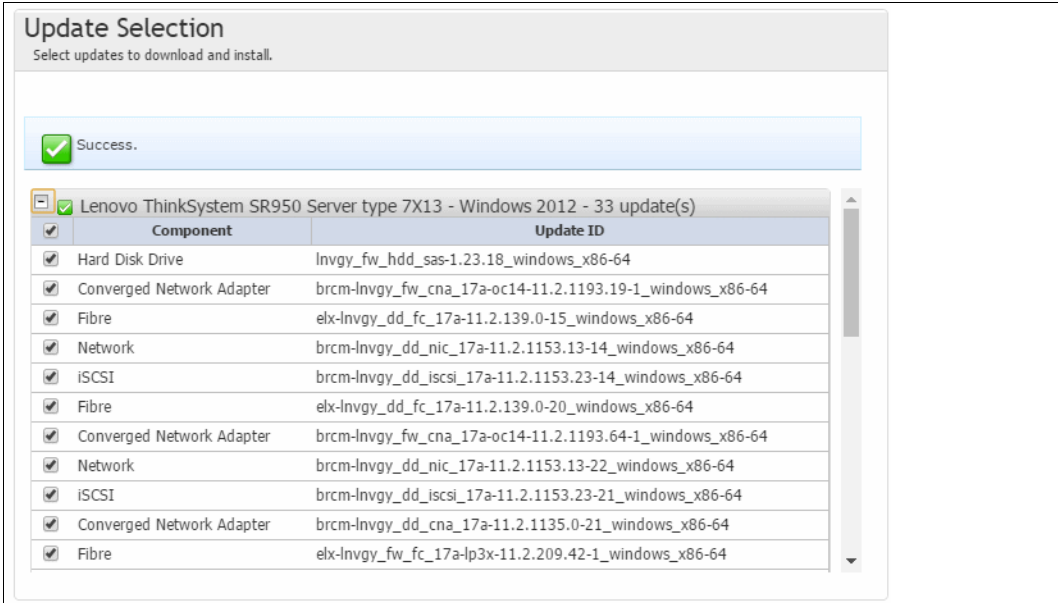


Figure 5-55 Available updates for the SR950

12. UpdateXpress will retrieve the appropriate files to update the server, Figure 5-56.




Figure 5-56 Acquiring selected updates for the SR950

- 13.If the download is successful, the Acquire Updates screen will display a success message, Figure 5-57.




Figure 5-57 Downloads available

- 14.When the download shows as successful, select **Next**.
- 15.The Finish screen will appear as seen in Figure 5-58. You can either view the logs of the download or select **Finish**. The program will then exit.




Figure 5-58 Successful completion screen

- 16.Restart the system to complete the update process.

5.3.2 XClarity Essentials Bootable Media Creator (BoMC)

The Lenovo XClarity Essentials Bootable Media Creator (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.

You can also create a Lenovo ServerGuide™ Setup and Installation CD for deploying Windows operating systems and updates on supported systems. For additional information on installing an operating system, refer to Chapter 6, “Operating system installation” on page 183.

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 Lenovo. Any firewalls must allow these ports so that BoMC can retrieve the files from the Lenovo support website.

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 the *Lenovo XClarity Essentials BoMC User 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.
2. Download the latest version of BoMC from the above web page and save it in the BoMC folder.
3. From a command line, enter the appropriate command to start the BoMC. The command depends on the operating system, as listed in Table 5-4.

Table 5-4 Commands for installing operating systems

Operating system	Command name
Microsoft Windows 7	lnvgy_utl_bomc_11.0_windows_i386.exe
Microsoft Windows 8	lnvgy_utl_bomc_11.0_windows_i386.exe
Microsoft Windows 2008 R2	lnvgy_utl_bomc_11.0_windows_i386.exe
Microsoft Windows 10	lnvgy_utl_bomc_11.0_windows_i386.exe
Microsoft Windows Server 2012	lnvgy_utl_bomc_11.0_windows_i386.exe
Microsoft Windows Server 2012 R2	lnvgy_utl_bomc_11.0_windows_i386.exe
Microsoft Windows Server 2016	lnvgy_utl_bomc_11.0_windows_i386.exe
Red Hat Enterprise Linux7.0 64-bit	lnvgy_utl_bomc_11.0_rhel7_x86-64.bin
SUSE Linux Enterprise Server 12 64-bit	lnvgy_utl_bomc_11.0_sles12_x86-64.bin

4. The Welcome window opens, as shown in Figure 5-59. Click **Next**.

Lenovo XClarity Essentials Bootable Media Creator

Lenovo

1. Welcome

2. Targeted Systems

3. Media Purpose

4. Acquire Location

5. HTTP Proxy

6. Target Directory

7. Media Format

8. Unattended Mode

9. Confirm Choices

10. Progress

11. Finish

Welcome to the Bootable Media Creator

Helps you to create bootable media for specified machine types and for selected tasks such as updating, diagnostics and deployment. You can create the media to use either a graphical user interface or text-based interface for performing the tasks. The descriptive name you enter here will be displayed on the screen when you start your bootable media, which can be useful to distinguish bootable media images that you create.

Descriptive name:

BootableMedia_20170906-033940

☐ Load in a previously saved configuration

☐ Check for the latest version of this tool

[Important notice regarding Requisite Drivers](#)

Version 11.0.15

© Copyright Lenovo 2014, 2017.

Previous Next

Figure 5-59 XClarity Essentials Bootable Media Creator (BoMC) welcome screen

5. Select the platform that you want to create bootable media for as shown in Figure 5-60 on page 167. Click **Next**.

1. Welcome

2. Targeted Systems

3. Acquire Location

4. HTTP Proxy

5. Target Directory

6. Media Format

7. Unattended Mode

8. Confirm Choices

9. Progress

10. Finish

Targeted Systems

Specify what systems this bootable media should support. Please notice that you can only select systems in the same category.

ThinkSystem Platforms

<input type="checkbox"/>	Lenovo ThinkSystem ST550 Server	7X10	<input checked="" type="checkbox"/>
<input type="checkbox"/>	Lenovo ThinkSystem SR950 Server	7X11	<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/>	Lenovo ThinkSystem SR950 Server	7X12	<input checked="" type="checkbox"/>
<input type="checkbox"/>	Lenovo ThinkSystem SR950 Server	7X13	<input checked="" type="checkbox"/>
<input type="checkbox"/>	Lenovo ThinkSystem SR850 Server	7X18	<input checked="" type="checkbox"/>
<input type="checkbox"/>	Lenovo ThinkSystem SR850 Server	7X19	<input checked="" type="checkbox"/>
<input type="checkbox"/>	Lenovo ThinkSystem SN550 Server	7X16	<input checked="" type="checkbox"/>
<input type="checkbox"/>	Lenovo ThinkSystem SN550 Server	7X17	<input checked="" type="checkbox"/>

System X Platforms

<input type="checkbox"/>	Product Family	Machine Type	Updates
<input type="checkbox"/>	BladeCenter HX5	1909	<input checked="" type="checkbox"/>
<input type="checkbox"/>	BladeCenter HX5	1910	<input checked="" type="checkbox"/>
<input type="checkbox"/>	BladeCenter HS22	1911	<input checked="" type="checkbox"/>
<input type="checkbox"/>	BladeCenter Server HS23	1929	<input checked="" type="checkbox"/>
<input type="checkbox"/>	BladeCenter HS22	1936	<input checked="" type="checkbox"/>
<input type="checkbox"/>	BladeCenter HS22V	1949	<input checked="" type="checkbox"/>
<input type="checkbox"/>	System x3100 M4	2582	<input checked="" type="checkbox"/>

Update List

Roll Back

Previous

Next

Figure 5-60 Platform selection for BoMC

6. Select the location of the updates that will be applied to the server, Figure 5-61.

If updates are local, select the path where they are located. If they need to be downloaded, select either the System Pack or individual updates to download. In this scenario, individual updates has been selected, however the recommended path is to use the System Pack to ensure all firmware gets updated. Click **Next**.

1. Welcome

2. Targeted Systems

3. Acquire Location

4. HTTP Proxy

5. Target Directory

6. Media Format

7. Unattended Mode

8. Confirm Choices

Acquire Location

You can choose to acquire updates, tools and bootable image from Lenovo support web site or from local directory.

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

Browse

Figure 5-61 Acquiring appropriate updates

7. If access to the Internet is required, enter in the Proxy details as shown in Figure 5-62. If access to the Internet is already available, select **Do not use a proxy** and click **Next**.

The screenshot shows a wizard interface with a sidebar on the left containing seven steps: 1. Welcome, 2. Targeted Systems, 3. Acquire Location, 4. HTTP Proxy (highlighted in blue), 5. Target Directory, 6. Media Format, and 7. Unattended Mode. The main content area is titled 'HTTP Proxy' and contains the following text: '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 Lenovo support web site.' Below this text are two radio buttons: 'Do not use proxy' (selected) and 'Use proxy'. The 'Use proxy' option is followed by input fields for 'Host Name' and 'Port'. There is also an unchecked checkbox for 'Authenticate using the following credentials', which is followed by input fields for 'User Name' and 'Password'. At the bottom of the main area is a 'Test Connection' button.

Figure 5-62 Proxy details

8. Select the file that will host the firmware that will be downloaded by BoMC from the Lenovo support site, as shown in Figure 5-63. Click **Next**.

The screenshot shows a wizard interface with a sidebar on the left containing four steps: 1. Welcome, 2. Targeted Systems, 3. Acquire Location, and 4. HTTP Proxy (highlighted in blue). The main content area is titled 'Target Directory' and contains the following text: '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.' Below this text is a 'Directory:' label followed by a text input field containing 'C:\BoMC\workingdir'. At the bottom of the main area is a 'Browse' button.

Figure 5-63 Select target directory

9. Select which type of media will be created (CD/DVD, USB or PXE) and a location to write the media to (directly to a device or locally). In this example, the media will be written locally. Click **Next**.

The screenshot shows a wizard interface with a sidebar on the left containing five steps: 1. Welcome, 2. Targeted Systems, 3. Acquire Location, 4. HTTP Proxy, and 5. Target Directory (highlighted in blue). The main content area is titled 'Media Format' and contains the following text: '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.' Below this text is a 'Device type:' label followed by a dropdown menu showing 'CD/DVD'. There are two radio buttons: 'Write directly to device' and 'Write to image file' (selected). The 'Write to image file' option is followed by a text input field containing 'C:\BoMC\workingdir\bootable.iso'. At the bottom of the main area is a 'Browse' button.

Figure 5-64 Type of media to be created

10. If you are using a TFTP or FTP server you can perform updates in unattended mode as seen in Figure 5-65. Otherwise, select **Do not use unattended mode**. Click **Next**.

1. Welcome

2. Targeted Systems

3. Acquire Location

4. HTTP Proxy

5. Target Directory

6. Media Format

7. Unattended Mode

Unattended Mode Configuration

Helps you to create 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

Specify a directory to save the unattended log files. This directory should start from the root. Also, when uploading to TFTP server, FTP server, NFS and Samba server, make sure the directory has been created and the anonymous access granted.

Figure 5-65 Optional unattended mode for updates

Note: If using unattended mode, you can upload the log files to one of the following:

- TFTP Server
- FTP Server
- NFS Share
- Samba Server
- USB Drive

If you are using RHEL 6.4 x64 operating system environment, then use the following commands:

- For Samba shares:

```
mount -t cifs //$address/$directory $mount_point -o user=anonymous, password=anonymous
```

- For NFS shares:

```
mount -t nfs $address:$directory $mount_point
```

- For FTP and TFTP uploads, ensure that the local file can be uploaded to the address and directory specified in the FTP or TFTP command.

11. Confirm the selected choices on the Confirm Choices screen as shown in Figure 5-66 on page 170. If changes are needed at this point, click the **Back** button and make the appropriate changes. If the selections are correct, click **Next** to continue.

Alternatively, you can save the configuration you have just created to be imported into the tool at a later date.

Confirm Choices

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

Machine Type	7X12
Update Location	Lenovo Support Web Site
Update Type	Latest Individual updates
Proxy IP Address	NO
Proxy Port Number	NO
Proxy User Name	NO
Media Purpose	Updates
Enable Task Auto Run	NO
Image File Name	C:\BoMC\workingdir\bootable.iso
Target Directory	C:\BoMC\workingdir
Overwrite Existing Image	NO
Include VMWare ESXI Update	NO
Use Unattend Mode	NO
Descriptive Name	BootableMedia_20170906-033940

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


Figure 5-66 Confirmation of selections

12. You will then see the progress of the downloads, similar to Figure 5-67. Use the + button next to entry expand each option and view the downloads.

1. Welcome
2. Targeted Systems
3. Acquire Location
4. HTTP Proxy
5. Target Directory
6. Media Format
7. Unattended Mode

Creation Progress

Your bootable media is being created.


Acquiring updates: Invgy_utl_lxce_ux_2.0.0_anyos_x86-64
16,574 of 221,803KB (2,071KB/Sec)
Acquiring 3 of 4 package bundle(s)

+ ☒ (type 7X12) - Updates - 14 update(s)
+ ☐ (type Requisite Package) - Updates - (No additional requisite required [Details...](#))
+ ☐ (type Tool) - UPDATE - 1 update(s)
☐ (type Tool) - Bootable Environment -

Figure 5-67 Progress screen for firmware download

13. Once the downloads have completed, BoMC will start the creation of bootable media, as seen in Figure 5-68 on page 171. This may take a moment to complete depending on the size of the download.




Figure 5-68 Select the devices that require firmware updates

14. After the creation has finished, click **Next** and then **Finish**. The ISO has now been created locally. You now have a bootable image with the updates.

You can mount the ISO image through the XCC as discussed in 6.2.1, “Installing an OS with XClarity Controller (XCC)” on page 185.

5.3.3 XClarity Essentials OneCLI

OneCLI is a command-line utility that is ideal for a non-GUI environment. It provides a command line interface to the same functions available with XClarity Essentials UpdateXpress, as discussed in 5.3.1, “XClarity Essentials UpdateXpress” on page 159.

OneCLI is available to download from the following support page:

<https://datacentersupport.lenovo.com/documents/lnvo-tcli>

OneCLI replaces the Advanced Settings Utility (ASU) used in previous generations of servers. The supported operating systems are listed in Table 5-5.

Table 5-5 Supported OS versions with OneCLI

Supported OS	OS Version
Microsoft Windows	Server 2016 (x86-64)
	Server 2012 R2 (x86-64)
	Server 2012 (x86-64)
	Server 2008 R2 (x86-64)
	Server 2008 (x86-32/x86-64)
	Microsoft Windows 7/8/10 Desktop Editions (for Acquire and Remote update function)
Red Hat	Red Hat Enterprise Linux 7 Server (x64) Editions (up to U3)
	Red Hat Enterprise Linux 6 Server (x86 & x64) Editions (up to U9)
	Red Hat Enterprise Linux 5 Server (x86 & x64) Editions (up to U10)

Supported OS	OS Version
SUSE	SUSE Linux Enterprise Server 12 (x64) (SP2)
	SUSE Linux Enterprise Server 11 (x86 & x64) (SP4)
	SUSE Linux Enterprise Server 10 (x86 & x64) (SP4)
Customized VMware ESXi	ESXi 6.5
	ESXi 6.0 (up to U3)
	ESXi 5.5 (up to U3)

Follow these steps to extract OneCLI for use:

1. Go to the above web page and download the package for the system you'll be running the tool on to a directory on your C drive or a memory key
2. Unzip the zip or tgz to the directory
3. Change to that directory
4. Run OneCli.exe (Windows) or ./OneCli (Linux)

XClarity Essentials OneCLI is now ready to use.

Before applying updates to the SR950 server using OneCLI, download the latest packages to apply to the server. This can be achieved using OneCLI commands from the command prompt. Typically, this would mean running one single OneCLI command to acquire the latest UXSP update (recommended update stack) from the Lenovo support site using the machine type (for SR950 use either: 7X11, 7X12 or 7X13) and operating system being used.

For example: To acquire the latest UXSP to c:\SR950\pkgs\, which can then be updated/installed to Windows Server 2012 R2, an SR950 type 7X11, use the following command:

- ▶ Linux:


```
./OneCli update acquire --mt 7X11 --ostype win2012r2 --dir c:\SR950\pkgs
```
- ▶ Windows

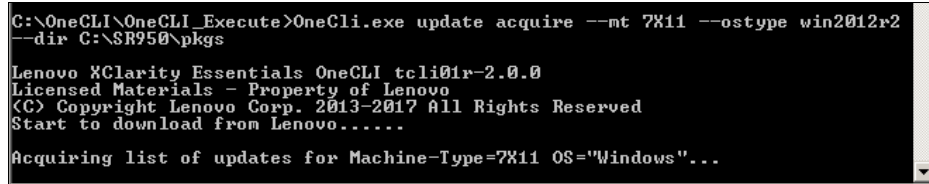

```
OneCli update acquire --mt 7X11 --ostype win2012r2 --dir c:\SR950\pkgs
```

Figure 5-69 on page 173 shows the download command. Valid choices for otype are:

- ▶ win2008, win2012, win2012r2, win2016
- ▶ rhel5, rhel6, rhel7
- ▶ sles10, sles11, sles12
- ▶ esxi5.0, esxi5.1, esxi5.5, esxi6.0, esxi6.5
- ▶ platform
- ▶ none

Tips:

- ▶ The otype none is used for operating system independent operations, such as hardware out-of-band updates.
- ▶ You can acquire updates for a OS different from the one where OneCLI runs. For example, you can run OneCLI on Windows to acquire a UXSP package for RHEL7.



```
C:\OneCLI\OneCLI_Execute>OneCli.exe update acquire --mt 7X11 --ostype win2012r2
--dir C:\SR950\pkgs

Lenovo XClarity Essentials OneCLI tcli01r-2.0.0
Licensed Materials - Property of Lenovo
(C) Copyright Lenovo Corp. 2013-2017 All Rights Reserved
Start to download from Lenovo.....

Acquiring list of updates for Machine-Type=7X11 OS="Windows"...
```

Figure 5-69 download updates with OneCLI

The command syntax is as follows:

OneCli <application> <command> [command option] [connect option]

To update the packages downloaded from the example above within the host OS, perform the following:

- ▶ For Linux
./OneCli update flash --dir c:\SR950\pkgs
- ▶ For Window
OneCli update flash --dir c:\SR950\pkgs

For additional information on OneCLI, including flashing the server out of band remotely, refer to the OneCLI User's Guide:

http://sysmgt.lenovofiles.com/help/topic/toolctr_cli_lenovo/onecli_r_sysxflexbc.html

Note: Other tools can be used to update firmware, but they are not in the scope of this guide. For example, users can update device drivers through the operating system. Windows Server and VMware have their own way to update drivers individually and on a large scale.

5.3.4 Lenovo XClarity Administrator

With Lenovo 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.

Note: 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 servers that are being managed by XClarity Administrator:

1. Click **Provisioning** → **Apply / Activate** under the Firmware Updates section, as shown in Figure 5-70 on page 174.

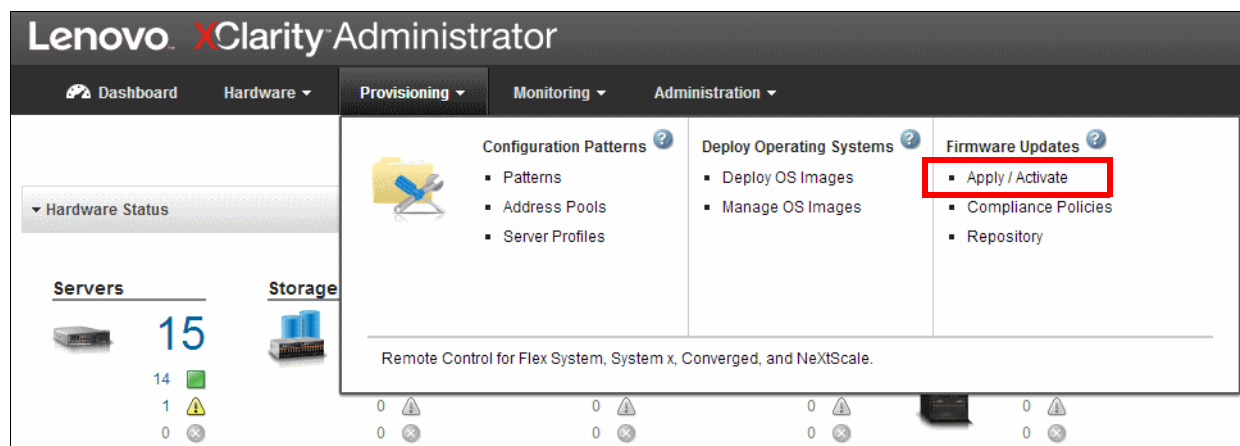


Figure 5-70 Apply/Activate firmware with LXCA

2. In the Firmware Updates: Apply / Activate window, select the required hardware components to update, and click the **Perform Updates** icon, as shown in Figure 5-71.

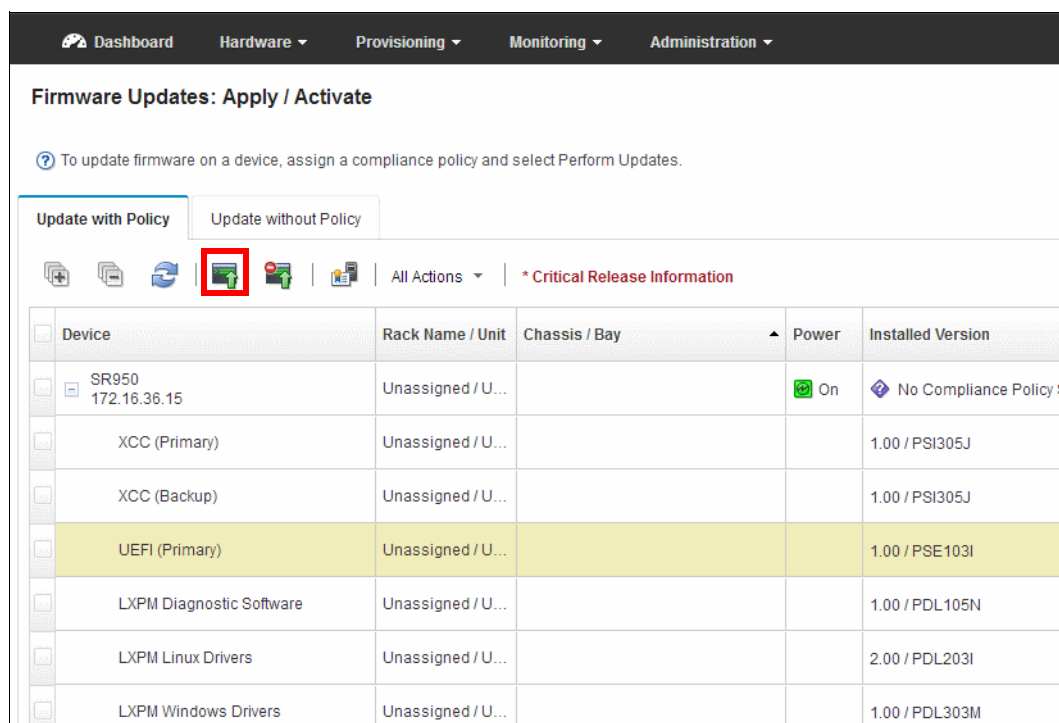


Figure 5-71 Apply updates

3. An Update Summary window opens, in which you can set Update Rules and Activation Rules. As shown in Figure 5-72 on page 175, 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.

Update Summary
Select your Update Rule and review your updates. Then click Perform Update.

Note: The update job will run in the background and might take several minutes to complete. Updates are performed as a job. You can go to the [Jobs](#) page to view the status of the job as it progresses.

* Update Rule: Stop all updates on error ?

* Activation Rule: Immediate activation ?

☐ Attempt to update components already in compliance

All Actions ▾ Filter

System	Chassis / Bay	Installed Version	Compliance Target
SR950 172.16.36.15			
Diagnostics		10.0 / DSALA2N	DSALA6L-10.1 / DSALA6L Invgv_fw_dsa_dsala6l-10.1_...

Simulate Update
Perform Update
Close

Figure 5-72 Firmware update parameters

- To start flashing the firmware, click **Perform Update**. When the Immediate activation option is chosen, you must confirm the operation by clicking **OK**, as shown in Figure 5-73.

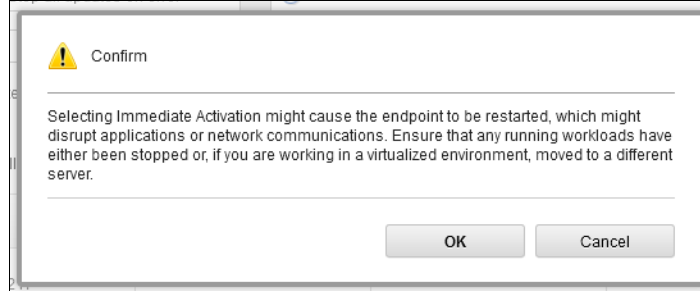


Figure 5-73 Confirmation for flashing firmware

- 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.

For additional information on XClarity Administrator, refer to 7.2, “Lenovo XClarity Administrator” on page 225.

5.4 Overview of firmware deployment tools: Using the right tool

The tools listed previously can be used to update most firmware for your server and the devices that are installed in the server.

Note: Lenovo typically releases firmware in bundles called UpdateXpress System Packs (UXSPs). To ensure that all of the firmware updates are compatible, you should update all firmware at the same time. If you are updating firmware for both Lenovo XClarity Controller and UEFI, update the firmware for Lenovo XClarity Controller first.

If you are updating firmware individually, it is recommended to do so in the following order:

1. Update XCC and reboot XCC
2. Update UEFI and reboot system to F1
3. Update LXPM
4. Update LXPM driver packs (Windows and Linux)
5. Update RAID adapter followed by any additional PCIe adapter firmware and reboot system
6. Update onboard Ethernet firmware via OS.
7. Reboot the server

Important: Ensure all firmware levels are at the latest levels. Do not update one device without updating all devices. This will ensure compatibility between all devices. The easiest and safest way to do this is to use the UXSPs provided by Lenovo.

There are several methods used by the Lenovo tools to update the server's firmware. These terminologies are discussed below.

► In-band update:

The installation or update is performed using a tool or application within an operating system that is executing on the server's core CPU.

► Out-of-band update:

The installation or update is performed via the XClarity Controller service processor, collecting the update, and then directing the update to the target subsystem or device. Out-of-band updates have no dependency on an operating system executing on the core CPU. However, most out-of-band operations do require the server to be in the S0 (Working) power state.

► On-Target update (local update):

The installation or update is initiated from an operating system executing on the server's operating system.

► Off-Target update: (remote update)

The installation or update is initiated from a computing device interacting directly with the server's XClarity Controller service processor.

► UpdateXpress System Packs (UXSPs):

UXSPs are bundled updates designed and tested to provide the interdependent level of functionality, performance, and compatibility. UXSPs are server machine-type specific and are built (with firmware and device driver updates) to support specific Windows Server, Red Hat Enterprise Linux (RHEL) and SUSE Linux Enterprise Server (SLES) operating system distributions. Machine-type-specific firmware-only UXSPs are also available.

See the following table to determine the best suited Lenovo tool to use for installing the firmware for the SR950 server.

Table 5-6 Comparison of Lenovo firmware deployment tools

Tool	In-band update	Out-of-band update	On-target update	Off-target update	GUI	CLI	Supports UXSPs
Lenovo XClarity Provisioning Manager Limited to core system firmware only. ^a	Yes	No	No	Yes	Yes	No	Yes
Lenovo XClarity Controller Supports core system firmware and most advanced I/O option firmware updates	No	Yes	No	Yes	Yes	Yes	No
Lenovo XClarity Essentials OneCLI Supports all core system firmware, I/O firmware, and installed operating system driver updates	Yes	Yes	No	No	No	Yes	Yes
Lenovo XClarity Essentials UpdateXpress Supports all core system firmware, I/O firmware, and installed operating system driver updates	Yes	Yes	No	No	Yes	No	Yes
Lenovo XClarity Administrator Supports core system firmware and I/O firmware updates	Yes	Yes	No	Yes	Yes	No	No

a. Core system firmware refers to the XClarity Controller (XCC) firmware and the XClarity Provisioning Manager (LXPM) / UEFI firmware

The firmware for the SR950 server can be downloaded from the following link:

<https://datacentersupport.lenovo.com/no/en/products/servers/thinksystem/sr950/7x12/downloads>

5.5 Backing up the server configuration

Once you have set up the server or if you have made changes to the setup, it is recommended that you back up the server configuration.

You can back up the following server components:

- Baseboard Management Controller (BMC)

You can back up the management controller configuration through the Lenovo XClarity Controller (XCC) interface. To back up the BMC with XCC go to **BMC Configuration** → **Backup and Restore** → **Backup BMC Configuration**, as seen in Figure 5-74 on page 178.

Lenovo SR950 System name: Export USERID 10:06 AM

▼ Backup BMC Configuration ?

1 Set Password 2 Backup Process 3 Download Backup File

Note that you will be asked for this password when you use the file to restore a configuration.

Set password ! Confirm password ! Start Backup Encrypt whole file Encrypt only sensitive data

Last backup: No backup has been requested

► Restore BMC from Configuration File ?

Last restore: No restore has been requested

► Reset BMC to factory default

Figure 5-74 Backup and Restore for BMC

Alternatively, you can use the **save** command from Lenovo XClarity Essentials OneCLI to create a backup of all configuration settings. For more information about the save command, see:

http://toolscenter.lenovofiles.com/help/topic/toolctr_cli_lenovo/onecli_r_save_command.html

- Operating system

Use your own operating-system and user-data backup methods to back up the operating system and user data for the server.

5.6 Tethering with XClarity Mobile

You can use the Lenovo XClarity Mobile application on Android or iOS mobile devices to monitor managed devices in one or more XClarity Administrator instances. The XClarity Mobile app is designed for “anywhere management” for Lenovo servers.

It allows administrators to check system inventory and status, alerts, events, jobs and notifications, and perform power actions on managed systems. Using the XClarity Mobile app, you can perform the following activities:

- Configure network settings and properties
- View the status summary of each connected XClarity Administrator.
- View the status summary of all managed devices.
- Display graphical views (maps) for chassis, rack servers, and storage devices.
- Monitor the detailed status of each managed device.
- Monitor the inventory of each managed device.
- Monitor audit events, hardware and management events, alerts, and jobs.
- Turn on or off the location LED on a managed device.
- Power on, power off, or restart a managed device.

- ▶ Trigger the collection of diagnostic data.
- ▶ Set up automatic problem notification through Call Home.
- ▶ Push event notifications to your mobile device (see Forwarding events to mobile devices).
- ▶ Send feedback about this mobile app to Lenovo Support.
- ▶ Connect your mobile device directly to a ThinkSystem server to manage the server using the XClarity Mobile app (for devices that support USB tethering)

The interface of the XClarity Mobile app is shown in Figure 5-75.

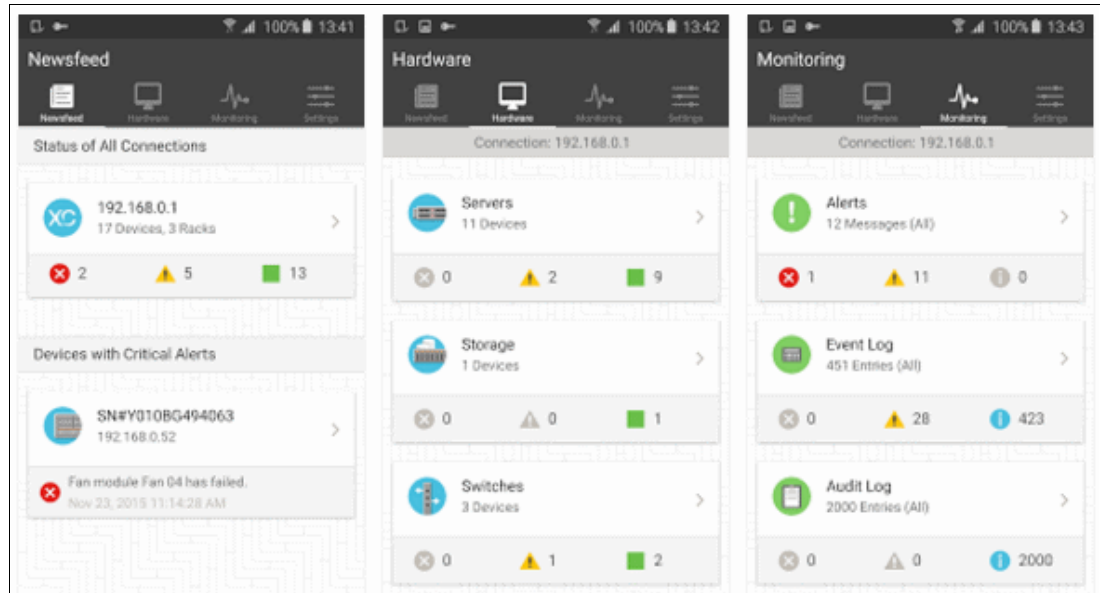


Figure 5-75 XClarity Mobile app interfaces

For additional information on monitoring the SR950 with the XClarity Mobile app, refer to 5.6.3, “Connecting XClarity Mobile directly to the SR950 server” on page 181.

You can also connect your mobile device directly to the server and then launch the Lenovo XClarity Mobile app and log in to the server’s baseboard management controller (BMC) using the same web and CLI credentials. This is discussed further in 5.6.3, “Connecting XClarity Mobile directly to the SR950 server” on page 181.

5.6.1 XClarity mobile app pre-requisites

The following are the pre-requisites for using the mobile app with the SR950 server.

- ▶ iOS tablets are supported at iPhone screen-resolution only. Android tablets are not currently supported.
- ▶ The following mobile operating systems are supported:
 - Android 5.0 or later
 - iOS 9, 10, or later

Note: iOS 8 is supported only for Lenovo XClarity Mobile v1.3.0 and earlier.

- ▶ Lenovo XClarity Mobile can connect to Lenovo XClarity Administrator 1.2.1 or later.
- ▶ Ensure that a network connection is available from your mobile device to the Lenovo XClarity Administrator instances. This might require the use of a VPN solution. See your network administrator for assistance.

- Import the CA certificate for each Lenovo XClarity Administrator instance.

Note: All connections to Lenovo XClarity Administrator use HTTPS. However, there must be a valid certificate chain before the connection is considered trusted and data can be passed to the mobile device. To create a trusted certificate chain, you must import the Lenovo XClarity Administrator self-signed certificate authority (CA) into the mobile device.

To import the self-signed CA certificate into the mobile device, follow these steps:

1. Download the CA certificate to a local system:
 - a. Connect to the XClarity Administrator instance using a web browser on your local system.
 - b. From the XClarity Administrator menu bar, click **Administration** → **Security** to display the Security page.
 - c. Click **Certificate Authority** under the Certificate Management section. The Certificate Authority page is displayed.
 - d. Click **Download Certificate Authority Root Certificate**.
 - e. Click **Save as der** or **Save as pem** to save the CA certificate as a DER or PEM file on your local system. PEM format works in most cases.
2. Transfer the CA certificate file to your mobile device, for example, via attaching it as an email you can view from the device, or through cloud storage (such as Dropbox), or via file transfer to the device.
3. Import the trusted CA certificate:
 - For Android users: Typically this is done by selecting **Settings** → **Security** → **Install from phone storage**, and then selecting the certificate file that you downloaded.
 - For iOS users: Open the email on your mobile device, and click the document link in the email to import the trusted CA certificate.

Note: For iOS 10.3 and later, imported certificates are not trusted by default. To trust the certificates, select **Settings** → **General** → **About** → **Certificate Trust Settings**, and then enable the certificate trust.

5.6.2 Installing and setting up XClarity Mobile app

To install and setup the XClarity Mobile app, follow these steps:

1. Download the XClarity Mobile app from iTunes App Store (iOS) or Google Play Store (Android).
2. Install the application by following the instructions on the mobile device.

Note: A mobile OS-level security code to unlock screen access is required to use the XClarity Mobile app. If one is not already set up, you are instructed to set one up during installation.

Once the application has been installed:

1. Open the application from your device.
2. Set a PIN code for the app.
3. Click Settings.
4. Add or edit connections to multiple XClarity Administrator instances using the automatic discovery or by providing an IP address and user credentials.

Uninstall old version first: For Android devices, you must uninstall Lenovo XClarity Mobile app v1.3.1 and earlier before installing v1.3.2.

5.6.3 Connecting XClarity Mobile directly to the SR950 server

As previously mentioned, you can connect the XClarity Mobile application directly to the SR950 server using a USB cable.

The actions available via a menu on the mobile app directly connected to the server include:

- ▶ Service
 - Sharing summary information using email or other means that is provided by the mobile device
 - Clearing the event and audit log
 - Downloading the event and audit log to the mobile-device local storage or transmitting the log by any means that is provided by the mobile device
 - Configuration and Initial Setup
 - Configuring server properties, such as location and contact information for initial setup
 - Viewing and changing the IPv4 and IPv6 BMC network interface settings
 - Specifying boot order and one-time boot settings
 - Instructing a selected Lenovo XClarity Administrator to manage the server after initial setup is complete
- ▶ Power Actions
 - Powering the server on or off, restarting the server, or triggering NMI
 - Resetting the BMC

To connect the server and mobile device together, complete the following steps to start managing and configuring the server:

1. Enable the front USB port for XCC access by performing one of the following actions:
 - Press and hold the blue ID location LED on the front panel of the server for at least 3 seconds until the light blinks every couple of seconds.
 - From the XCC CLI, run the **usbfp** command
 - From the XCC web interface, click **BMC Configuration** → **Network** → **Front Panel USB Port Management**.
2. Connect your phone USB cable to the front panel USB port on the SR950 server, as indicated in Figure 5-76 on page 182.




Figure 5-76 Connecting to the USB port

3. On your mobile device, enable USB tethering.
 - a. For iOS, click **Settings** → **Cellular** → **Personal Hotspot**.
 - b. For Android, click **Settings** → **Mobile hotspot and tethering** → **USB tethering**.
4. On your mobile device, launch the Lenovo XClarity Mobile app.
5. If automatic discovery is disabled, click **Discovery** on the USB Discovery page to connect to the server's XCC management controller and collect information, including inventory, health, firmware, network configuration, and a list of the latest active events.

Operating system installation

This chapter provides an overview of the OS installations for the SR950 server.

Topics in this chapter are:

- ▶ 6.1, “Options for OS installation” on page 184
- ▶ 6.2, “Installing an OS with Lenovo tools” on page 185
- ▶ 6.3, “Operating system installation examples” on page 200
- ▶ 6.4, “Installing an OS on the M.2 adapter” on page 220
- ▶ 6.5, “Bootting from SAN” on page 221

At the time of this writing, the SR950 server is supported by following operating systems:

- ▶ Microsoft Windows Server 2012 R2
- ▶ Microsoft Windows Server 2016
- ▶ Microsoft Windows Server 2019
- ▶ Microsoft Windows Server, version 1709
- ▶ Microsoft Windows Server, version 1803
- ▶ Red Hat Enterprise Linux 6.10 x64
- ▶ Red Hat Enterprise Linux 6.9 x64*
- ▶ Red Hat Enterprise Linux 7.3
- ▶ Red Hat Enterprise Linux 7.4
- ▶ Red Hat Enterprise Linux 7.5
- ▶ Red Hat Enterprise Linux 7.6
- ▶ SUSE Linux Enterprise Server 11 x64 SP4
- ▶ SUSE Linux Enterprise Server 12 SP2
- ▶ SUSE Linux Enterprise Server 12 SP3
- ▶ SUSE Linux Enterprise Server 12 SP4
- ▶ SUSE Linux Enterprise Server 12 Xen SP4
- ▶ SUSE Linux Enterprise Server 15
- ▶ SUSE Linux Enterprise Server 15 Xen
- ▶ 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

*Note: Support for RHEL 6.9 on the SR950 is limited to up to 4 processors. Configurations greater than 4 processors are not supported.

For the latest information about the specific versions and service levels that are supported and any other prerequisites, see the Operating System Interoperability Guide:

<https://lenovopress.com/osig#servers=sr950-7x11-7x12-7x13>

6.1 Options for OS installation

There are several options that are available to install an operating system (OS) on the SR950 server.

The main methods discussed here include using:

- ▶ **Lenovo XClarity Controller (XCC)**
The XCC can be used to remotely mount and install an operating system on the server. This is ideal for installing OSes on a single server.
- ▶ **Lenovo XClarity Provisioning Manager**
Lenovo XClarity Provisioning Manager (LXPM) can also be used to install an operating system on a single server.
- ▶ **Lenovo XClarity Administrator**
Lenovo XClarity Administrator is ideal for installing and managing multiple servers at once. Up to 28 server OSes can be deployed concurrently using this tool.
- ▶ **XClarity Essentials Bootable Media Creator**
XClarity Essentials Bootable Media can be used to create bootable images on several forms of media, including CD, DVD, ISOimage, PXE files, and USB flash drive.
- ▶ **Local USB port**
Use the SR950's local USB port on the front of the server to install the OS.

Table 6-1 lists the Lenovo OS deployment tools for installing an OS on the SR950 server. This table can help determine which tool may be best suited to install an OS for a particular situation.

Table 6-1 Comparison of Lenovo OS deployment tools

	XClarity Controller (XCC)	XClarity Provisioning Manager	XClarity Administrator	XClarity Essentials Bootable Media Creator	Local USB port
Remote installation available	Yes	Yes	Yes	Yes	No
Physical presence required	No	No	No	No	Yes
Graphical user interface	Yes	Yes	Yes	Yes	Yes
Command line interface	Yes	No	No	Yes	No

Note: Before you can install an OS on the SR950 server, the RAID controller must have initialized both controller and drives.

- ▶ To install an OS with LXPM, refer to 5.1.7, “Configuring RAID with LXPM” on page 135.
- ▶ To install an OS with XCC, refer to 5.2.5, “Configuring RAID with XCC” on page 153.

6.2 Installing an OS with Lenovo tools

This section discusses the installation of an operating system using the following tools

- ▶ 6.2.1, “Installing an OS with XClarity Controller (XCC)” on page 185
- ▶ 6.2.2, “Installing an OS with Lenovo XClarity Provisioning Manager (LXPM)” on page 190
- ▶ 6.2.3, “Installing an OS with XClarity Administrator (LXCA)” on page 195
- ▶ 6.2.4, “Lenovo XClarity Bootable Media Creator (BoMC)” on page 200
- ▶ 6.2.5, “Installing an OS with the local USB port” on page 200

6.2.1 Installing an OS with XClarity Controller (XCC)

The SR950 includes the XClarity Controller (XCC) Enterprise version as standard, which enables a remote presence (KVM) and remote media (ISO and IMG image files) to install an operating system on the server.

You must be logged into the XCC with a user ID that includes Supervisor access. You can check your level of access by clicking on the ‘USERID’ at the top right hand corner of the screen, as shown in Figure 6-1.

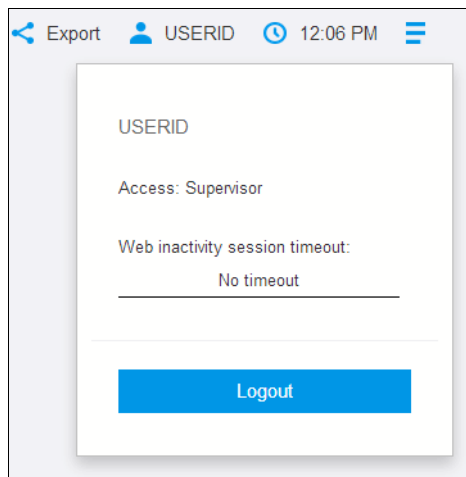


Figure 6-1 XCC access level

To log into XCC, refer to “Default local network access to XCC” on page 143.

Complete the following steps to mount an ISO and install an operating system with the XCC:

1. Connect to XCC by entering the IP address of XCC into a supported web browser. See the “Supported web browsers with XCC” on page 143 for a list of supported browsers.
2. From the XCC interface, go to **Server Configuration** → **Boot options**.
3. Highlight the CD/DVD ROM option from the **Available devices** menu and click the arrow to move it to the **Boot order you want to configure** box.
4. Click on the arrows to move the CD/DVD ROM option to the top of the menu, as shown in Figure 6-2 on page 186.

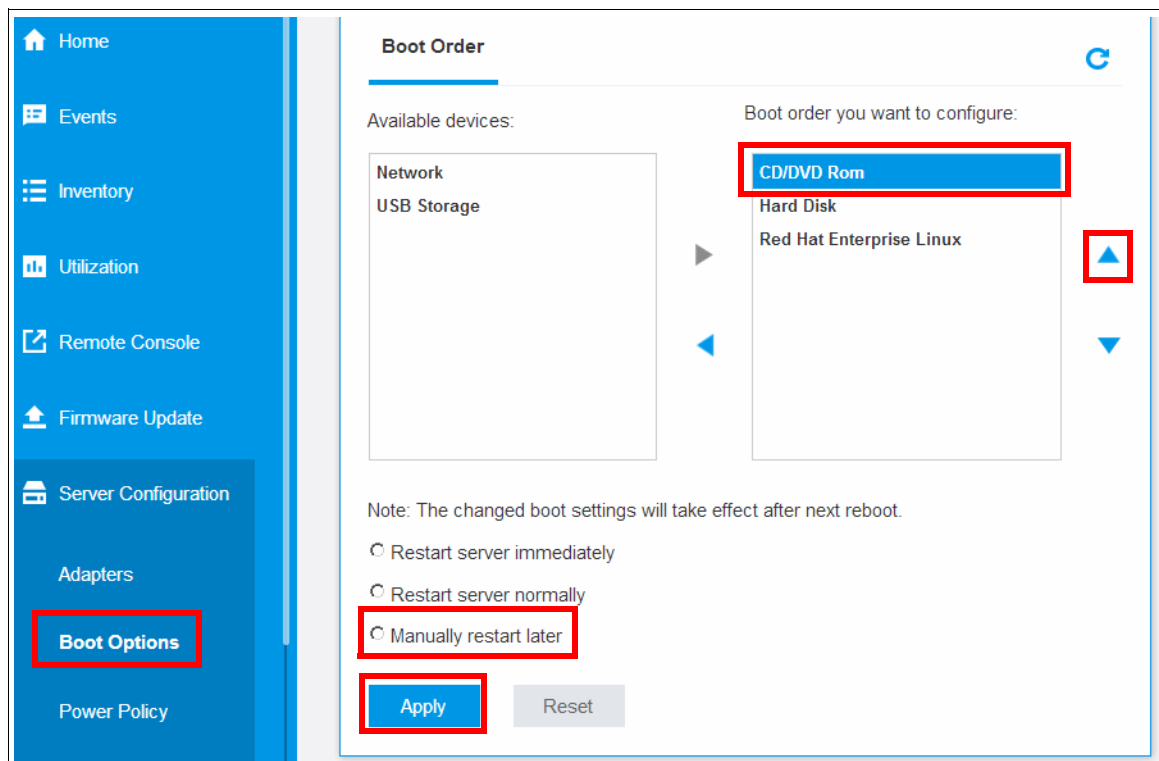


Figure 6-2 XCC boot options - boot from CD/DVD for remote ISO mount

5. Click the **Manually restart later** radio button. Click the Click **Apply** and **OK** to the confirmation prompt.

Mounting an ISO image

6. From the XCC interface, click the **Remote Control** option either via the bottom of the home page in XCC or via the left hand menu, as shown in Figure 6-3 on page 187.

Note: from the Remote Console option in XCC you can open the **Latest Failure Screen**. This is useful in the event that you need to provide additional information to Lenovo support on a failure.

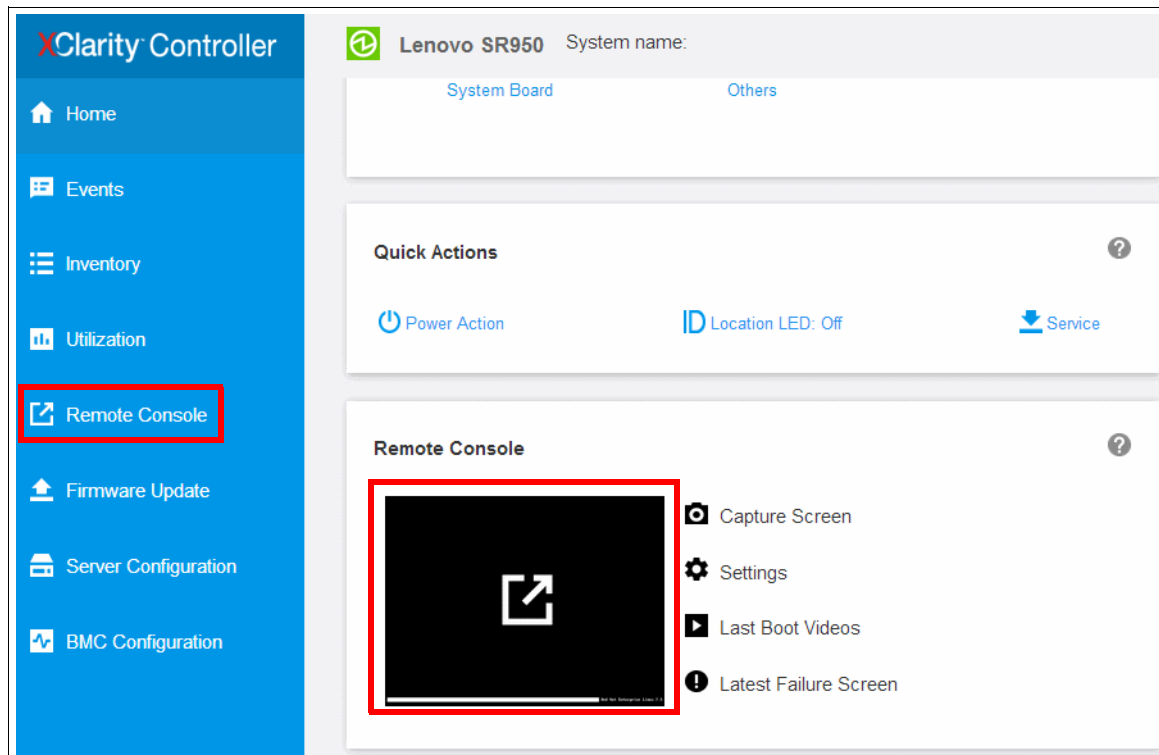


Figure 6-3 Remote console access in XCC

The remote console settings pop-up screen will appear. If you want to allow other users remote control access during your session, check the radio button for **Multi-user Mode**. Up to 6 con-current users can connect at the one time.

7. Click **Launch Remote Console**. This will launch the HTML5-based remote console.

Browser popup blocker: You will need to disable the browser pop-up blocker if it is enabled. If pop-ups are not enabled, the remote KVM session will not open.

The remote KVM session will launch as shown in Figure 6-4.

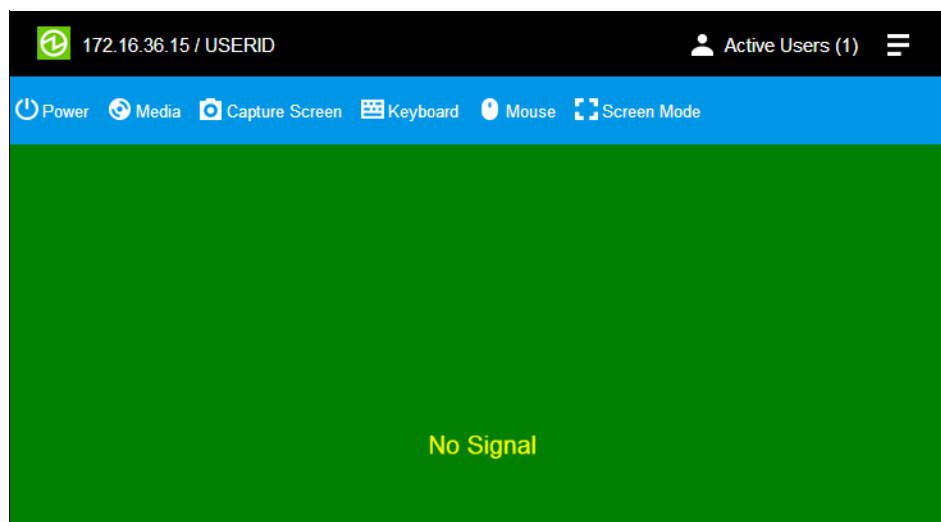


Figure 6-4 Remote KVM session

8. The ISO image will now need to be mounted. Click the **Media** option at the top of the remote console screen as shown in Figure 6-5.




Figure 6-5 Mount an ISO via the Media option

9. The virtual media mounting screen will load. There are three methods to mount an ISO:
- Locally
 - From a network (CIFS, NFS or HTTPFS)
 - From Remote Disc on Card (RDOC).

In this example, the ISO is local. Click the **Activate** button to enable mounting the ISO, as shown in Figure 6-6.

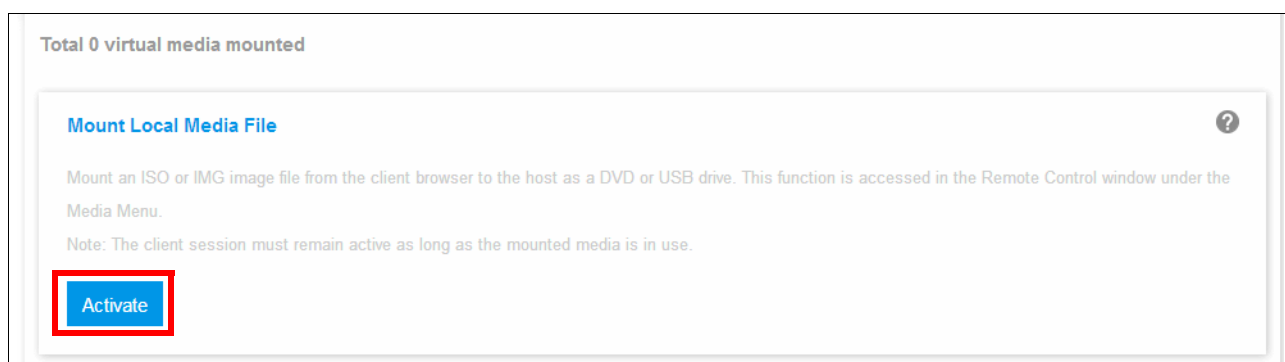


Figure 6-6 Activate mounting

10. Click the **Browse** button and browse to the local ISO to be installed on the SR950 and click **Open**.
11. Click the **Mount all local media** button as shown in Figure 6-7.

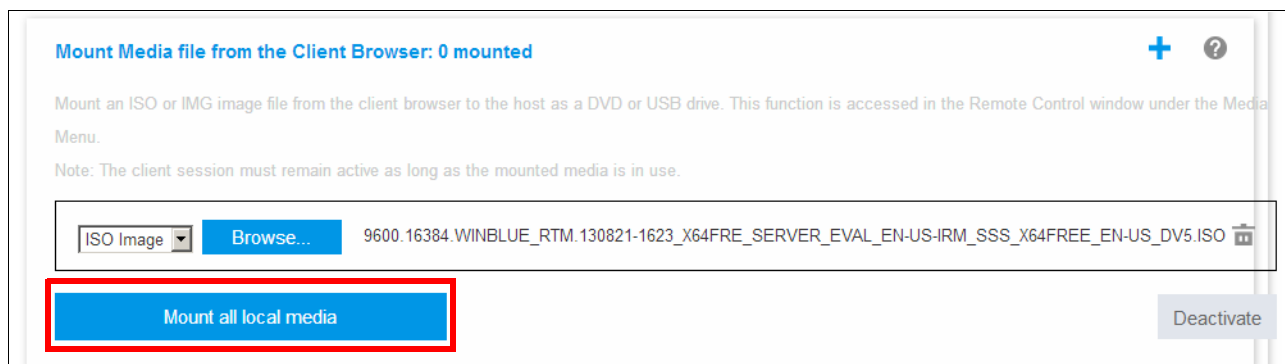


Figure 6-7 Mounting the ISO file

Once the ISO has been successfully mounted, the media will display a tick to confirm it is ready to be deployed as shown in Figure 6-8 on page 189.

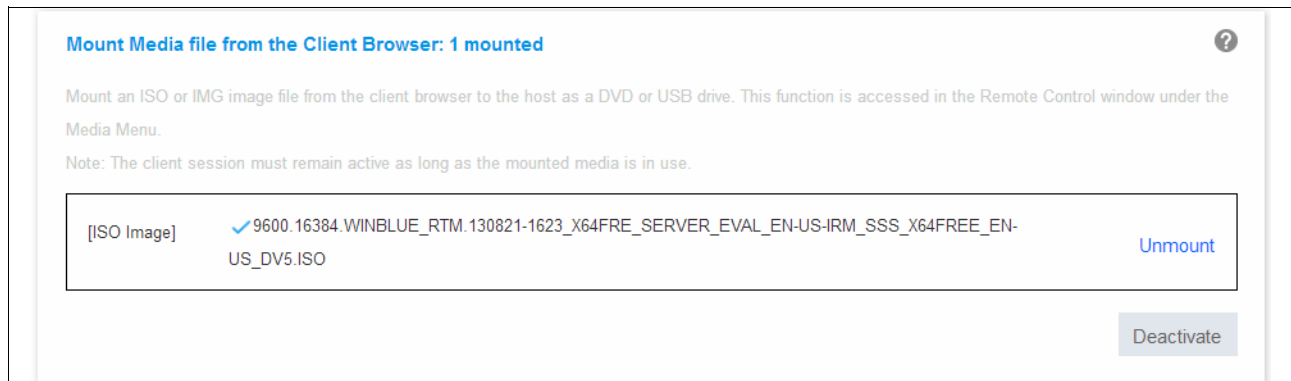


Figure 6-8 Media ready for deployment

12. Scroll to the bottom of the virtual media mounting screen and use the drop down menu under the **Select one virtual media to boot at next restart**. Select the OS you just mounted.




Figure 6-9 Select the virtual media to load

13. To clarify that the OS has been mounted and ready to boot to, you can confirm its presence in XCC under **Boot options**, as shown in Figure 6-10. The server is now ready to be rebooted and will load the ISO.




Figure 6-10 ISO image mounted

14. You can reboot the server from the media mounting screen in XCC, as shown in Figure 6-9 or you can reboot the server from the remote console screen power menu as shown in Figure 6-11 on page 190.

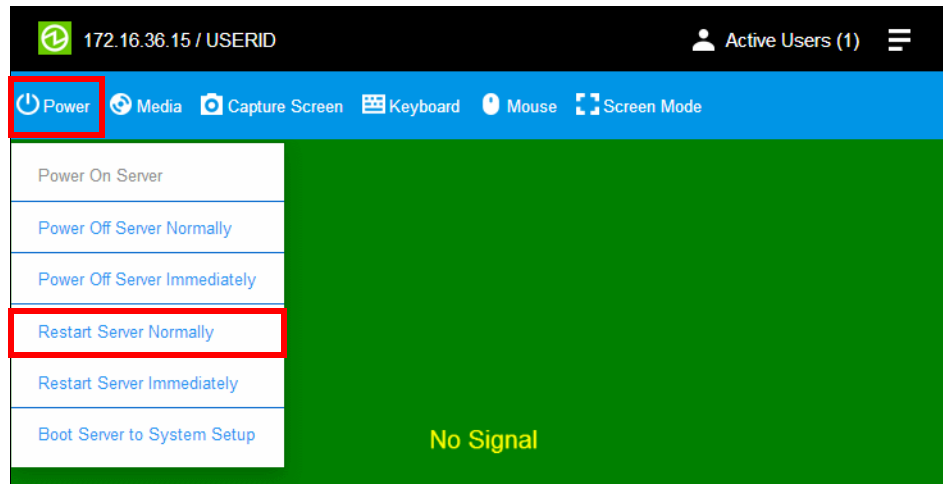


Figure 6-11 Restart the server to boot to the ISO

15. Click **OK** to the confirmation message to restart the server. When the ISO loads, follow the operating system prompts to install the OS.

For specific OS installation examples, refer to 6.3, “Operating system installation examples” on page 200.

6.2.2 Installing an OS with Lenovo XClarity Provisioning Manager (LXPM)

Lenovo XClarity Provisioning Manager (LXPM) allows installation of an operating system directly from its graphical user interface (GUI). To access LXPM, refer to 5.1.1, “Accessing Lenovo XClarity Provisioning Manager” on page 120.

There are two ways in which to install an OS via LXPM:

- **Manual install:** Install the operating system manually. This means you can save the drivers to a USB drive for installation at a later time. In some instances (for example, Linux installations), the storage controller drivers must be mounted separately from the OS at the same time as the OS is mounted and installed. See the SUSE and RHEL installation examples in 6.3, “Operating system installation examples” on page 200, for further details of this.
- **Automatic install:** Install the operating system in unattended mode. After the installation is complete, all drivers are automatically installed as well and no further action is required.

Tip: Before installing an OS via LXPM, load the ISO remotely with XCC. Refer to “Mounting an ISO image” on page 186 for instructions.

LXPM manual OS installation

To install an OS manually with LXPM, follow these steps:

1. Click on the **OS Installation** option from the left hand menu of LXPM.
2. Select **Manual Install** as shown in Figure 6-12 on page 191.




Figure 6-12 Manual OS installation selection with LXPM

3. Click **Next**.
4. Insert a USB drive into the front USB port of the SR950.
5. Select **OS Family** that will be installed on the server
6. Select the **OS Edition** that is mounted
7. Select the **USB location to save drivers** to, as shown in Figure 6-13. This will create a local driver repository on the USB device which can be used during an OS installation or installed later.

Tip: Alternatively you can load and apply the drivers via the mount media option after the OS has been mounted.




Figure 6-13 Select the location of the USB drive to save drivers

8. Refer to the next section, “LXPM automatic OS installation” steps 6 onwards for completing the OS install.

LXPM automatic OS installation

To complete the OS installation using the Manual install option in LXPM, follow these steps:

1. From the LXPM interface, click the **OS Installation** option from the left hand menu, as shown in Figure 6-14.

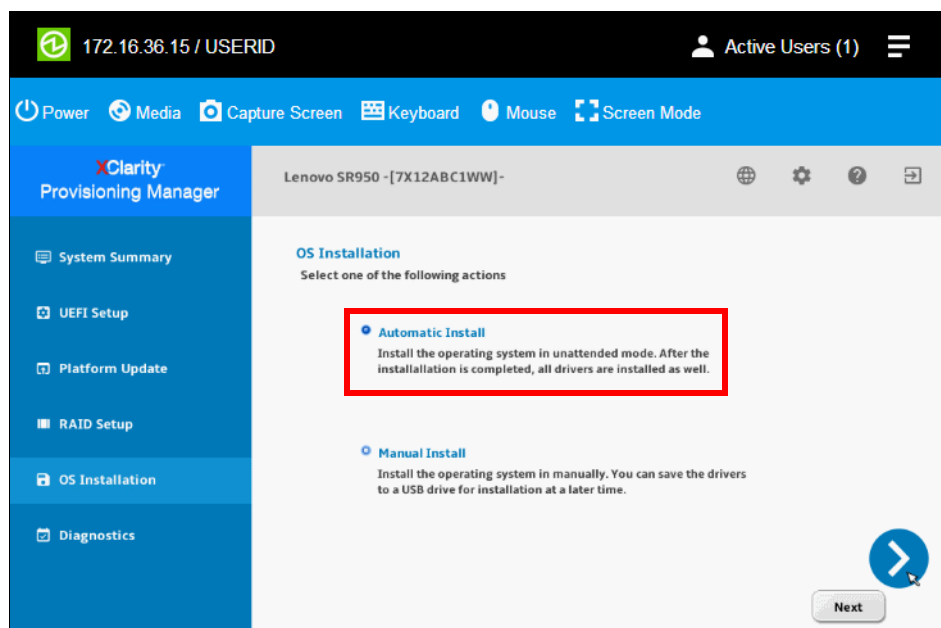


Figure 6-14 Automatic OS installation selection with LXPM

2. Click **Next**.
3. Select the **OS Family** and **OS Edition**. Ensure these selections match the OS you are installing.
4. Select where the **OS Image** will be installed from. In this example, the ISO is local as it has been mounted remotely, as explained in “Mounting an ISO image” on page 186.
5. Select **Automatic Install** and click **Next**.
6. Select the drive where the OS will be installed and booted from. In this example, the OS is being installed on an NVMe adapter as shown in Figure 6-15 on page 193.




Figure 6-15 Drive selection screen

7. Select the **OS Family**, **OS Edition**, and location of the **OS Image** that will be installed as shown in Figure 6-16. Note, the installation will fail if the OS and OS edition selected here does not match the OS and edition being installed on the SR950. Click **Next**.

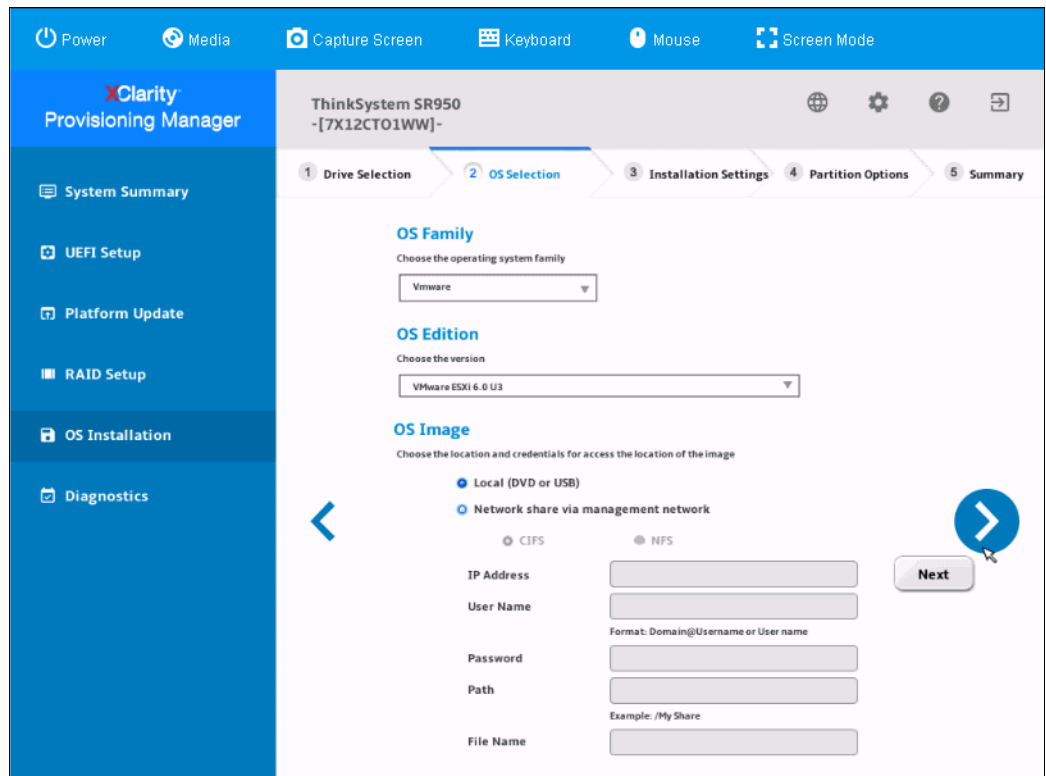


Figure 6-16 OS Selection screen

- Enter the **Host Name** and **Root Password** for the OS being installed as shown in Figure 6-17. Enter the Networking for an IP address and DNS settings if required. Click **Next**.

Figure 6-17 Installation settings

- Ensure the data enter is correct on the summary page as shown in Figure 6-18 and click **Next** to begin the installation of the OS.

Figure 6-18 Summary

10. The Operating System will begin to load as shown in Figure 6-19. Follow the prompts of the OS installation to complete the install.

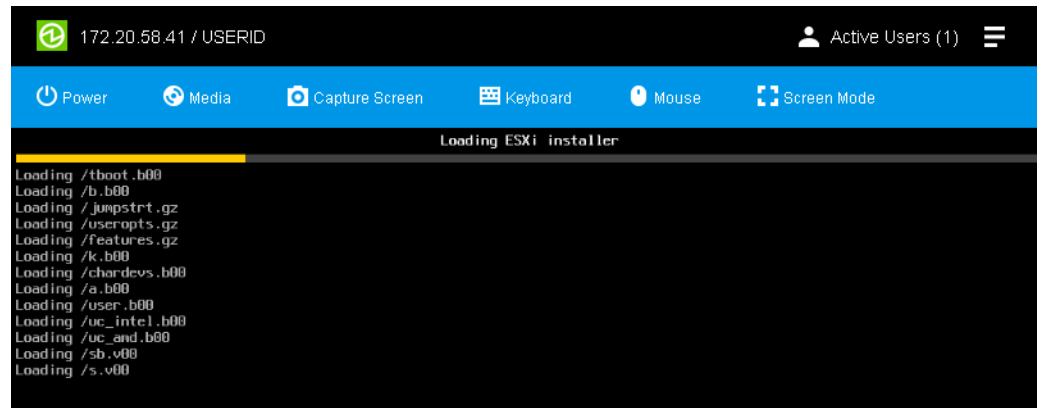


Figure 6-19 OS install

For specific OS installation examples, refer to 6.3, “Operating system installation examples” on page 200.

6.2.3 Installing an OS with XClarity Administrator (LXCA)

You can use Lenovo XClarity Administrator (LXCA) to deploy an operating-system image to up to 28 servers concurrently.

The Pro version, XClarity Pro, is required for OS deployments. For details, see the Lenovo XClarity Administrator product guide:

<https://lenovopress.com/tips1200#lenovo-xclarity-pro>

Note: The target drive must be initialized before the OS can be deployed with LXCA.

Steps to install using LXCA are as follows:

1. Before beginning the installation, validate that the management server network port being used to attach the data network is configured to be on the same network as the data network ports on the server(s).
2. From the Lenovo XClarity Administrator menu bar, click **Provisioning** → **Deploy OS images** as shown in Figure 6-20 on page 196 to display the Deploy Operating System: Deploy OS images page.

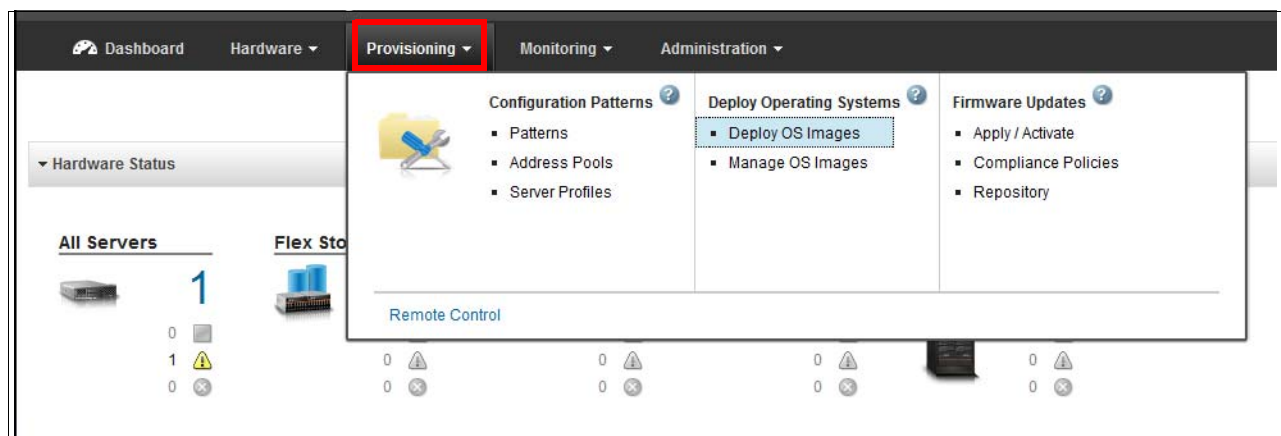


Figure 6-20 XClarity Administrator deploy images

3. Select the server or multiple servers to which the operating system is to be deployed as shown in Figure 6-21. You can deploy an operating system on up to 28 servers at one time.




Figure 6-21 Select servers and images

4. If you must change the OS image, select **Change Selected** → **Image to deploy**, as shown in Figure 6-22.




Figure 6-22 Change OS if required

5. Choose the required OS image from list of available images, as shown in Figure 6-23 on page 197.




Figure 6-23 Select OS image

6. If required, click **Change Selected** → **Network Settings** to configure network settings for the managed servers.
7. If a license is required, click the **License Key** icon and specify the license key to use to activate the operating system after it is installed as shown in Figure 6-24.




Figure 6-24 License key

8. The License Key activation screen appears, as shown in Figure 6-25. Select the license type as follows:
 - **Use the volume license key defined in Global Setting.** Use this to enter global settings to be used by all servers
 - **Use the following retail license key.** Use this for activating individual servers with a retail vendor key

If you are using a retail license, enter the license key as shown. Click **OK**.




Figure 6-25 License key activation

Note: Lenovo XClarity Administrator supports default volume-license keys for Windows installations and individual retail keys for both Windows and VMware ESXi.

9. For each server, select the preferred storage location where you want to deploy the operating system image from the **Storage** column as shown in Figure 6-26 on page 198.




Figure 6-26 Select where the OS will be installed to

Tip: To ensure that operating system deployments are successful, detach all storage from the managed server except the storage that is chosen for the operating-system deployment.

10. Verify that the deployment status for all selected servers is **Ready**.

If the server is not ready for an OS deployment, then click the **Not Ready** link to get information to help resolve the problem. If the network settings are not valid, click **Change Selected** → **Network Settings** to configure the network settings.

11. You will be prompted to enter an administrator password before deploying, as shown in Figure 6-27.

Set Credentials

Operating system credentials must be set.

You must set the operating system credentials before you deploy an operating system.

Credentials

Set the credentials to be used on the deployed operating systems.

RHEL/ESXi

User: root


Password:

Confirm Password:

Figure 6-27 Set the OS administrator password

12. Once the administrator password has been set, confirm the credentials and click the **Deploy images** icon to initiate the Operating System deployment as shown in Figure 6-28 on page 199.

Deploy OS Images

 **Operating systems on the selected compute nodes will be overwritten.**

[Show Details](#) ✕

Summary

Active Domain Directory

Confirm the following selections and click 'Deploy' to schedule the images for deployment.

[Expand All](#) [Collapse All](#)

▼ SR950 esxi6.5_Ubjesxi6.5_Ubxs86_64-Install-Virtualization

Bay:

Not set

Storage:

Local Disk

Host Name:

node20202020202

MAC:

00:0A:F7:26:D5:A0

IP Address:

172.16.32.22

Subnet Mask:

255.255.255.0

Gateway:

172.16.32.254

DNS 1:

Not set

DNS 2:

Not set

MTU:



1500

Deploy

Cancel

Figure 6-28 Deploy OS

13. The OS deployment starts and an information message opens, as shown in Figure 6-29.

 OS deployment to compute node ("SR950:20202020202020202020202020202020") has been started. [Hide Details](#) 

FQXHMFC0120I

It might take a few minutes to deploy the images. You can monitor progress from the Jobs list.

Figure 6-29 OS deploying

You can monitor the OS deployment progress from the Jobs page, as shown in Figure 6-30.

[illegible]

Figure 6-30 Monitor the OS deployment in the Jobs section

For additional information on Lenovo XClarity Administrator, refer to 7.2, “Lenovo XClarity Administrator” on page 225.

Download XClarity Administrator from the following link:

<http://www3.lenovo.com/us/en/data-center/software/systems-management/xclarity/>

6.2.4 Lenovo XClarity Bootable Media Creator (BoMC)

The Lenovo XClarity Essentials Bootable Media Creator (BoMC) provides a tool for creating a bootable image for supported media (CD, DVD, ISO image, USB flash drive, or PXE files) to create a ServerGuide Setup and Installation CD for deploying Windows operating systems and updates on the SR950.

For information on creating a bootable media CD, refer to 5.3.2, “XClarity Essentials Bootable Media Creator (BoMC)” on page 165.

Download BoMC from the following link:

<https://support.lenovo.com/no/en/solutions/lncv-bomc>

6.2.5 Installing an OS with the local USB port

You can use the local USB port on the front of the SR950 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 the following web pages:

- ▶ Installing Windows from a USB flash drive:
<http://technet.microsoft.com/en-us/library/dn293258.aspx>
- ▶ How to create a bootable USB drive to install SLES:
<http://www.novell.com/support/kb/doc.php?id=3499891>
- ▶ 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
- ▶ 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

6.3 Operating system installation examples

This section provides example installations of the following OSes using XClarity Provisioning Manager (LXPM) UEFI installations:

- ▶ 6.3.1, “Installing Windows Server 2012 R2 and 2016” on page 201
- ▶ 6.3.2, “Installing SUSE Linux Enterprise Server 12 SP2” on page 205
- ▶ 6.3.3, “Installing Red Hat Enterprise Linux 7.3” on page 211
- ▶ 6.3.4, “Installing Lenovo Customized VMware ESXi 6.5” on page 216

All of the following examples make use of the remote mounting function to deploy the operating systems remotely to the SR950 server.

Tip: Ensure the disks and controllers have been initialized and configured as explained in Chapter 5, “Preparing the hardware” on page 119.

6.3.1 Installing Windows Server 2012 R2 and 2016

To install Windows Server on the SR950 server using LXPM, use the following procedure:

1. Connect to XCC, launch a remote console session and mount the Windows ISO image to the SR950 server as explained in “Mounting an ISO image” on page 186.
2. Enter into LXPM. For information on how to access LXPM, refer to 5.1.1, “Accessing Lenovo XClarity Provisioning Manager” on page 120.
3. Select **OS Installation** from the side menu on the LXPM interface.
4. Select **Automatic Install**.

Note: If a Manual installation is selected, then the Windows drivers will not be installed. You can optionally mount the drivers after the OS has been mounted by clicking the + sign, as shown in Figure 6-31 for the drivers to be installed during the OS install.

If Automatic Install is selected then the drivers are slipstreamed into the OS installation automatically and no further action is required.

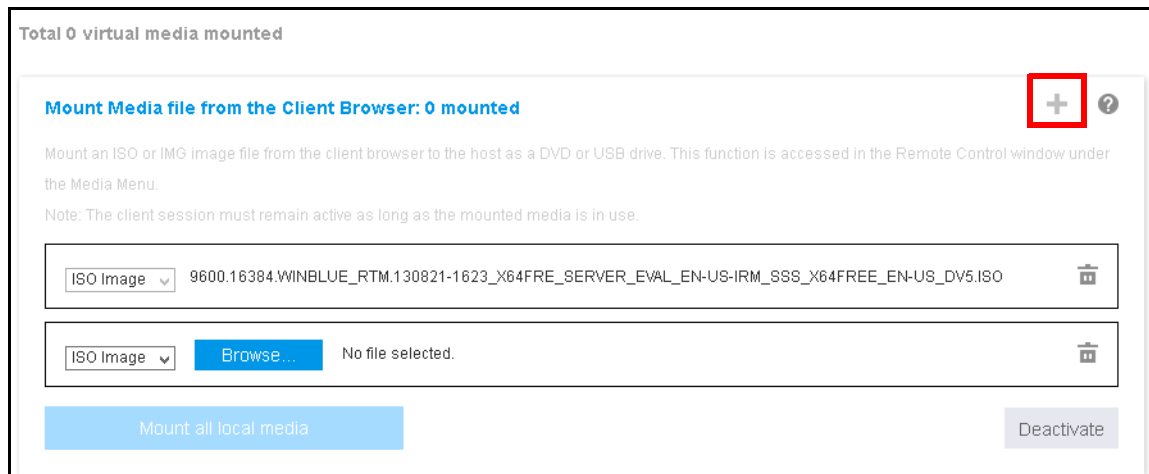


Figure 6-31 ISO and drivers mounted in LXPM

5. Enter in the details for Drive Selection and click **Next**.
6. Select the **OS Family**, **OS Edition** and **OS Image** location from the following screen and click **Next**.
7. Select the Installation type, Boot mode, Time Zone, Language (optional), License Key and NetBIOS host name, as shown in Figure 6-32 on page 202 and click **Next**.

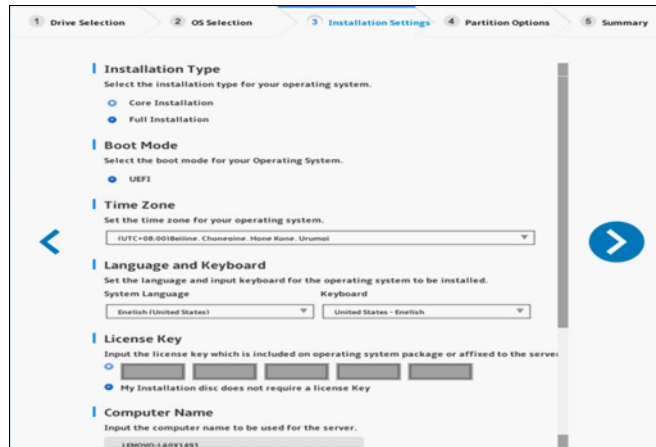


Figure 6-32 Windows installation settings

8. Make adjustments to the Partitions if required as shown in Figure 6-33.

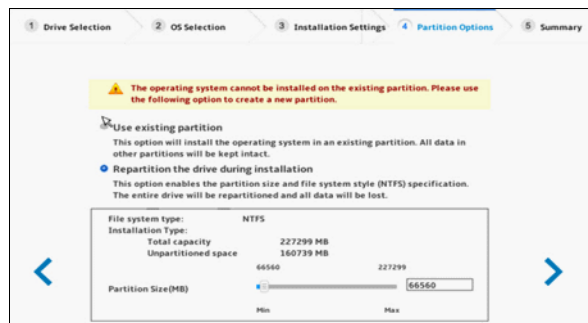


Figure 6-33 Windows partition settings

9. Clarify the settings are correct in the **Summary** tab and click **Next** to begin the installation as shown in Figure 6-34. The Windows installation will now begin.

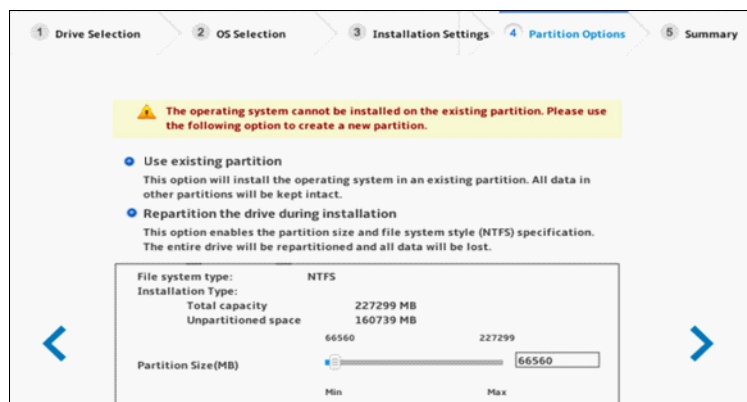


Figure 6-34 Summary

10. Read and approve the License Agreement (NOTE: System will stay on this screen until the License agreement is accepted) as shown in Figure 6-35 on page 203.

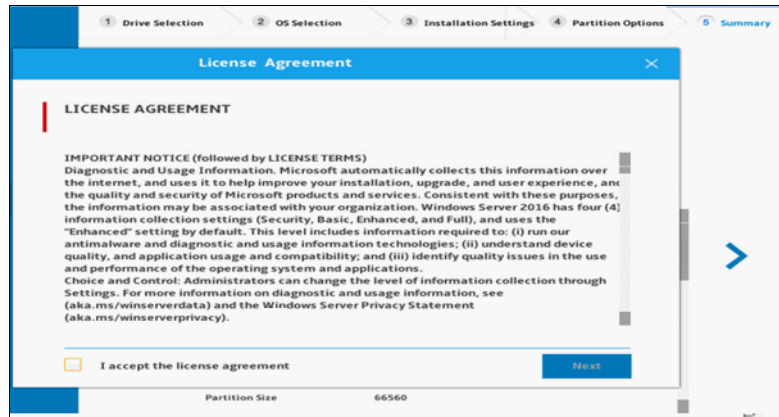


Figure 6-35 Windows license agreement

11. The OS installation will load as shown in Figure 6-36. This may take some time to load.




Figure 6-36 Windows installation loading screen

12. If the drivers are mounted separately from the OS then during the installation the base OS will not find any drives. Select **Load Drivers** when prompted and navigate to the drive containing the driver. The drivers will then load as shown in Figure 6-37 on page 204.

Note: Ensure you have mounted the OS ISO before you mount the driver packages in LXPM as shown in Figure 6-31 on page 201.




Figure 6-37 Windows drivers loading screen

13. The installation will continue after the drivers are loaded as shown in Figure 6-38.




Figure 6-38 OS Windows installation

14. After the Windows operating system installation completes, you will see that the appropriate OS drivers have been installed for the hardware as shown in Figure 6-39 on page 205.




Figure 6-39 OS drivers in the windows install

6.3.2 Installing SUSE Linux Enterprise Server 12 SP2

To install SUSE Linux Enterprise Server (SLES) operating system on the SR950 server using LXPM, follow these steps:

1. Download the SUSE SLES installation ISO from the following link:

<https://www.suse.com/products/server/download/>

2. Download the ISO drivers for the storage controller. The drivers can be downloaded from the following link:

<https://datacentersupport.lenovo.com/no/en/products/servers/thinksystem/sr950/7x12/downloads>

- a. Select the **OS SUSE Linux Enterprise Server 12** from the top drop-down menu.
- b. Expand the **Controller, RAID, SATA, SAS, SCSI and iSCSI** menu.
- c. Download the appropriate driver ISO for your controller. Example:

- For 530-x and 930-x controllers, use this driver ISO:
`megaraid_sas-07.700.26.00-1-sles12sp2.x86_64.iso`
- For 430-x class SAS HBAs, use this driver ISO:
`mpt3sas-20.00.00.00-4-sles12sp2.x86_64.iso`

In this installation example the 430-x class SAS HBA is installed and requires the following file as shown in Figure 6-40 on page 206

`lnvgy_dd_mpt3sas_20.00.00.00-4_sles12_x86-64.tgz`

- d. The driver ISO file is contained in the TAR file. Extract the TAR files (.tgz) to a local file on your computer. Note that the driver file may also be an image (IMG) file. In this example, it is an ISO file. In the instance that it is an IMG file, then mount it the same as you would mount the ISO file.

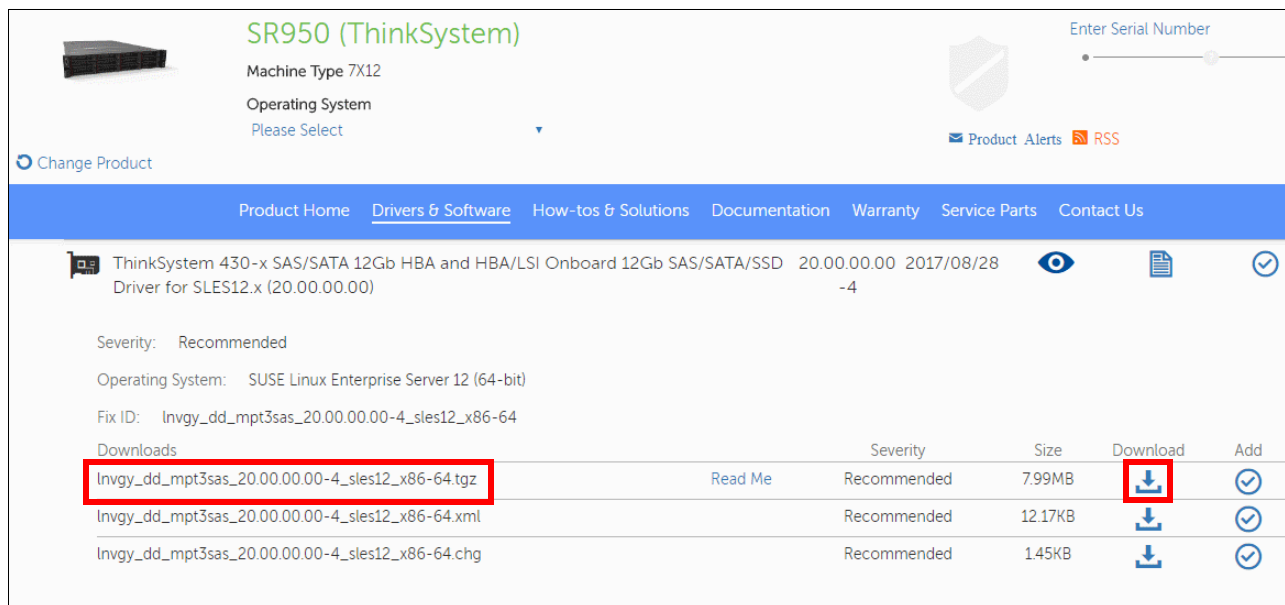


Figure 6-40 Download the SR950 controller driver(s)

3. Connect to XCC, launch a remote console session and mount the SUSE SLES installation ISO image to the SR950 server as explained in “Mounting an ISO image” on page 186.
4. Before mounting the OS, click the + sign to mount the driver ISO as well. Ensure the drivers are mounted after the OS has been mounted.
5. Click **Mount all local media** to mount the OS and drivers as shown in Figure 6-41.

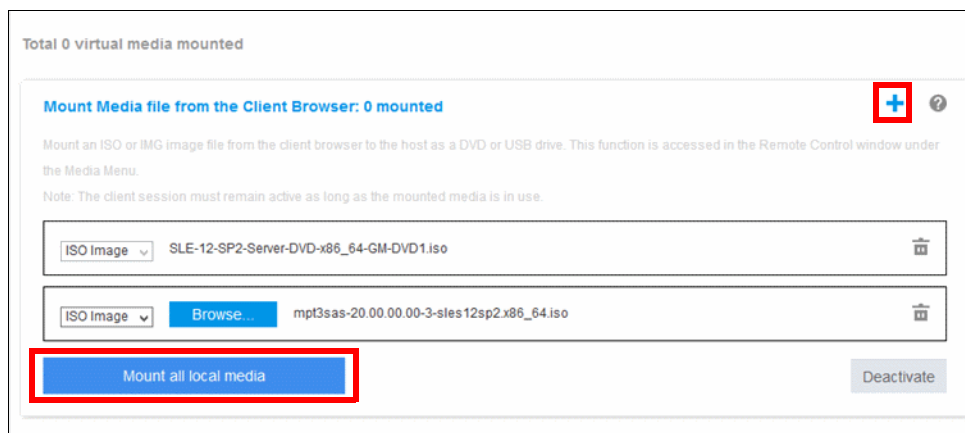


Figure 6-41 Mount OS ISO and driver ISO

6. Enter into LXPM. For information on how to access LXPM, refer to 5.1.1, “Accessing Lenovo XClarity Provisioning Manager” on page 120.
7. Select **OS Installation** from the side menu on the LXPM interface.
8. Select **Automatic Install** and click through until the installation begins as shown in Figure 6-42 on page 207.

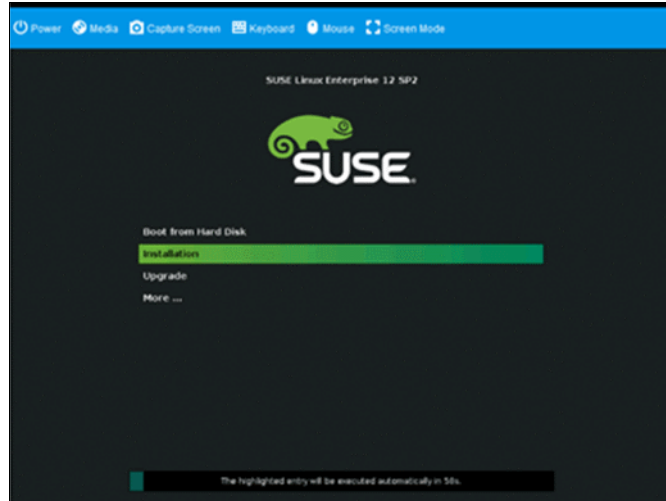


Figure 6-42 SUSE installation screen

9. Accept the EULA and click **Next** to begin the system hardware scan by the installation wizard.
10. Either register or skip the registration on the following screen. For this installation the registration was skipped as shown in Figure 6-43. Select **OK** then **Next**.




Figure 6-43 Registration option

11. The **Add On Product** screen will appear. This is where the updated storage drivers previously downloaded will be applied. Click **I would like to install an additional Add On product** and select **DVD**. Then click **Next**, as shown in Figure 6-44 on page 208.

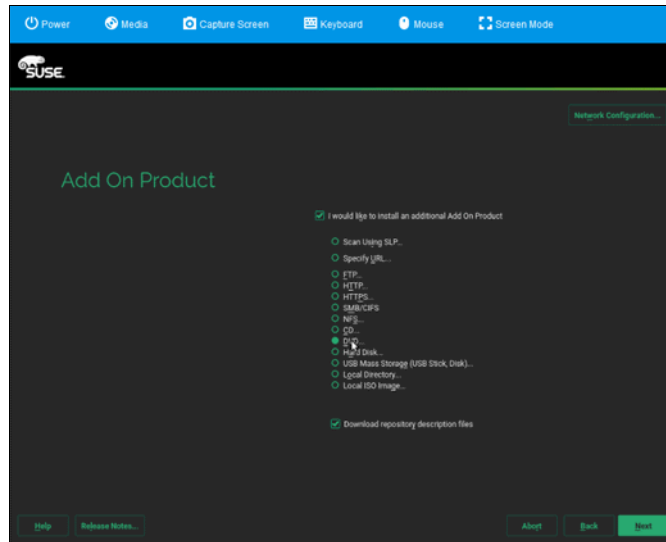


Figure 6-44 Add on product option: select DVD

12. The system will search for the mounted ISOs. There will be two devices listed (the ISO for the OS and the ISO for the drivers) as shown in Figure 6-45. Select the bottom device, as we mounted the driver ISO file second, then click **Next**.

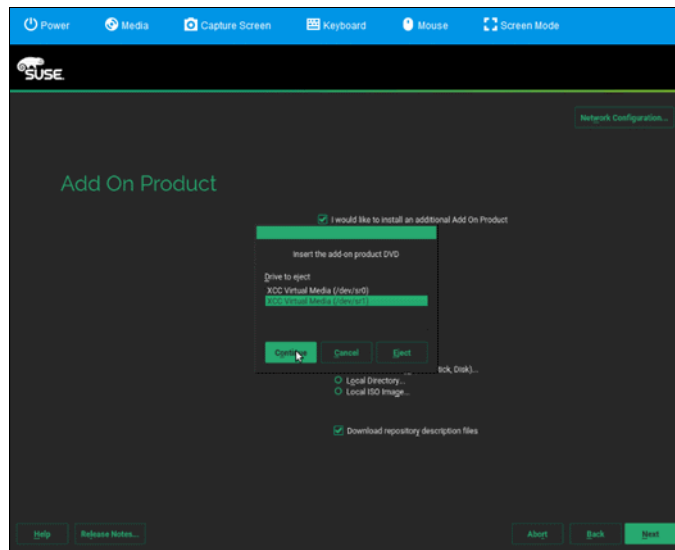


Figure 6-45 Add on product: storage driver ISO

13. A new repository will be added as shown in Figure 6-46 on page 209. This repository is the ISO containing the updated storage drivers.




Figure 6-46 New repository being created

14. Once the repository has loaded the drivers, you can add and delete new repositories from the following screen as shown in Figure 6-47. The controller drivers are now installed. Click **Next** if there are no additional drivers to add.

Note: If you are unable to find a particular driver for a specific device, check with the SUSE SolidDriver Program accessed from the following link:

<https://drivers.suse.com/>

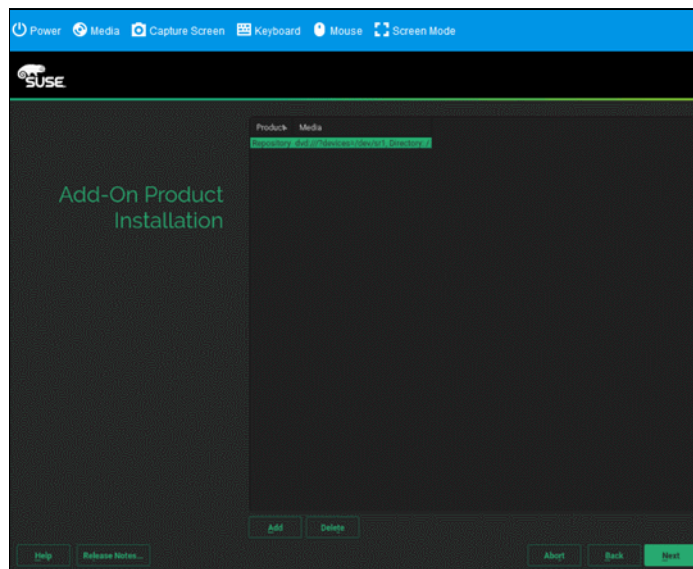


Figure 6-47 Add on product for additional drivers

15. Select the System Role: in this example the system role is **Default System** as shown in Figure 6-48 on page 210. Click **Next**.

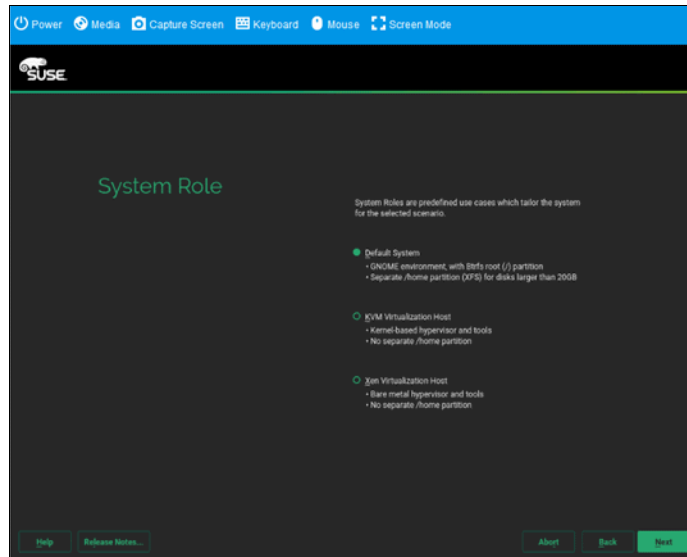


Figure 6-48 System role selection

16. From the following screen, select **Create Partition Setup** as shown in Figure 6-49.

Note: If no drives are present either the wrong controller driver ISO was selected or the driver files were not correctly loaded. Start again from step 2 if this is the case.




Figure 6-49 Partitioning option

17. Select the desired drive to install the OS on and click **Next** as shown in Figure 6-50 on page 211.

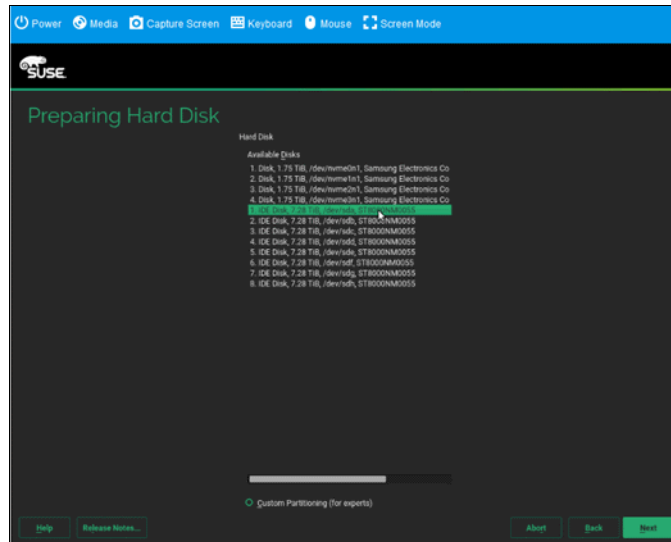


Figure 6-50 Select where the OS will be installed

18. Follow the installation wizard prompts and make the selections best suited for your environment. When the installation is complete, the server will be rebooted.

6.3.3 Installing Red Hat Enterprise Linux 7.3

To install Red Hat Enterprise Linux (RHEL) operating system on the SR950 server using LXPM, follow these steps:

1. Ensure you have the RHEL 7.3 installation ISO. If this has previously been brought and you need to re-download the ISO then download it from the following link:
<https://access.redhat.com/downloads>
2. Download the ISO drivers for the storage controller. The drivers can be downloaded from the following link:

<https://datacentersupport.lenovo.com/no/en/products/servers/thinksystem/sr950/7x12/downloads>

- a. Select the **Red Hat Enterprise Linux 7 (64-bit)** from the top drop-down menu.
- b. Expand the **Controller, RAID, SATA, SAS, SCSI and iSCSI** menu.
- c. Download the appropriate driver ISO for your controller. Example:

- For 530-x and 930-x controllers, use the driver ISO found in:

`lnvgy_dd_sraidmr_7.700.26.00-1_rhel7_x86-64.tgz`

- For 430-x class SAS HBAs, use this driver ISO:

`mpt3sas_20.00.00.00-4_rhel7_x86-64.iso`

In this installation example the 430-x class SAS HBA is installed and requires the following file as shown in Figure 6-51 on page 212:

`lnvgy_dd_mpt3sas_20.00.00.00-4_rhel7_x86-64.tgz` .

The driver ISO file is contained in the TAR file. Extract the TAR files (.tgz) to a local file on your computer. Note that the driver file may also be an image (IMG) file. In this example, it is an ISO file. In the instance that it is an IMG file, then mount it the same as you would mount the ISO file.

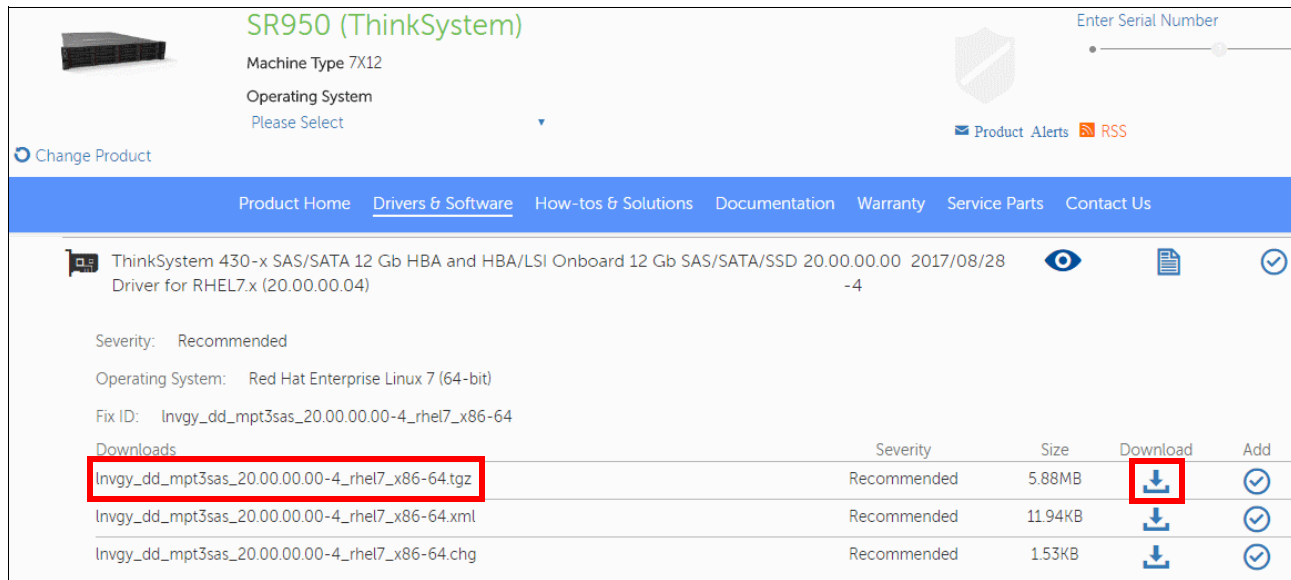


Figure 6-51 Download the RHEL drivers for the SR950 controller

3. Connect to XCC, launch a remote console session and mount the RHEL installation ISO image to the SR950 server as explained in “Mounting an ISO image” on page 186.
4. Before mounting the OS, click the + sign to mount the driver ISO as well. Ensure the drivers are mounted after the OS has been mounted.
5. Click **Mount all local media** to mount the OS and drivers as shown in Figure 6-52.

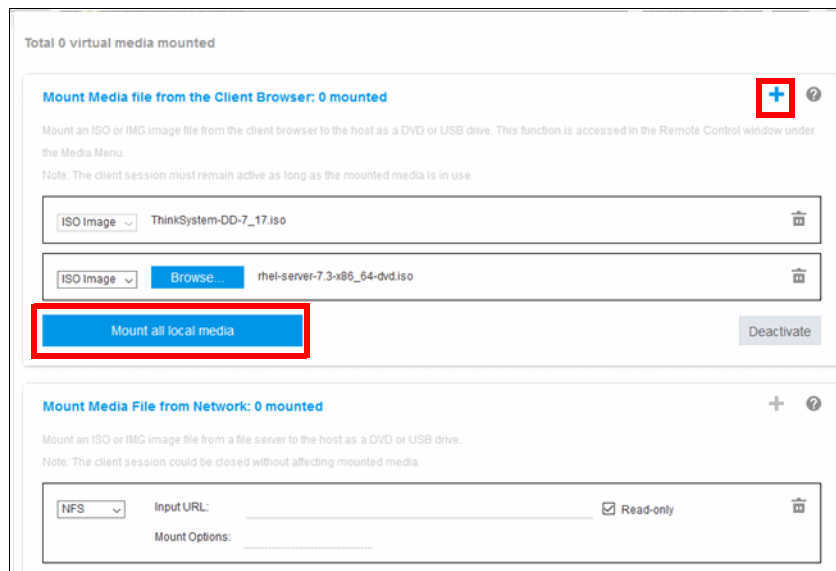


Figure 6-52 Mount OS ISO and driver ISO

6. Close the mount media screen and hit **F12**.
7. In the **Boot Devices Manager** screen, select **UEFI: XCC Virtual Media 0399**, as shown in Figure 6-53 on page 213.




Figure 6-53 Boot device manager

8. When RHEL starts, select the first option **Install Red Hat Enterprise Linux 7.3** from the initial install screen.
9. Press the Tab key to get to the full configuration boot options menu, as shown in Figure 6-54. This is necessary to add the grub installation line that will point to the 2nd ISO driver file that was mounted previously.

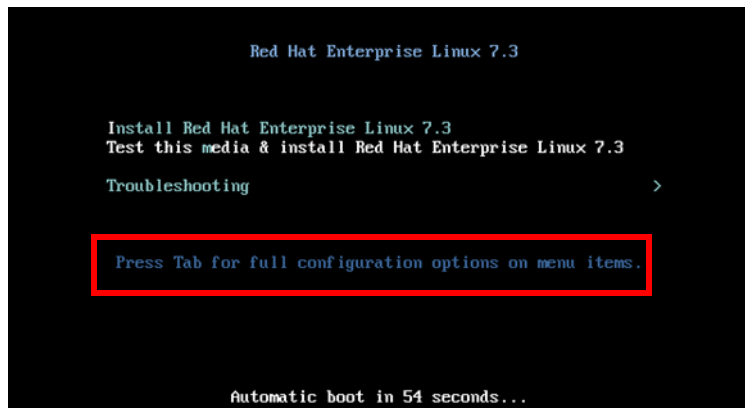


Figure 6-54 Press tab to access the configuration menu

10. Use the keyboard down arrow key to access the second line and add the device location of the controller drivers. In this example the driver was mounted remotely so it is denoted as **/dev/sr1**, as shown in Figure 6-55 on page 214.

Note: In this example we have pointed to `inst.dd=/dev/sr1`. This is because it was the second ISO we mounted. The first ISO mounted was the RHEL installation ISO and would point to `/dev/sr0`.

```
setparams 'Install Red Hat Enterprise Linux 7.3'

linuxefi /images/pxeboot/vmlinuz inst.stage2=hd:LABEL=RHEL-7.3\x20Server.x86_64 quiet inst\
.dd=/dev/sr1_
initrd /images/pxeboot/initrd.img

Press Ctrl-x to start, Ctrl-c for a command prompt or Escape to discard edits and
return to the menu. Pressing Tab lists possible completions.
```

Figure 6-55 Add the location of the driver ISO

11. Press **Ctrl-X** to continue the RHEL OS installation.

During the installation, you will notice that /dev/sr0 and /dev/sr1 are mounted and examined for installation source and drivers, as shown in Figure 6-56.

```
Power Media Capture Screen Keyboard Mouse Screen Mode

21.3651911 i8042: No controller found
OK | Started Plymouth Boot Screen.
OK | Reached target Paths.
OK | Reached target Basic System.
OK | Started Device-Mapper Multipath Device Controller.
Starting Open-iSCSI...
OK | Started Open-iSCSI.
Starting dracut initqueue hook
29.7397471 dracut-initqueue[951]: mount: /dev/sr0 is write-protected, mounting read-only
33.5885601 dracut-initqueue[951]: DD: Examining /dev/sr1
34.0205031 dracut-initqueue[951]: mount: /dev/sr1 is write-protected, mounting read-only
35.2643601 dracut-initqueue[951]: DD: Extracting: knod-mpt3sas
36.6575691 dracut-initqueue[951]: DD: Extracting: knod-megaraid_sas
37.0504701 dracut-initqueue[951]: DD: WARNING: DD repo content not a file: /media/DD-1/hwr_broa/rpms/x86_64/repodata
37.0520501 dracut-initqueue[951]: DD: WARNING: DD repo content not a file: /media/DD-1/hba_avag/rpms/x86_64/repodata
37.9102401 dracut-initqueue[951]: DD: Examining /dev/sr1
38.4888231 dracut-initqueue[951]: mount: /dev/sr1 is write-protected, mounting read-only
39.5739151 dracut-initqueue[951]: DD: Extracting: knod-mpt3sas
40.6644931 dracut-initqueue[951]: DD: Extracting: knod-megaraid_sas
41.0705321 dracut-initqueue[951]: DD: WARNING: DD repo content not a file: /media/DD-2/hba_avag/rpms/x86_64/repodata
41.0717281 dracut-initqueue[951]: DD: WARNING: DD repo content not a file: /media/DD-2/hwr_broa/rpms/x86_64/repodata
OK | Started dracut initqueue hook.
OK | Reached target Remote File Systems (Pre).
OK | Reached target Remote File Systems.
Starting dracut pre-mount hook...
OK | Started dracut pre-mount hook.
OK | Reached target Initrd Root File System.
Starting Reload Configuration from the Real Root...
OK | Started Reload Configuration from the Real Root.
```

Figure 6-56 RHEL installation

12. A GUI will appear as shown in Figure 6-57 on page 215. Select the **Installation Destination** option to select where the RHEL installation will be installed.

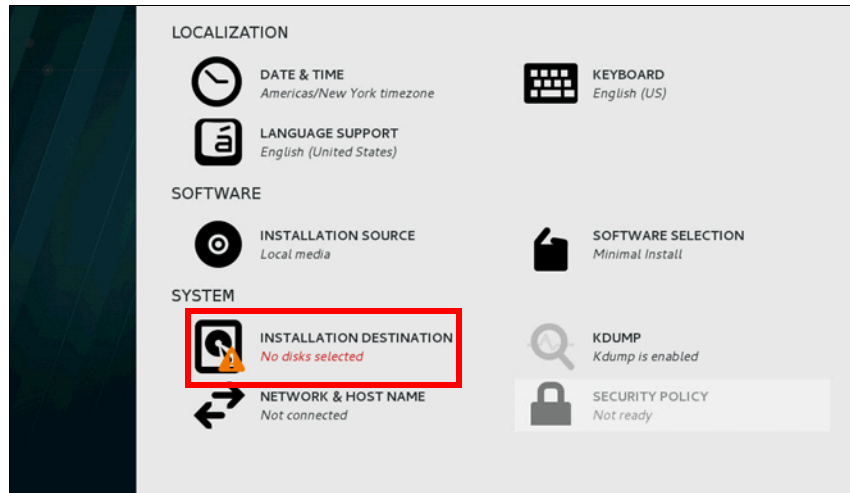


Figure 6-57 Installation destination option

13. If the driver ISO has loaded correctly then the installation process will see the local drives available to install to, as shown in Figure 6-58. Select the drive or device where the OS is to be installed and click **Done**.

Notes:

- For drives that are connected to a RAID adapter, the name of the RAID adapter will be printed on the drive.
- If an M.2 boot adapter is installed, it will be named `Lite0n`. The M.2 boot adapter does not require drivers loaded at the installation time to be used as the OS installation target.

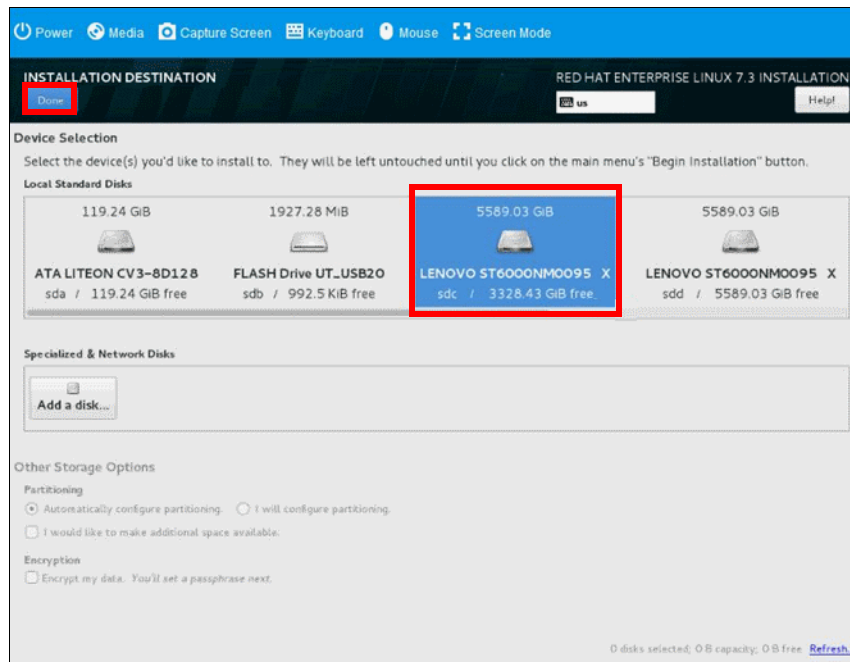


Figure 6-58 Installation destination

14. Follow the prompts for the remainder of the installation as required and reboot the server upon completion.

Note: It is also recommended that you update the device driver for the onboard Ethernet. These device drivers are located in the `intc-lnvgy_dd_nic_22.2-a_rhel7_x86-64.tgz` package, downloaded from the following link:

<https://datacentersupport.lenovo.com/no/en/products/servers/thinksystem/sr950/7x12/downloads>

6.3.4 Installing Lenovo Customized VMware ESXi 6.5

To install the VMware vSphere ESXi 6.5 operating system on the SR950 server using LXPM, follow the steps below:

1. Download the Lenovo Customized version of VMware vSphere from the following link:
https://my.vmware.com/web/vmware/info/slug/datacenter_cloud_infrastructure/vmware_vsphere/6_5#custom_iso

Note: The newest edition of the Lenovo Customized VMware vSphere ESXi has the necessary drivers slipstreamed into the ISO and they do not need to be applied separately.

2. Connect to XCC, launch a remote console session and mount the ESXi ISO image to the SR950 server as explained in “Mounting an ISO image” on page 186.

Figure 6-59 shows the Lenovo Customized VMware vSphere ESXi 6.5 version mounted to the server.

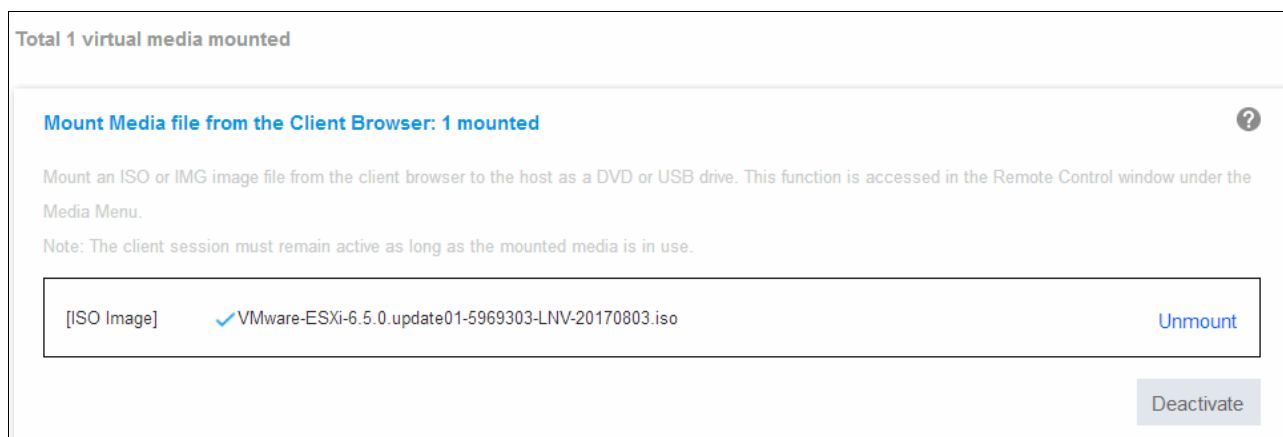


Figure 6-59 ESXi 6.5 mounted

3. Enter into LXPM. For information on how to access LXPM, refer to 5.1.1, “Accessing Lenovo XClarity Provisioning Manager” on page 120.
4. Select **OS Installation** from the side menu on the LXPM interface.
5. Select **Automatic Install** as shown in Figure 6-60 on page 217 and click **Next**.

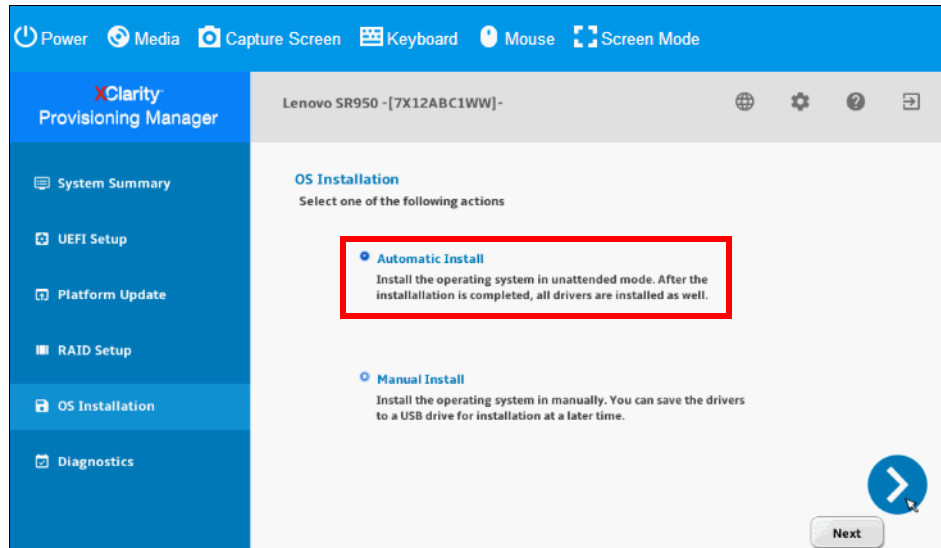


Figure 6-60 Automatic installation selection

6. Select the Drive where the OS will be installed on as shown in Figure 6-61 and click **Next**.

Important: If installing on a disk drive, ensure that the controller and disk drive have been initialized. If they have not been initialized then you can not write to the disk. To initialize, create a RAID array and partitions, refer to Chapter 5, “Preparing the hardware” on page 119.

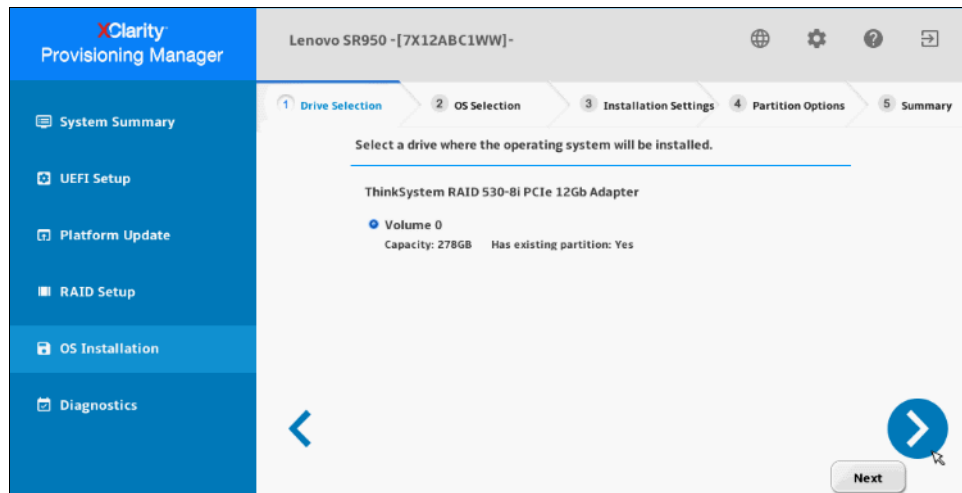


Figure 6-61 Select the drive where the OS will be installed

7. Select the **OS Family**, **OS Edition** and **OS Image** location from the following screen as shown in Figure 6-62 on page 218. In this example, the ISO has been mounted remotely and the SR950 server will see the media as a **Local DVD or USB**. Click **Next**.

Note: The edition of the OS mounted needs to match the edition selected here. If there is a mismatch between the two, the OS installation will fail.

Lenovo SR950 -[7X12ABC1WW]-

1 Drive Selection 2 OS Selection 3 Installation Settings 4 Partition Options 5 Summary

OS Family
Choose the operating system family
Vmware

OS Edition
Choose the version
Vmware ESXi 6.5 b

OS Image
Choose the location and credentials for access the location of the image

Local (DVD or USB)
Network share via management network

CIFS NFS

IP Address
User Name
Password
Path
File Name

Format: Domain@Username or User name
Example: /My Share

Next

Figure 6-62 OS Selection screen

- Set the name and password for the ESXi installation. Set the IP address to either obtain automatically or use a specific IP address. Lastly, set the DNS to obtain automatically or enter in the specific details of the DNS server and click **Next** as shown in Figure 6-63.

ThinkSystem SR950
-[7X12CT01WW]-

1 Drive Selection 2 OS Selection 3 Installation Settings 4 Partition Options 5 Summary

Input the computer name to be used for the server.
LENOVO-LA0X1052

Root Password
Set root password for this server.
Root Password
Confirm Password

Advanced

Ethernet Controller
Emulex VFAS ML2 Dual Port 10GbE SFP #0

IP Address Settings
Obtain an IP address automatically
Use the following IP address

IP Address
Subnet Mask
Default Gateway

DNS Settings
Obtain DNS server address automatically
Use the following DNS server address

Preferred DNS Server
Alternate DNS Server

Next

Figure 6-63 Windows installation settings

9. Make adjustments to the Partitioning if required and click **Next** to the **Summary** page as shown in Figure 6-64.

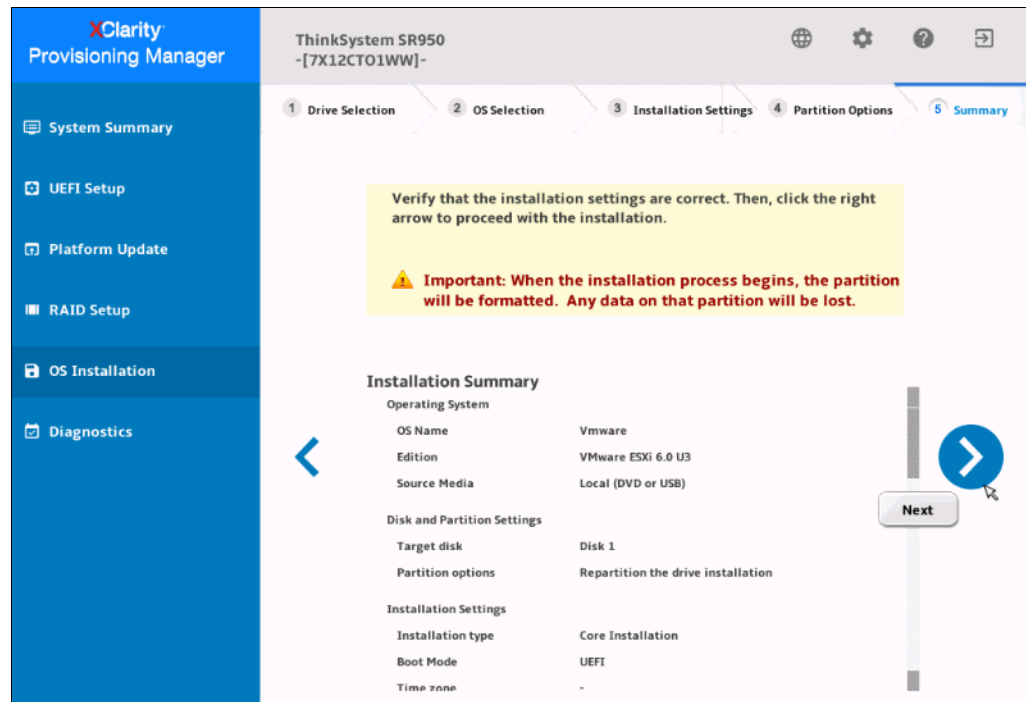


Figure 6-64 Summary of installation

10. Click **Next** to start the installation. The installation will begin as shown in Figure 6-65.

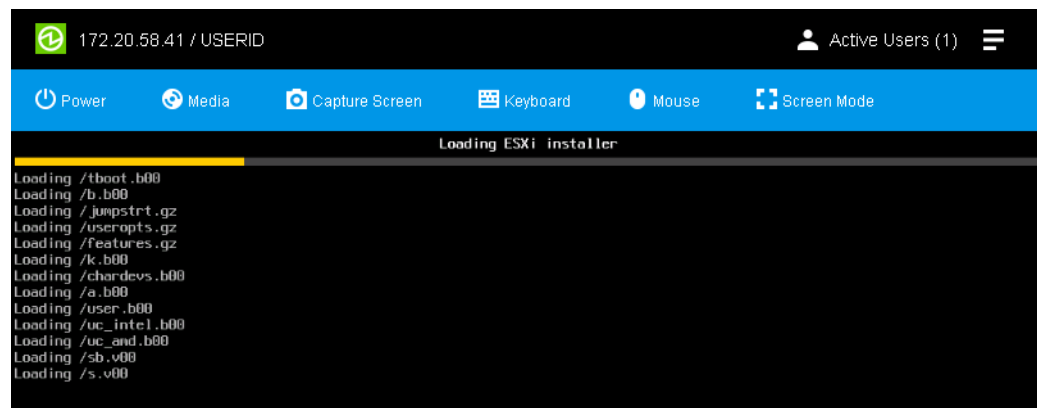


Figure 6-65 ESXi installer

11. Follow the prompts for the remainder of the installation as required and reboot the server upon completion.

Note: The Lenovo Customized VMware ESXi ISO images can be downloaded from the following links:

- ▶ Lenovo ESXi 6.5 versions at the Download VMware vSphere version 6.5 page - https://my.vmware.com/web/vmware/info/slug/datacenter_cloud_infrastructure/vmware_vsphere/6_5#custom_iso
- ▶ Lenovo ESXi 6.0 versions at the Download VMware vSphere version 6.0 page -

https://my.vmware.com/web/vmware/info/slug/datacenter_cloud_infrastructure/vmware_vsphere/6_0#custom_iso

6.4 Installing an OS on the M.2 adapter

The server supports one or two M.2 form-factor SATA drives for use as an operating system boot solution. One or two drives are supported which are installed on an adapter. The adapter is then installed in a dedicated slot on the system board.

When two M.2 drives are used, the drives are configured by default as a RAID-1 mirrored pair for redundancy. The drives also supports Secure Boot to ensure that only a digitally signed operating system can be used.

To install an OS on the M.2 adapter follow one of the procedures outlined in 6.2, “Installing an OS with Lenovo tools” on page 185.

When prompted for a drive selection to install the OS, ensure the **LITEON** option is selected, as shown in the example using Lenovo XClarity Provisioning Manager (LXPM) in Figure 6-66. The LITEON is the M.2 boot drive. Continue the OS installation as normal.

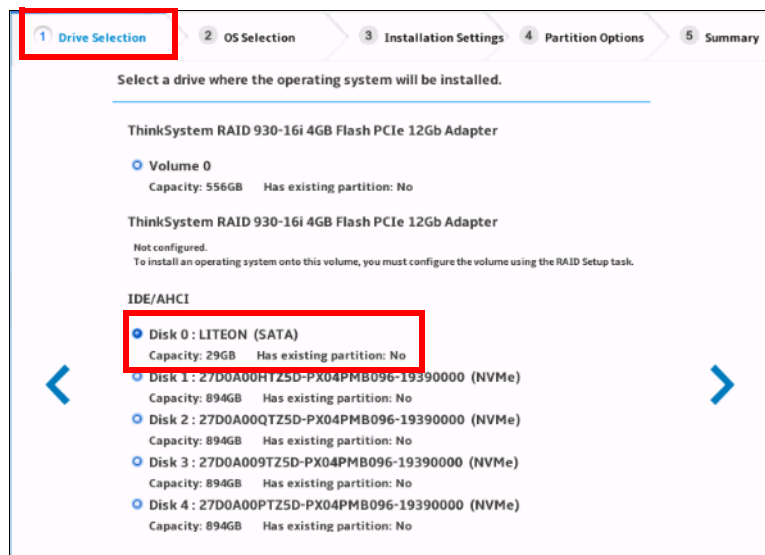


Figure 6-66 Drive selection for M.2 adapter

Notes:

- ▶ Although the adapter bus itself is faster than a SATA drive, the 7N47A00129 and 7N47A00130 M.2 SATA drives will perform at their max rated speed of 6Gbps. If you require fast read/writes from a drive then it is recommended that the M.2 adapter is only used for booting from and not reading and writing data to (such as event logs).
- ▶ Some Lenovo M.2 drives are designated as boot drives only; they do not have the endurance rating for normal R/W application workloads. Consult the M.2 product guide for details:

<https://lenovopress.com/lp0769-thinksystem-m2-drives-adapters>

6.5 Booting from SAN

Boot from SAN (also known as SAN Boot) is a technique that is used when the node in the chassis has no local disk drives or boot adapters (for example, M.2 adapter). 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 LXPM (UEFI) recognizes the adapter. Enter into LXPM, select **UEFI Setup** → **System Settings** → **NIC Configuration**. The Adapters and UEFI Drivers panel opens, as shown in Figure 6-67.




Figure 6-67 Adapters that are listed in LXPM-UEFI

- ▶ 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**.
- ▶ Set the HBA as the first boot order from the XCC as shown in Figure 6-68 on page 222. To access the XCC refer to 5.2.1, “Accessing XCC” on page 142.




Figure 6-68 Ensure HBA is first listed in boot order

- ▶ Remove all devices from this boot order list that might not host an OS. The optimal minimum configuration is CD/DVD and Hard Disk 0.
- ▶ 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
- ▶ *Windows Boot from Fibre Channel SAN – Overview and Detailed Technical Instructions for the System Administrator:*
<https://support.microsoft.com/en-us/help/305547/support-for-booting-from-a-storage-area-network-san>
- ▶ *vSphere Storage* document from VMware:
<https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.storage.doc/GUID-81722422-A131-4E6A-905C-6B3CA0E32D99.html>

Managing the hardware

Maintaining and managing systems efficiently is essential to business IT operations. In this chapter, we discuss the Lenovo applications that are available to manage and maintain the SR950 server.

Topics in this chapter are:

- ▶ 7.1, “Management tools overview” on page 224
- ▶ 7.2, “Lenovo XClarity Administrator” on page 225
- ▶ 7.3, “Lenovo XClarity Integrators” on page 236
- ▶ 7.4, “Lenovo XClarity Energy Manager” on page 240
- ▶ 7.5, “Lenovo XClarity Controller management functions” on page 263
- ▶ 7.6, “Lenovo XClarity Essentials” on page 271
- ▶ 7.7, “LCD system information display panel” on page 272

7.1 Management tools overview

There are a number of Lenovo tools available for managing the SR950 server, including:

- ▶ Lenovo XClarity Administrator
- ▶ Lenovo XClarity Integrator
- ▶ Lenovo XClarity Energy Manager
- ▶ Lenovo XClarity Provisioning Manager
- ▶ Lenovo XClarity Essentials
- ▶ Lenovo XClarity Controller

Table 1 shows a comparison of the capabilities of each of the tools.

Table 1 Lenovo tools for managing the SR950 server

Function	XClarity Administrator	XClarity Integrator	XClarity Energy Manager	XClarity Provisioning Manager	XClarity Essentials ^a	XClarity Controller
Multiple systems management	Yes	Yes	Yes	No	Yes	No
Operating system deployment	Yes	No	No	Yes	No	No
Firmware updates ^b	Yes	Yes	No	Yes ^c	Yes	Yes
System configuration	Yes	Yes	No	Yes	Yes	Yes
Events / alerts	Yes	Yes	Yes			Yes
Inventory / Log	Yes	Yes		Yes ^d	Yes	Yes
Power management	No	Yes ⁵	Yes	No	No	No

a. Lenovo XClarity Essentials includes Lenovo XClarity Essentials OneCLI, Lenovo XClarity Essentials Bootable Media Creator, and Lenovo XClarity Essentials UpdateXpress.

b. Most options can be updated through the Lenovo tools. Some options, such as GPU firmware or Omni-Path firmware require the use of vendor tools.

c. Firmware updates are limited to core updates which include updates for Lenovo XClarity Provisioning Manager, Lenovo XClarity Controller firmware, and UEFI only. Firmware updates for optional devices, such as adapters, are not supported.

d. Provides limited inventory information

This chapter discusses each of the tools capabilities in more detail on the following pages:

- ▶ 7.2, “Lenovo XClarity Administrator” on page 225
- ▶ 7.3, “Lenovo XClarity Integrators” on page 236
- ▶ 7.4, “Lenovo XClarity Energy Manager” on page 240
- ▶ 7.5, “Lenovo XClarity Controller management functions” on page 263
- ▶ 7.6, “Lenovo XClarity Essentials” on page 271
- ▶ 7.7, “LCD system information display panel” on page 272

7.2 Lenovo XClarity Administrator

Lenovo XClarity Administrator (LXCA) is a centralized resource management solution that reduces complexity, speeds response, and enhances the availability of Lenovo server systems and solutions. LXCA provides agent-free hardware management for Lenovo ThinkSystem servers, including the SR950 server.

LXCA 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.

LXCA 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

In this section, we describe key functions and tasks of LXCA that are relevant to the SR950 system. This section includes the following topics:

- ▶ 7.2.1, “Discovering the SR950 server” on page 226
- ▶ 7.2.2, “Remote control” on page 230
- ▶ 7.2.3, “Hardware monitoring” on page 232
- ▶ 7.2.4, “Updating firmware” on page 233
- ▶ 7.2.5, “Operating system deployment” on page 233
- ▶ 7.2.6, “Call home service” on page 233

The Lenovo XClarity Administrator dashboard is shown in Figure 7-1 on page 226.

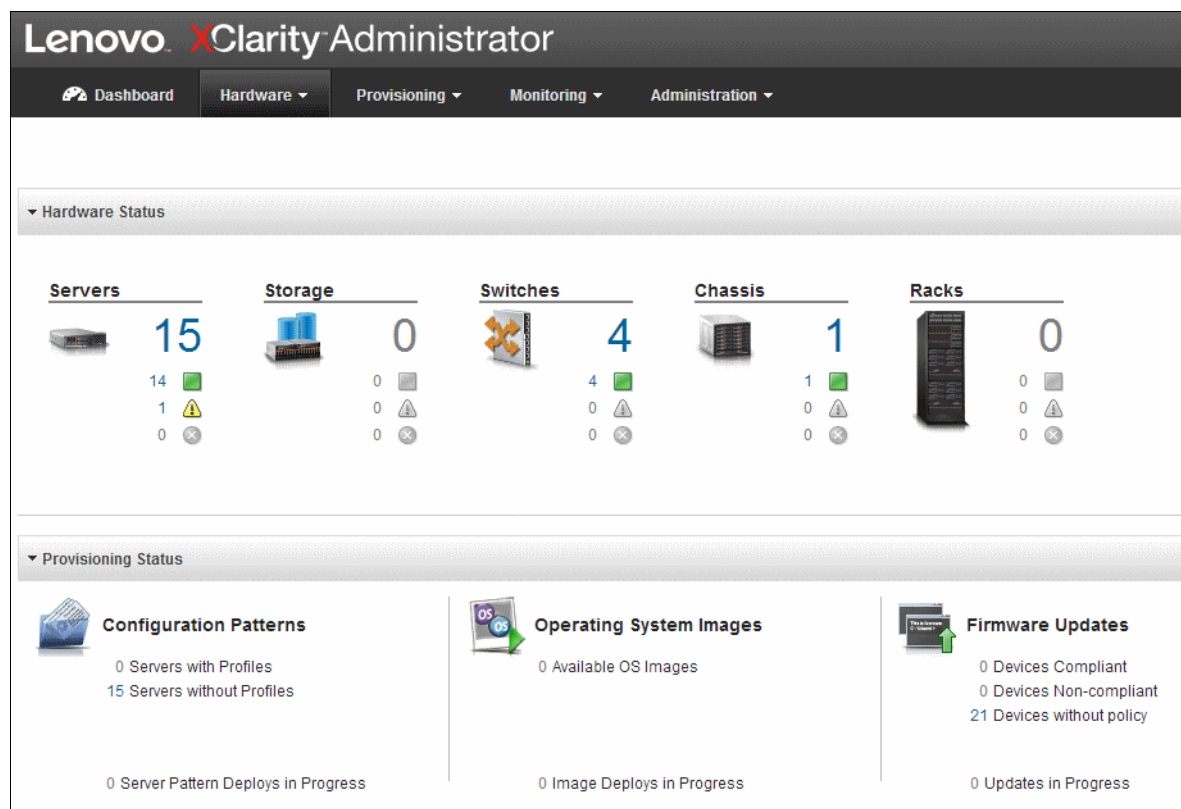


Figure 7-1 Lenovo XClarity Administrator dashboard

7.2.1 Discovering the SR950 server

After the XCC is configured on the network, (refer to “Configuring a dedicated or shared network port” on page 144), you can discover the system from within LXCA. Once XCC has been added to the LXCA console, you can perform management tasks against it.

Complete the following steps to discover and add the XCC to the LXCA console:

1. Log in to the LXCA web interface by browsing to the following web page, 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://172.16.36.15`

If you are logging in for the first time, the **Discover and Manage New Devices** page is displayed, as shown in Figure 7-2.

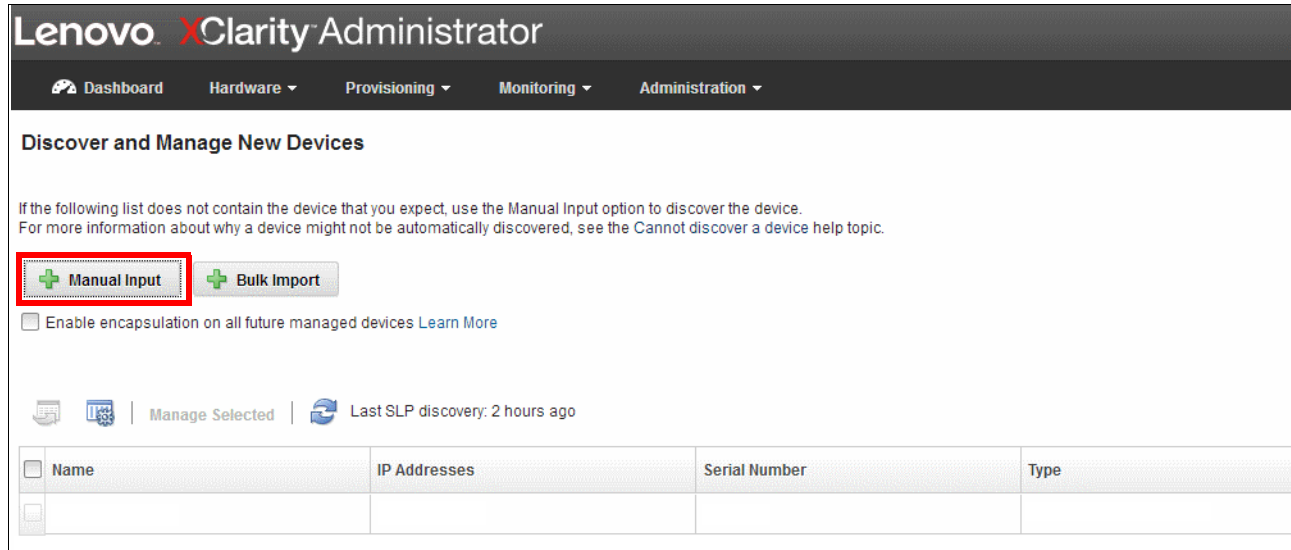


Figure 7-2 Discover and Manage New Devices page

2. If you are not logging in for the first time, the dashboard will display. To get to the **Discover and Manage New Devices** page, click **Hardware** → **Discover and Manage New Devices**, as shown in Figure 7-3.



Figure 7-3 Discover and Manager New Devices

3. To discover and start managing the SR950, click the **Manual Input** option as shown in Figure 7-4 on page 228. You can discover a single server or multiple servers at once.

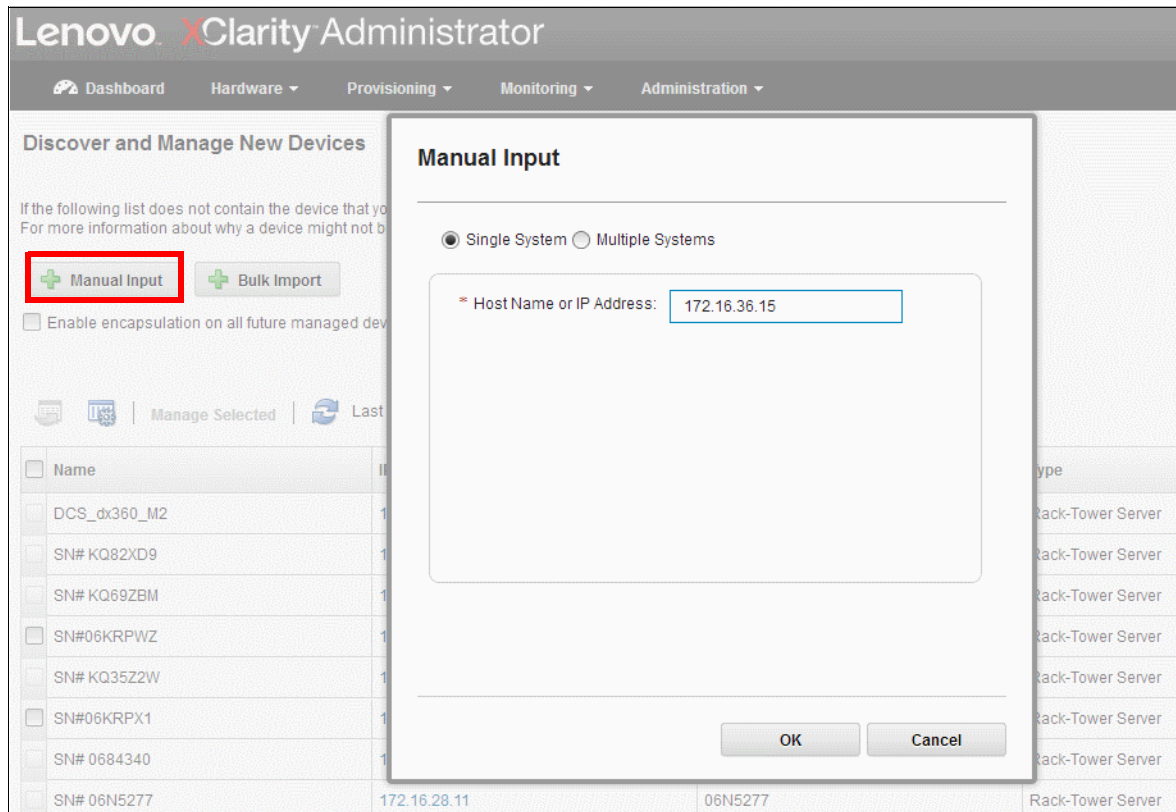


Figure 7-4 Discovering the SR950

- Once the server has been discovered, LXCA will prompt you to enter in the XCC user ID and password credentials, as shown in Figure 7-5. For additional information on the XCC, refer to 5.2, “Preparing the hardware with XClarity Controller” on page 142.

Manage

is going to be managed.

Specify server management module credentials

USERID

☒ Do not create a recovery account and leave all local users enabled.
☐ Create a recovery account and disable all local users.
☐ Force management even if the system is being managed by another instance of Lenovo® XClarity Administrator
When force management, need to use the Recovery-id management.

This will perform the following actions against server:

1. Configure NTP clients to use the NTP settings from Lenovo® XClarity Administrator
2. Configure for centralized user management

Figure 7-5 Input the XCC user ID and password credentials

5. When the credentials have been entered, click the **Manage** button to start managing the server. LXCA will look for the IP address on the network as shown in Figure 7-6 and add it to the managed server list.

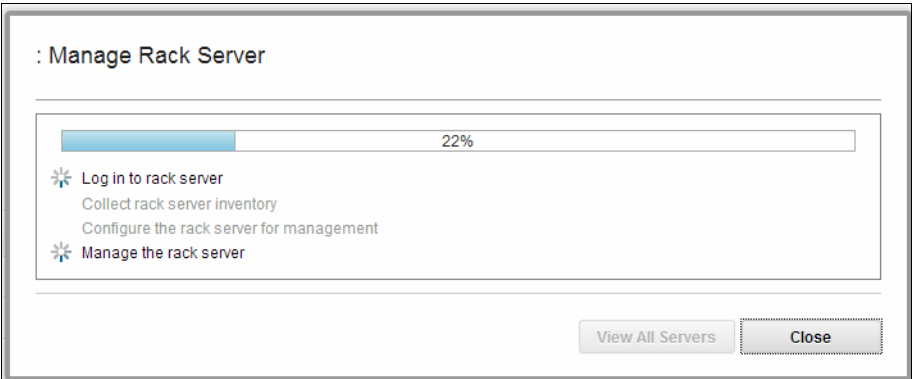


Figure 7-6 Discovering the SR950 server

6. When the discovery process finishes and LXCA discovers the XCC, you see a confirmation message, as shown in Figure 7-7.




Figure 7-7 Successful discovery

7. Close the box and click **Hardware** → **Servers**, or click the **View All Servers** button to see the SR950 in the managed servers list, as shown in Figure 7-8.

The screenshot shows the Lenovo XClarity Administrator interface. The top navigation bar includes 'Dashboard', 'Hardware', 'Provisioning', 'Monitoring', and 'Administration'. The main section is titled 'Servers'. Below the title, there are icons for various server actions and a 'Filter By' dropdown. The main content is a table with the following columns: 'Server', 'Status', 'Power', 'IP Addresses', 'Rack Name/Unit', 'Chassis/Bay', and 'Product'. The first row of the table shows a server named 'SR950' with a status of 'Normal' (indicated by a green square icon) and power status 'On' (indicated by a green square icon). The IP addresses listed are '172.16.36.15, 169.254.95...'. The rack name/unit is 'Unassigned / Unassigned', and the chassis/bay is 'Not Applicable / Not Applicable'.

Server	Status	Power	IP Addresses	Rack Name/Unit	Chassis/Bay	Product
SR950	Normal	On	172.16.36.15, 169.254.95...	Unassigned / Unassigned	Not Applicable / Not Applicable	Lenovo

Figure 7-8 Discovered servers list

8. Set the servers information **Actions** → **Inventory** → **Edit Properties**, as shown in Figure 7-9 on page 230.

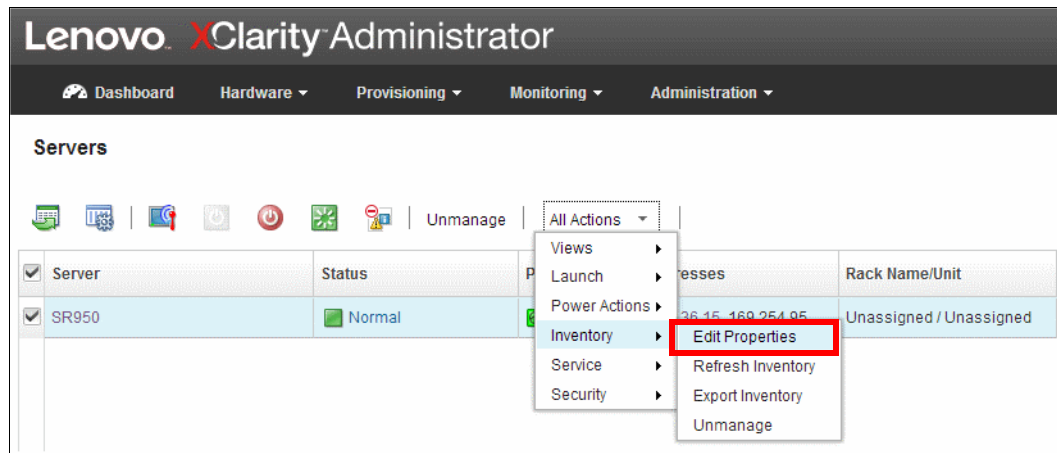


Figure 7-9 Edit properties of the server

The properties box will display as shown in Figure 7-10.

The 'Edit Properties: SR950' dialog box is shown. It contains a warning message: 'Some of the information below will be saved on the device and some will be saved in Lenovo SR950 inventory. It might take a few minutes for your updates to appear.' Below this, there are input fields for:

- Name: SR950
- Support Contact: Lenovo Professional Services
- Location: 1
- Room: (empty)
- Rack: (empty)
- Lowest Rack Unit: 0
- Description: Rack 26, U16

 At the bottom, there are 'Save' and 'Cancel' buttons.

Figure 7-10 Edit the properties of the server

9. Edit the required properties and click **Save**. The SR950 server is now ready to be managed by LXCA.

7.2.2 Remote control

XClarity Administrator provides remote control features for all managed hardware. To open a remote control session to the SR950 server, use the following procedure:

1. Select **Hardware** → **Servers**, select the server from the list, and click **All Actions** → **Launch** → **Launch Remote Control**, as shown in Figure 7-11 on page 231.

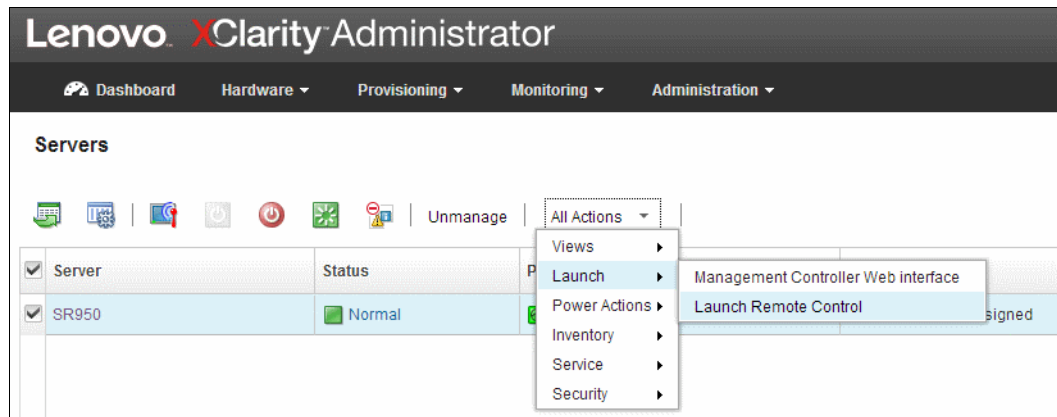


Figure 7-11 Launch remote control session

2. Select either a **Single User** or **Multiple User** from the prompt. Click **Launch remote console**, as shown in Figure 7-12.




Figure 7-12 Single or multiple user console

LXCA also allows you to launch XCC directly from the LXCA interface. To open XCC for the SR950 server, follow these steps:

1. Select **Hardware** → **Servers**, select the server from the list.
2. Click **All Actions** → **Launch** → **Management Controller Web interface**, as shown in Figure 7-13.

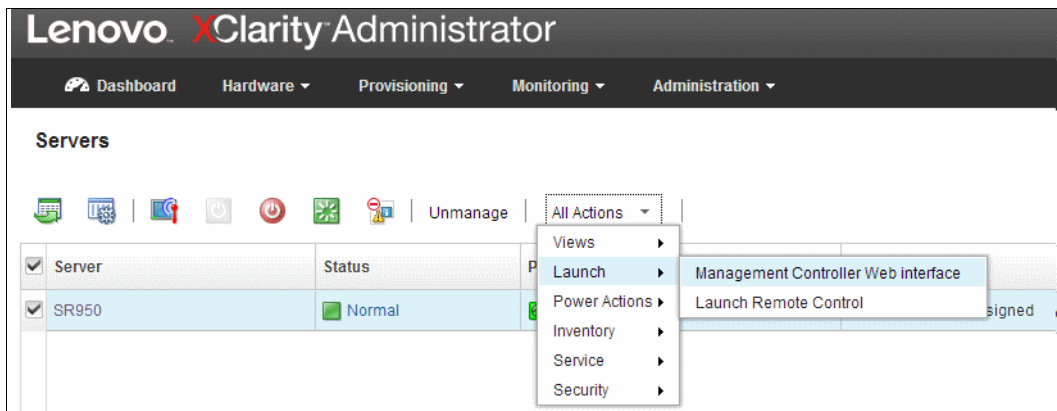


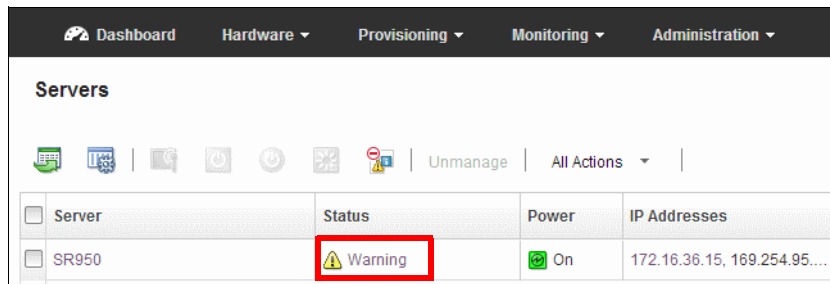
Figure 7-13 Launch the SR950 XCC

The XCC will now launch in a separate web browser window. Ensure pop-ups are not disabled for the browser.

7.2.3 Hardware monitoring

LXCA provides a centralized view of events and alerts that are generated from managed endpoints, such as the SR950 server. When an issue is detected by the XCC, an event is passed to LXCA. That event is then displayed in the alerts list that is available within the user interface.

To look at a specific alert select **Hardware** → **Servers** and click on the **Warning** (for example) listed under the servers **Status**, as shown in Figure 7-14.

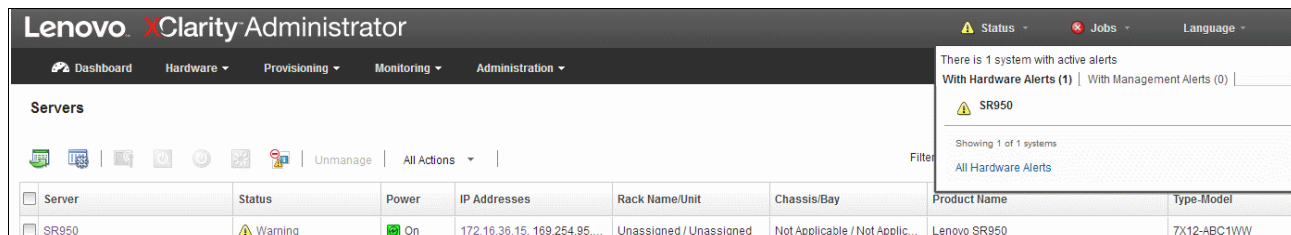


The screenshot shows the 'Servers' page in the LXCA interface. A table lists server status. The 'SR950' server is highlighted with a red box around the 'Warning' status icon and text.

Server	Status	Power	IP Addresses
SR950	Warning	On	172.16.36.15, 169.254.95...

Figure 7-14 Warning status

Alternatively you can look at all hardware alerts for all servers by selecting **Status All** → **Hardware Alerts**, as shown in Figure 7-15.

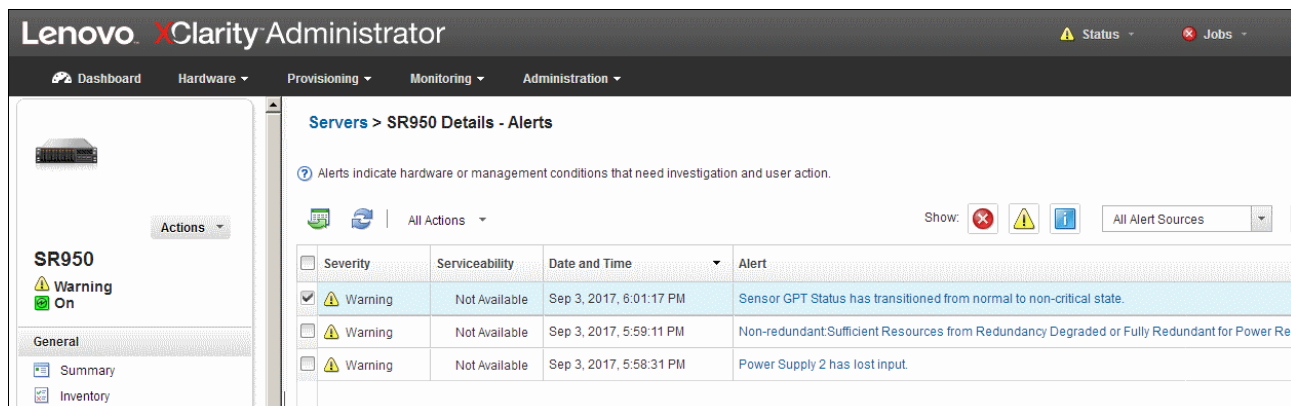


The screenshot shows the 'Hardware Alerts' menu in the LXCA interface. A dropdown menu is open, showing 'All Hardware Alerts' and 'Showing 1 of 1 systems'.

Server	Status	Power	IP Addresses	Rack Name/Unit	Chassis/Bay	Product Name	Type-Model
SR950	Warning	On	172.16.36.15, 169.254.95...	Unassigned / Unassigned	Not Applicable / Not Applicable	Lenovo SR950	7X12-ABC1WW

Figure 7-15 All hardware alerts menu

Figure 7-16 shows the **Alerts** page for the SR950. In this example, a power supply has lost power and a warning message is being displayed.



The screenshot shows the 'Alerts' page for the SR950 server. A table lists alerts, including 'Sensor GPT Status has transitioned from normal to non-critical state' and 'Power Supply 2 has lost input'.

Severity	Serviceability	Date and Time	Alert
Warning	Not Available	Sep 3, 2017, 6:01:17 PM	Sensor GPT Status has transitioned from normal to non-critical state.
Warning	Not Available	Sep 3, 2017, 5:59:11 PM	Non-redundant:Sufficient Resources from Redundancy Degraded or Fully Redundant for Power Res
Warning	Not Available	Sep 3, 2017, 5:58:31 PM	Power Supply 2 has lost input.

Figure 7-16 Warning message on the SR950

You can also open the system event log of the server, check light path states or power consumption. For example, to open the Power Consumption History graph, browse to Power and Thermal of the Status and Health section, as shown in Figure 7-17 on page 233.

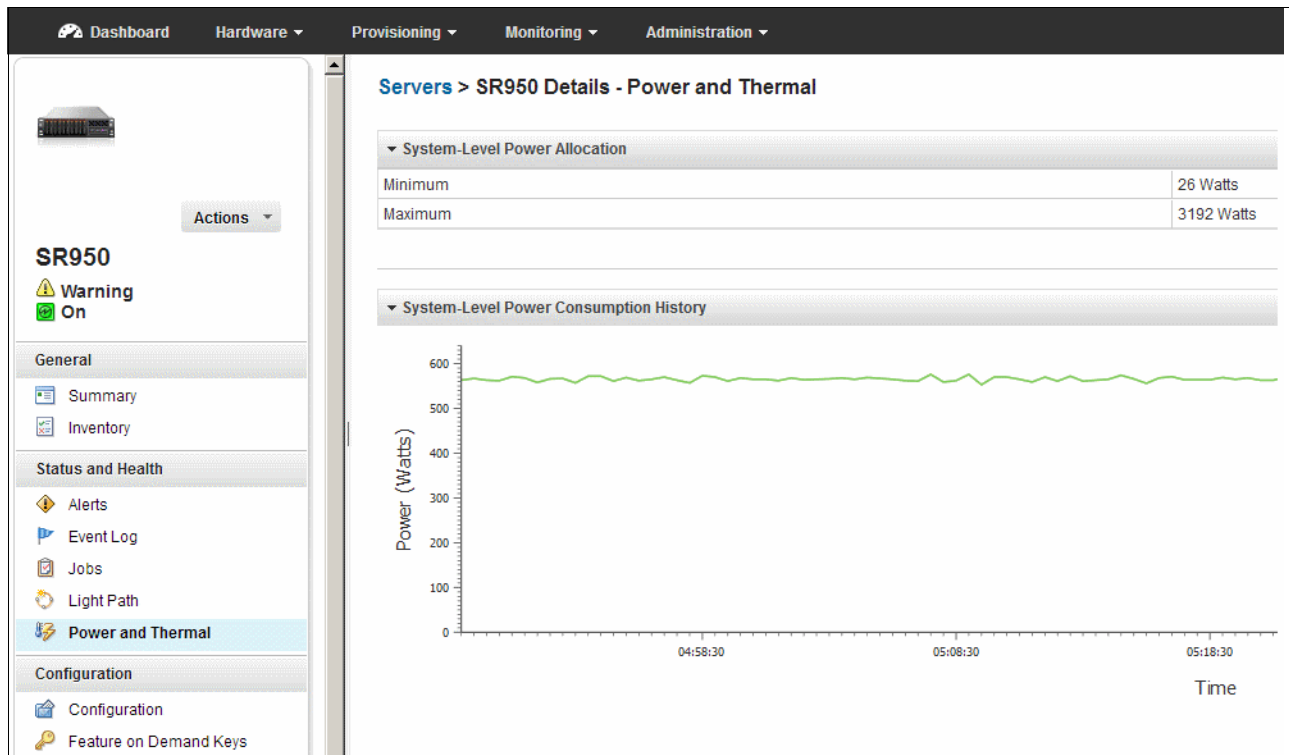


Figure 7-17 Power and Thermal details of the SR950

7.2.4 Updating firmware

LXCA can be used to update the firmware of the SR950 server. All components can be updated using this tool.

For information on how to apply firmware to the server using LXCA, refer to 5.3.4, “Lenovo XClarity Administrator” on page 173 for additional details.

7.2.5 Operating system deployment

LXCA 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 LXCA, refer to 6.2.3, “Installing an OS with XClarity Administrator (LXCA)” on page 195.

LXCA also integrates with managers, such as VMware vSphere and Microsoft System Center. This capability is discussed in 7.3, “Lenovo XClarity Integrators” on page 236.

7.2.6 Call home service

You can configure LXCA 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 ThinkSystem server. This automatic support notification is called *Call Home*.

Note: When call home is configured and enabled in LXCA, call home is disabled on all managed chassis and servers to avoid duplicate problem records being created.

LXCA 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 LXCA:

1. From the LXCA menu bar, click **Administration** → **Service and Support**, as shown in Figure 7-18.

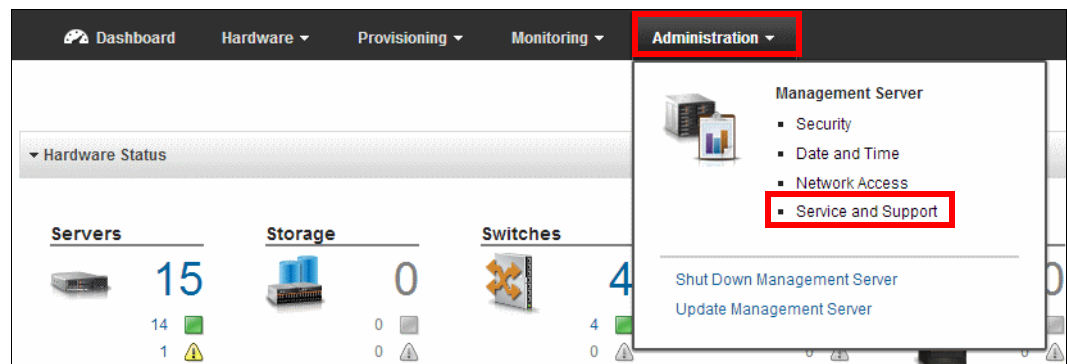


Figure 7-18

2. Click the **Call Home Configuration** tab as shown in Figure 7-19 on page 235.

Service and Support

From this page, you can download diagnostic files and collect diagnostics from managed devices. You can also work with Call Home.

Management Server Files | Endpoint Service Data | Device Actions | **Call Home Configuration** | Service Forwarders | Problem Record Status | Server Logging

From this page, you can create a service forwarder for Call Home that automatically sends service data for any managed device to Lenovo Support when certain serviceable events occur. You can enable the Default Call Home service forwarder from the Service Forwarder tab.

Default Call Home Forwarder

Forwarder State: **Not Set**

Configure Call Home

* Contact Name

* Company Name

* Country

* Email

* Phone Number

* Street Address

* City

* State or Province

Figure 7-19 Call Home Configuration screen

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 downloads and additional information for Lenovo XClarity Administrator see the following links:

- Download LXCA
<https://datacentersupport.lenovo.com/no/en/documents/lnvo-lxcaupd>
- Lenovo XClarity Administrator Planning and Implementation Guide
<http://lenovopress.com/sg248296>
- Online documentation
http://sysmgmt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2Faug_product_page.html

7.3 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 non-disruptive server updates
- ▶ Perform non-disruptive 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 (free download, support requires XClarity Pro license)
<https://datacentersupport.lenovo.com/documents/LNVO-VMWARE>
- ▶ Lenovo XClarity Integrator for VMware vRealize Orchestrator (free download)
<https://datacentersupport.lenovo.com/documents/LNVO-VMRO>
- ▶ Lenovo XClarity Integrator for VMware vRealize Automation (free download)
<https://marketplace.vmware.com/vsx/solutions/xclarity-integrator-for-vrealize-automation>
- ▶ Lenovo XClarity Integrator for VMware vRealize Log Insight (free download)
<https://marketplace.vmware.com/vsx/solutions/lenovo-networking-content-pack-for-vmware-vrealize-log-insight>
- ▶ Lenovo XClarity Integrator for Microsoft System Center (free download, support requires XClarity Pro license)
<https://datacentersupport.lenovo.com/documents/LNVO-MANAGE>

Support for Lenovo XClarity Integrators for VMware vCenter and Microsoft System Center is included in Lenovo XClarity Pro version.

Note: The Lenovo XClarity Integrator requires a Lenovo customized ESXi version. The latest Lenovo customized version can be downloaded from this website:

https://my.vmware.com/web/vmware/info/slug/datacenter_cloud_infrastructure/vmware_vsphere/6_5#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 the following tables.

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)

Lenovo XClarity Integrator per managed server, for Microsoft System Center or VMware vCenter for United States, Asia Pacific, Canada, and Japan	Part number
Lenovo XClarity Integrator for MSSC, Per Managed Server w/1Yr S&S	00MT275
Lenovo XClarity Integrator for MSSC, Per Managed Server w/3Yr S&S	00MT276
Lenovo XClarity Integrator for MSSC, Per Managed Server w/5Yr S&S	00MT277
Lenovo XClarity Integrator for MSSC, w/IMMV2ADv, Per Managed Server w/1Yr S&S	00MT278
Lenovo XClarity Integrator for MSSC, w/IMMV2ADv, Per Managed Server w/3Yr S&S	00MT279
Lenovo XClarity Integrator for MSSC, w/IMMV2ADv, Per Managed Server w/5Yr S&S	00MT280
Lenovo XClarity Integrator for VMware vCenter, Per Managed Server w/1Yr S&S	00MT281
Lenovo XClarity Integrator for VMware vCenter, Per Managed Server w/3Yr S&S	00MT282
Lenovo XClarity Integrator for VMware vCenter, Per Managed Server w/5Yr S&S	00MT283
Lenovo XClarity Integrator f/VMware vCtr w/IMMv2Adv, Per Managed Server w/1Yr S&S	00MT284
Lenovo XClarity Integrator f/VMware vCtr w/IMMv2Adv, Per Managed Server w/3Yr S&S	00MT285
Lenovo XClarity Integrator f/VMware vCtr w/IMMv2Adv, Per Managed Server w/5Yr S&S	00MT286

Table 7-5 Lenovo XClarity Integrator part numbers per managed server (EMEA and Latin America)

Lenovo XClarity Integrator per managed server, for Microsoft System Center or VMware vCenter for Europe Middle East, Africa, and Latin America	Part number
Lenovo XClarity Integrator for MSSC, Per Managed Server w/1Yr S&S	00MT287
Lenovo XClarity Integrator for MSSC, Per Managed Server w/3Yr S&S	00MT288
Lenovo XClarity Integrator for MSSC, Per Managed Server w/5Yr S&S	00MT289
Lenovo XClarity Integrator for MSSC, w/IMMV2ADv, Per Managed Server w/1Yr S&S	00MT290
Lenovo XClarity Integrator for MSSC, w/IMMV2ADv, Per Managed Server w/3Yr S&S	00MT291
Lenovo XClarity Integrator for MSSC, w/IMMV2ADv, Per Managed Server w/5Yr S&S	00MT292
Lenovo XClarity Integrator for VMware vCenter, Per Managed Server w/1Yr S&S	00MT293
Lenovo XClarity Integrator for VMware vCenter, Per Managed Server w/3Yr S&S	00MT294
Lenovo XClarity Integrator for VMware vCenter, Per Managed Server w/5Yr S&S	00MT295
Lenovo XClarity Integrator f/VMware vCtr w/IMMv2Adv, Per Managed Server w/1Yr S&S	00MT296
Lenovo XClarity Integrator f/VMware vCtr w/IMMv2Adv, Per Managed Server w/3Yr S&S	00MT297
Lenovo XClarity Integrator f/VMware vCtr w/IMMv2Adv, Per Managed Server w/5Yr S&S	00MT298

Lenovo XClarity Integrator provides ease of management in the following ways:

- ▶ **Non-disruptive server updates:** Eliminate maintenance downtime through non-disruptive 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 7-20.




Figure 7-20 Start a firmware update with upward integration module (UIM)

2. Any virtual machines (VMs) that are running on the first server are moved from the first server to another server, as shown in Figure 7-21.




Figure 7-21 UIM evacuates first servers VMs for updating

3. The server is taken offline and the updates are applied.
4. 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 7-22.




Figure 7-22 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:

- A RAS event that is detected on the server (see Figure 7-23) is reported up the chain to Microsoft System Center, VMware vCenter or XClarity.




Figure 7-23 RAS event reports up the chain

2. The server is emptied and the workload resumes on a different server, as shown in Figure 7-24.




Figure 7-24 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, v5.1.0:
<https://support.lenovo.com/documents/LNVO-VMWARE>
- ▶ Lenovo XClarity Integrator Offerings for Microsoft System Center Management Solutions:
<https://support.lenovo.com/documents/LNVO-MANAGE>

7.4 Lenovo XClarity Energy Manager

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:

<https://www3.lenovo.com/us/en/data-center/software/systems-management/xclarity-energy-manager/>

The following section discusses licensing, system requirements, and how to download, install, and set up Energy Manager.

7.4.1 Licensing

Energy Manager requires a license to be purchased per device for Lenovo and non-Lenovo servers. Table 7-6 shows the ordering part number information for a single-device, 5-node and 50-device licenses.

Note: The SR950 ships standard with XClarity Controller Enterprise, which includes a license for XClarity Energy Manager.

Table 7-6 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 shown in Figure 7-25.

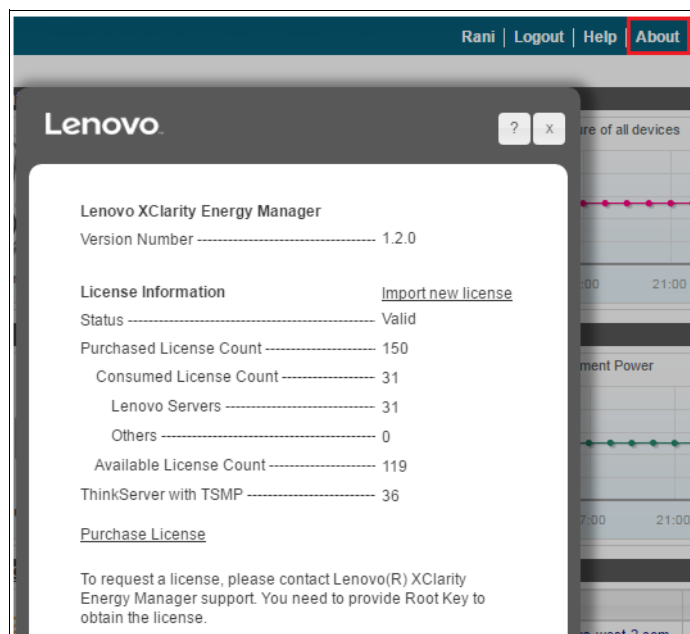


Figure 7-25 Energy Manager version and licensing

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.

7.4.2 System requirements

The Energy Manager Server where Energy Manager is installed. The following operating systems are supported:

- ▶ Microsoft Windows Server 2008
- ▶ Microsoft Windows Server 2008 R2
- ▶ Microsoft Windows Server 2012
- ▶ Microsoft Windows Server 2012 R2
- ▶ Microsoft Windows Server 2016
- ▶ Red Hat Enterprise Linux 6
- ▶ Red Hat Enterprise Linux 7
- ▶ SUSE Linux Enterprise Server 11
- ▶ SUSE Linux Enterprise Server 12 (

It is recommended to install the software on a system with at least:

- ▶ A dual-core processor of 2.6 GHz or higher
- ▶ 8 GB RAM
- ▶ 80 GB of hard drive space

To access the Energy Manager interface, the following web browsers are supported:

- ▶ Mozilla Firefox 27
- ▶ Google Chrome 31
- ▶ Microsoft Internet Explorer 11
- ▶ Microsoft Edge

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.

7.4.3 Download and install

You can download the Lenovo XClarity Energy Manager from the following web page:

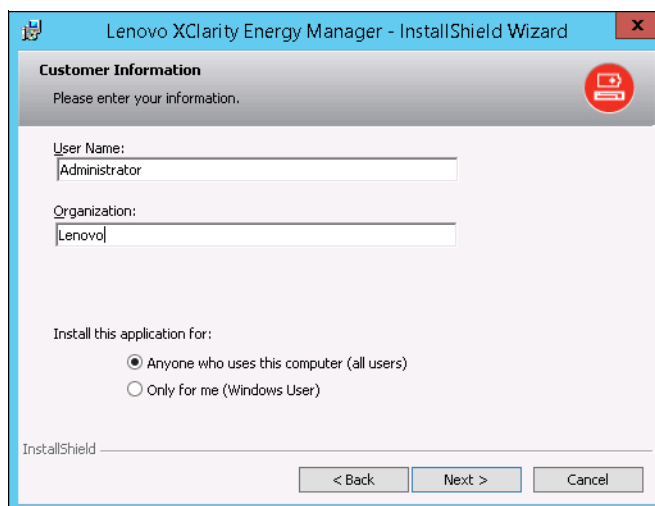
<https://support.lenovo.com/downloads/ds101160>

For Windows, run the installer and follow the wizard for the installation. For Linux, decompress the package and then launch the executable for the installation. The software will be installed on the system 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 7-26. Click **Next**.



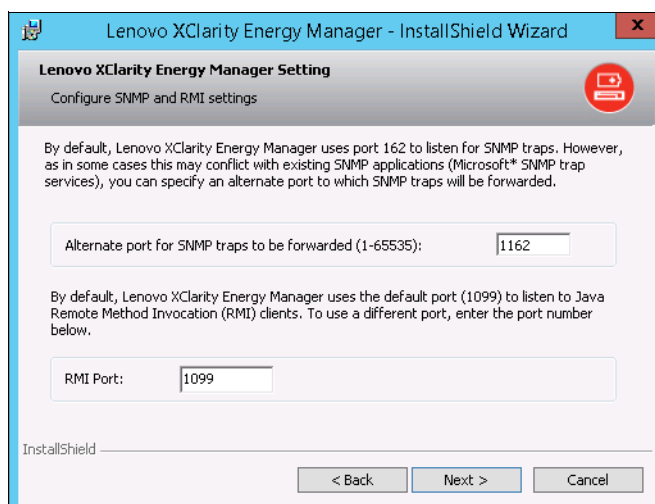
The screenshot shows the 'Customer Information' window of the 'Lenovo XClarity Energy Manager - InstallShield Wizard'. The window has a title bar with the application name and a close button. Below the title bar is a header area with the text 'Customer Information' and 'Please enter your information.' To the right of the header is a red circular icon with a white arrow pointing right. The main area contains two text input fields: 'User Name:' with the value 'Administrator' and 'Organization:' with the value 'Lenovo'. Below these fields are two radio buttons for installation scope: 'Anyone who uses this computer (all users)' (selected) and 'Only for me (Windows User)'. At the bottom left is the 'InstallShield' logo. At the bottom right are three buttons: '< Back', 'Next >', and 'Cancel'.

Figure 7-26 Customer information

5. Specify the directory where the software is to be installed and click **Next**.
6. Specify the SNMP and Java RMI ports to be used in Figure 7-27. 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 7-27 displays the SNMP and RMI settings page.

Click **Next** once the ports have been set.



The screenshot shows the 'Lenovo XClarity Energy Manager Setting' window of the 'Lenovo XClarity Energy Manager - InstallShield Wizard'. The window has a title bar with the application name and a close button. Below the title bar is a header area with the text 'Lenovo XClarity Energy Manager Setting' and 'Configure SNMP and RMI settings'. To the right of the header is a red circular icon with a white arrow pointing right. The main area contains two paragraphs of text. The first paragraph explains that by default, the software uses port 162 for SNMP traps, but an alternate port can be specified. Below this text is a text input field labeled 'Alternate port for SNMP traps to be forwarded (1-65535):' with the value '1162'. The second paragraph explains that by default, the software uses port 1099 for RMI clients, and a different port can be entered. Below this text is a text input field labeled 'RMI Port:' with the value '1099'. At the bottom left is the 'InstallShield' logo. At the bottom right are three buttons: '< Back', 'Next >', and 'Cancel'.

Figure 7-27 SNMP and RMI port settings

7. Specify the port to be used for the web service that will be used to access the Energy Manager interface in Figure 7-28. 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.

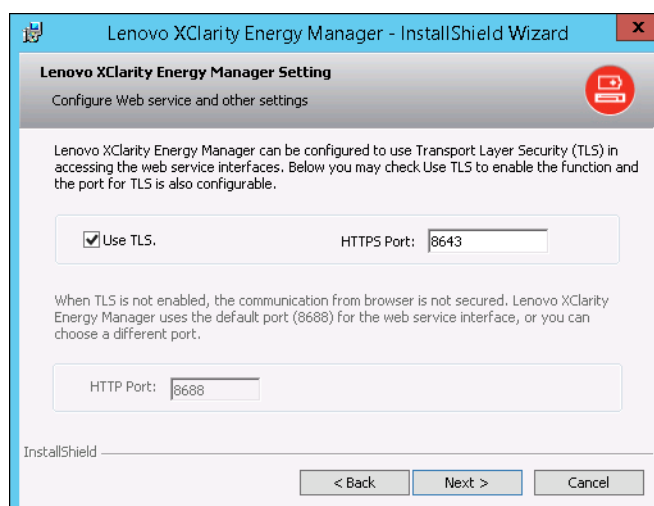


Figure 7-28 Configure web service

8. Set the sampling frequency and granularity of data as shown in Figure 7-29.
- 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 7-29 Set frequency and granularity of data

9. In Figure 7-30, 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 7.4.5, “Setting up Energy Manager” on page 249 for information on the data center hierarchy.

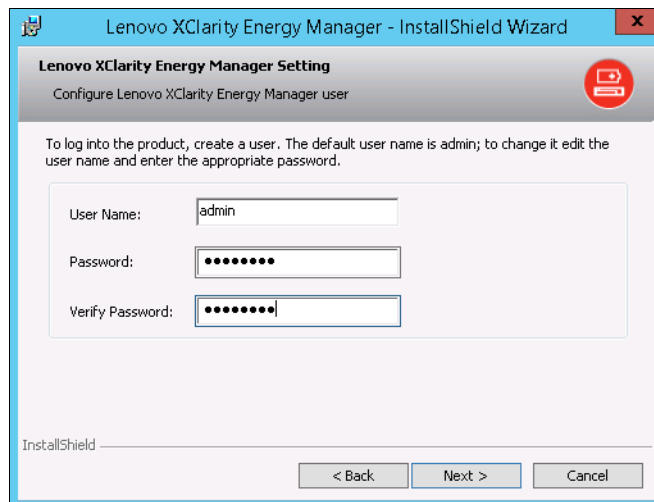


Figure 7-30 Set user name and password

10. Energy Manager uses a Java keystore (JKS) to manage security certificates. Set the JKS username and password as shown in Figure 7-31. 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.

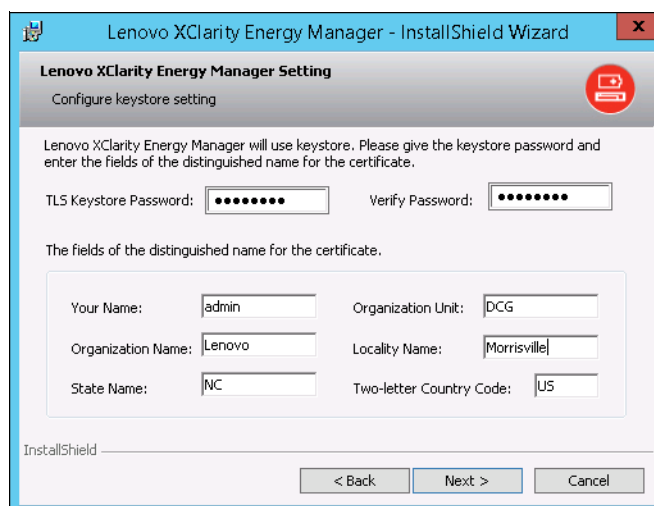


Figure 7-31 Keystore access

11. Energy Manager has an embedded database server as shown in Figure 7-32, you can set the database attributes including username, an open/unused port, password and installation path. Once set, click **Next**.

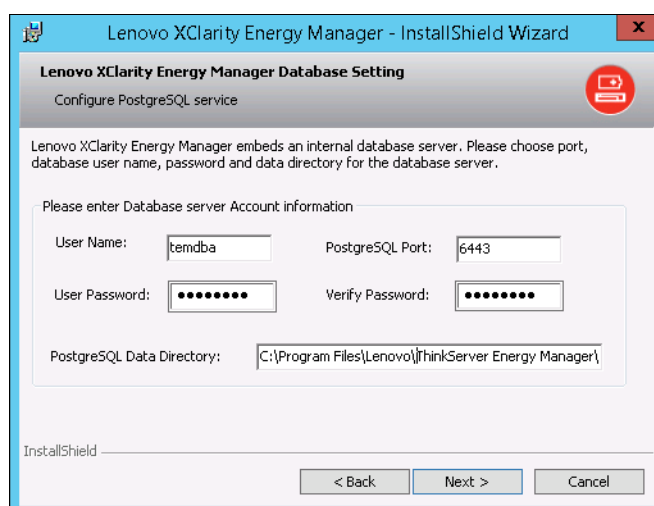


Figure 7-32 Embedded database installation

12. 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 7-33.

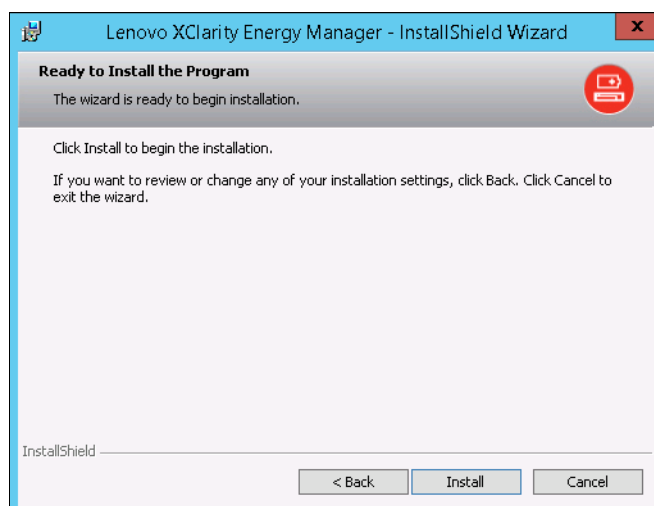


Figure 7-33 Begin installation

13. 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.

7.4.4 Launching 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 supported 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 7-28 on page 244. 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 system.

Once launched the login page will appear as shown in Figure 7-34. Enter the user ID and password that was set during installation (step 9 on page 245).

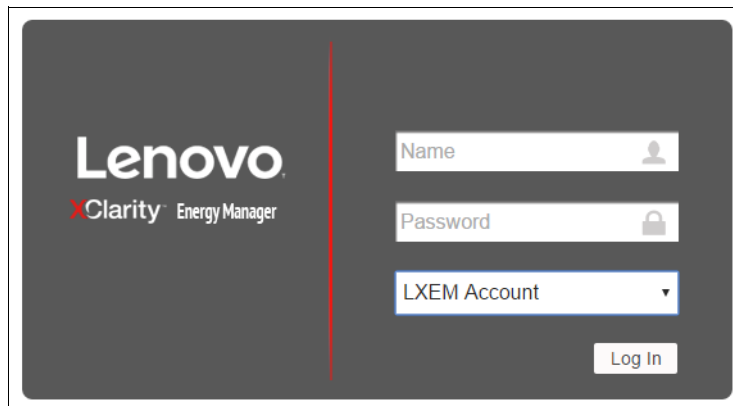
The image shows the login screen for the Lenovo XClarity Energy Manager. On the left side, there is the Lenovo logo above the text 'XClarity Energy Manager'. A vertical red line separates this from the login form on the right. The form contains three input fields: 'Name' with a user icon, 'Password' with a lock icon, and 'LXEM Account' with a dropdown arrow. Below these fields is a 'Log In' button.

Figure 7-34 Energy Manager login screen

Once logged in to Energy Manager, the dashboard is displayed, Figure 7-35. 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 7-35, the gadgets do not currently display any data because we have not yet discovered any devices.

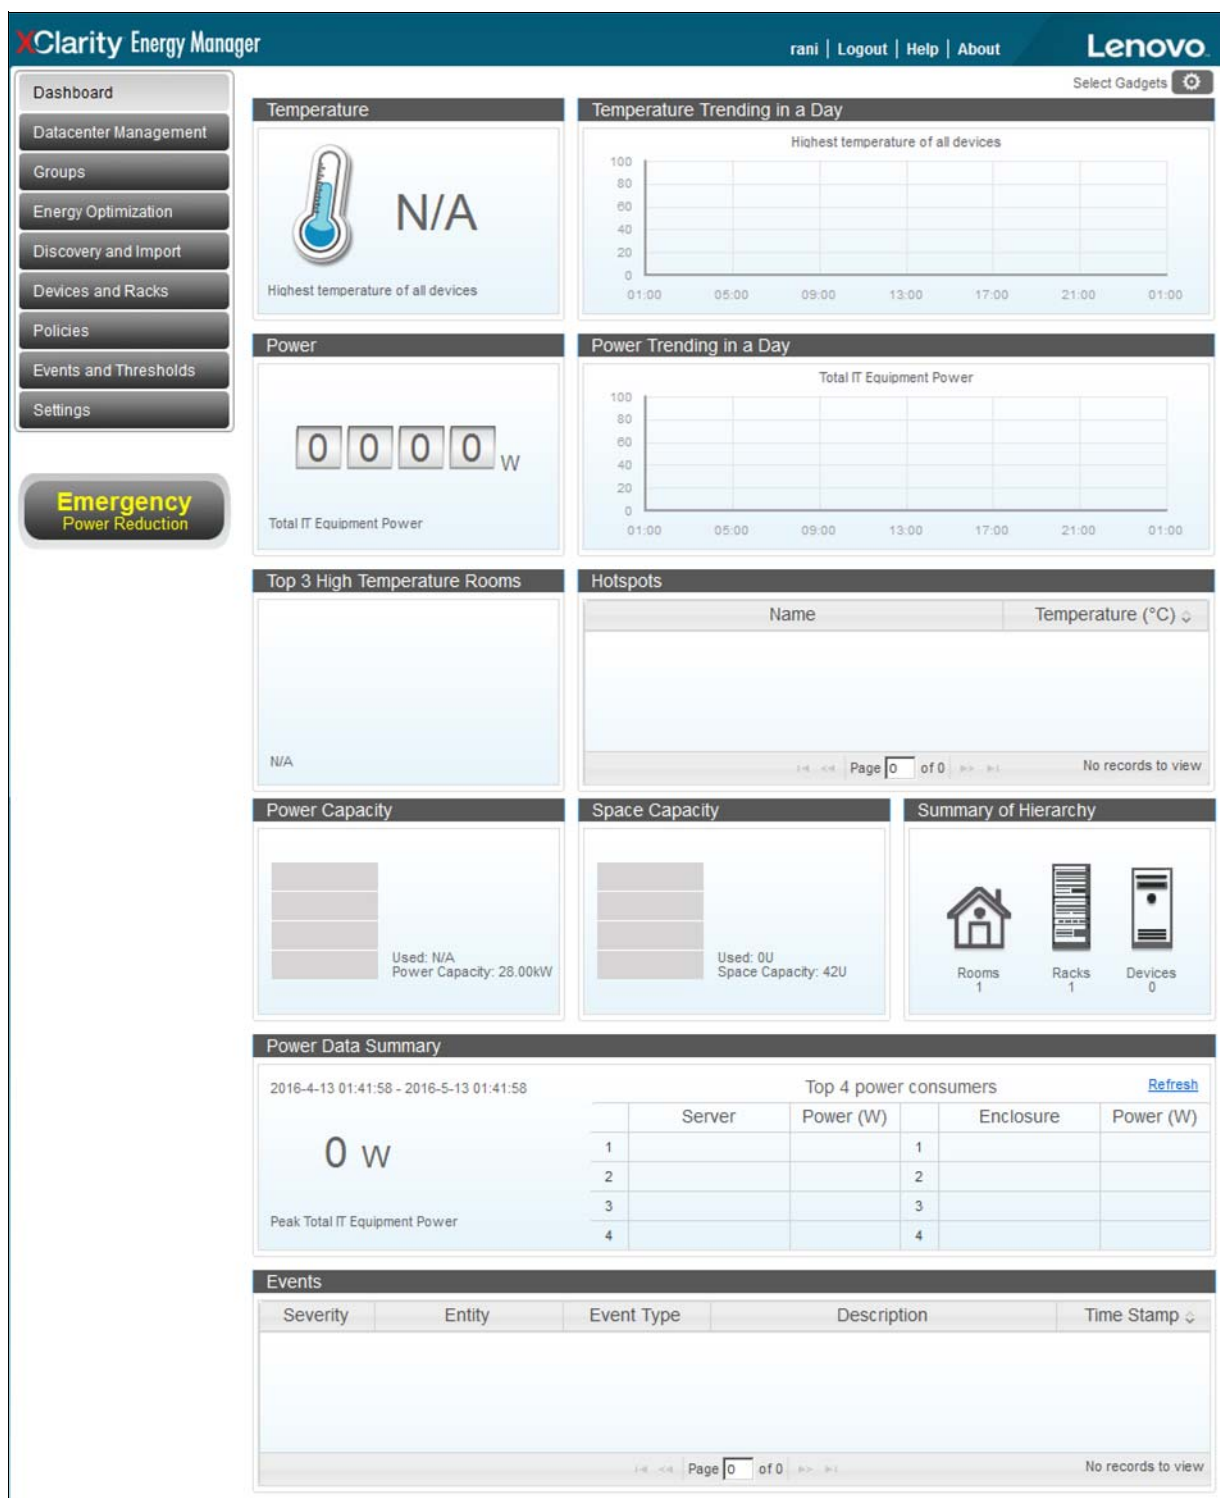


Figure 7-35 Lenovo XClarity Energy Manager Dashboard

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 7-36.

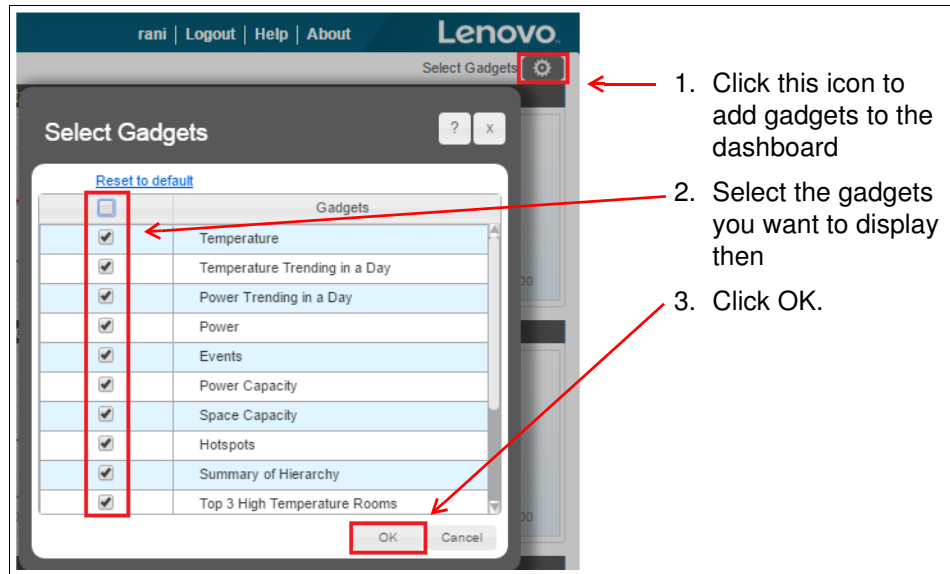


Figure 7-36 Customizing the Dashboard with gadgets

7.4.5 Setting up Energy Manager

To start monitoring and managing power and temperature of devices with 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 shown in Figure 7-37.

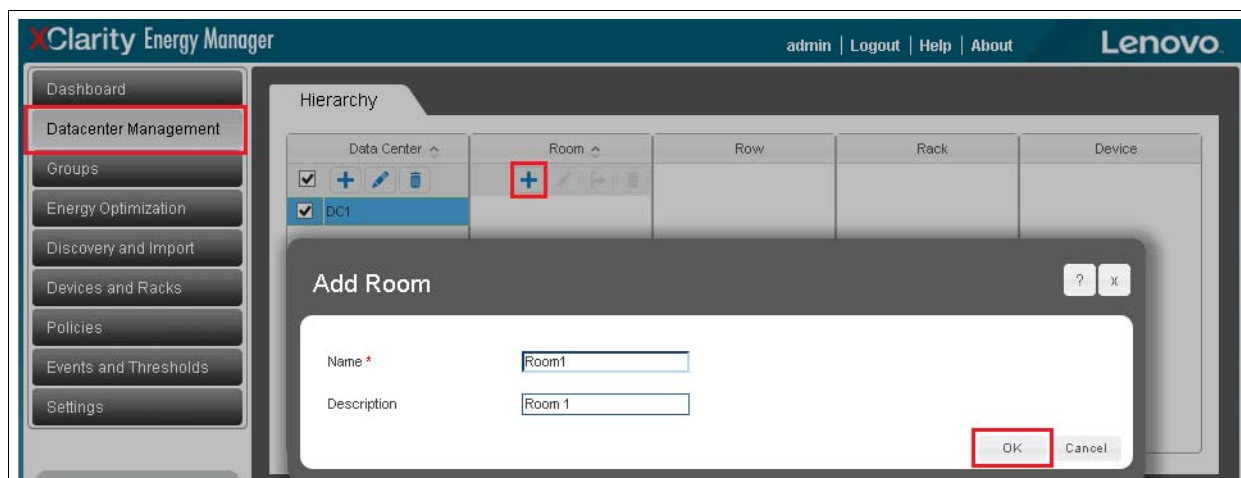


Figure 7-37 Creating a hierarchy

When adding a rack to the hierarchy, enter the total power available (in Watts) within that rack as shown in Figure 7-38. 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://datacentersupport.lenovo.com/no/en/documents/1nvo-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.

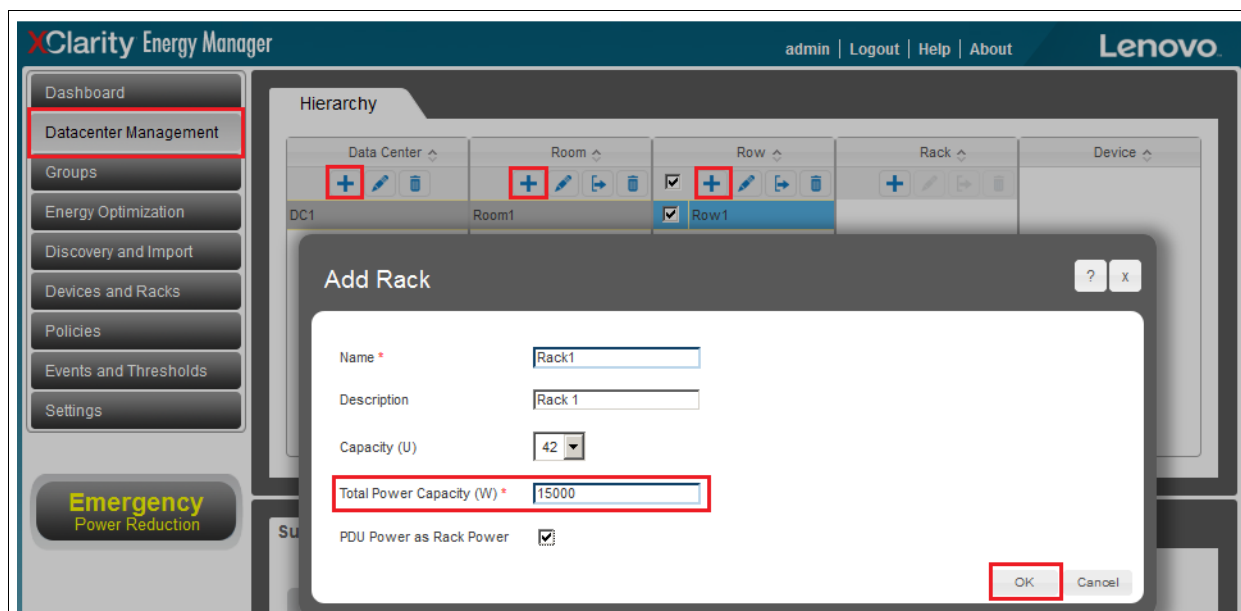


Figure 7-38 Adding a rack to the hierarchy

Discovering devices with Energy Manager

There are two ways to discover devices and add them to the hierarchy.

- **Discovery:**

To discover devices in the network, start from **Discovery and Import** on the left side navigation menu as shown in Figure 7-39.

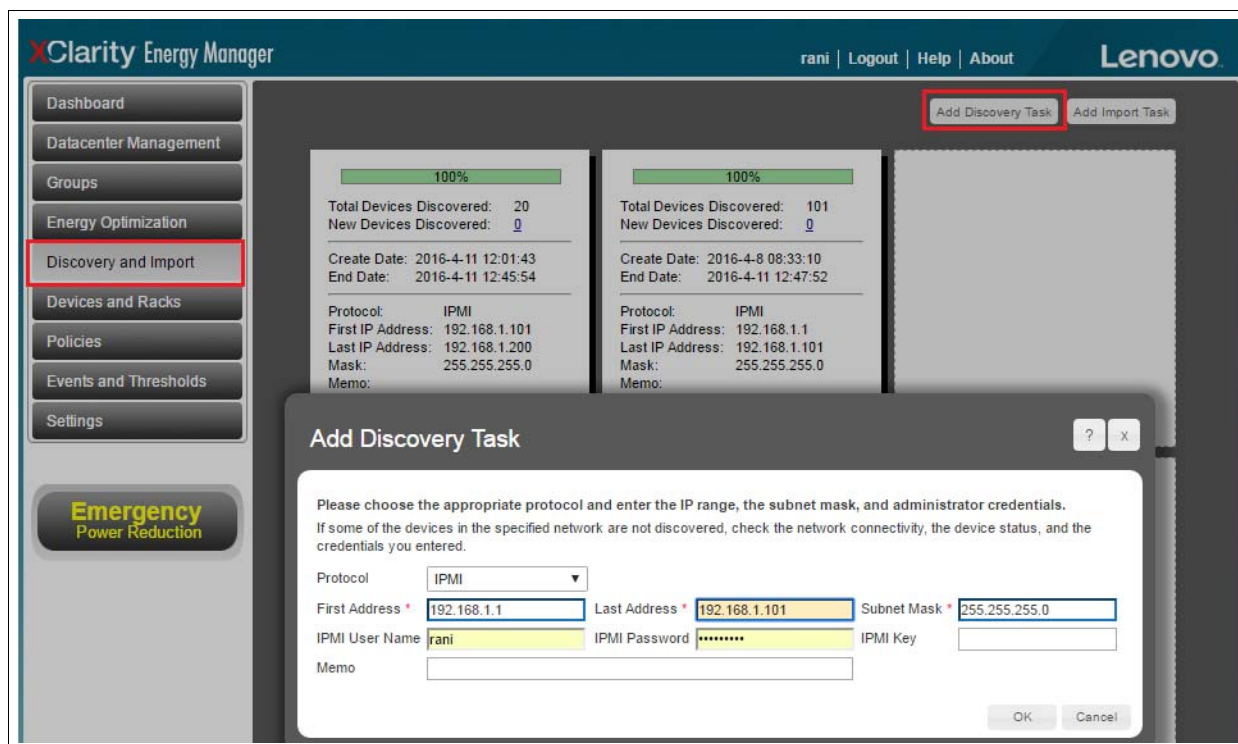


Figure 7-39 Discovery and import

- **Import devices (or an entire hierarchy) from a spreadsheet**

To import devices or the hierarchy from an Excel XLS file, start from the **Import** button in **Devices and Racks** as shown in Figure 7-40.

For the Excel file requirements, refer to the *Lenovo XClarity Energy Manager User Guide*:

<http://support.lenovo.com/us/en/downloads/ds101160>

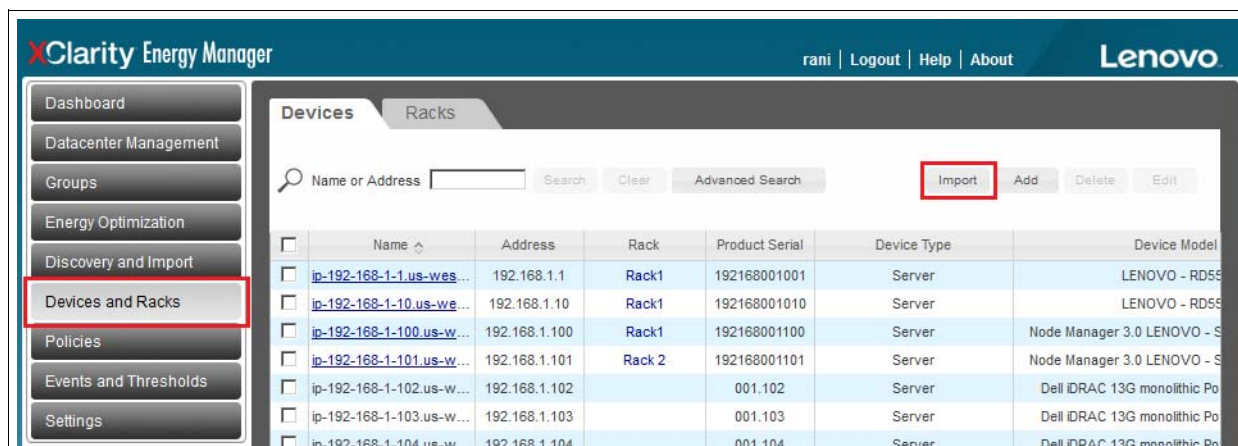


Figure 7-40 Discovery and import from Excel file

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 shown in Figure 7-41.

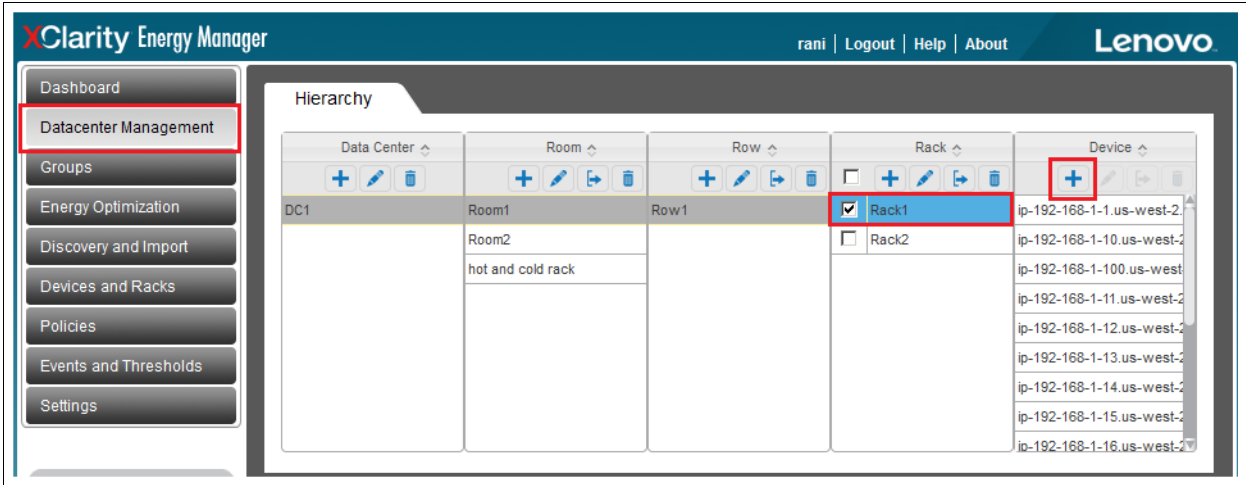


Figure 7-41 Adding discovered devices to a rack

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 7-42.

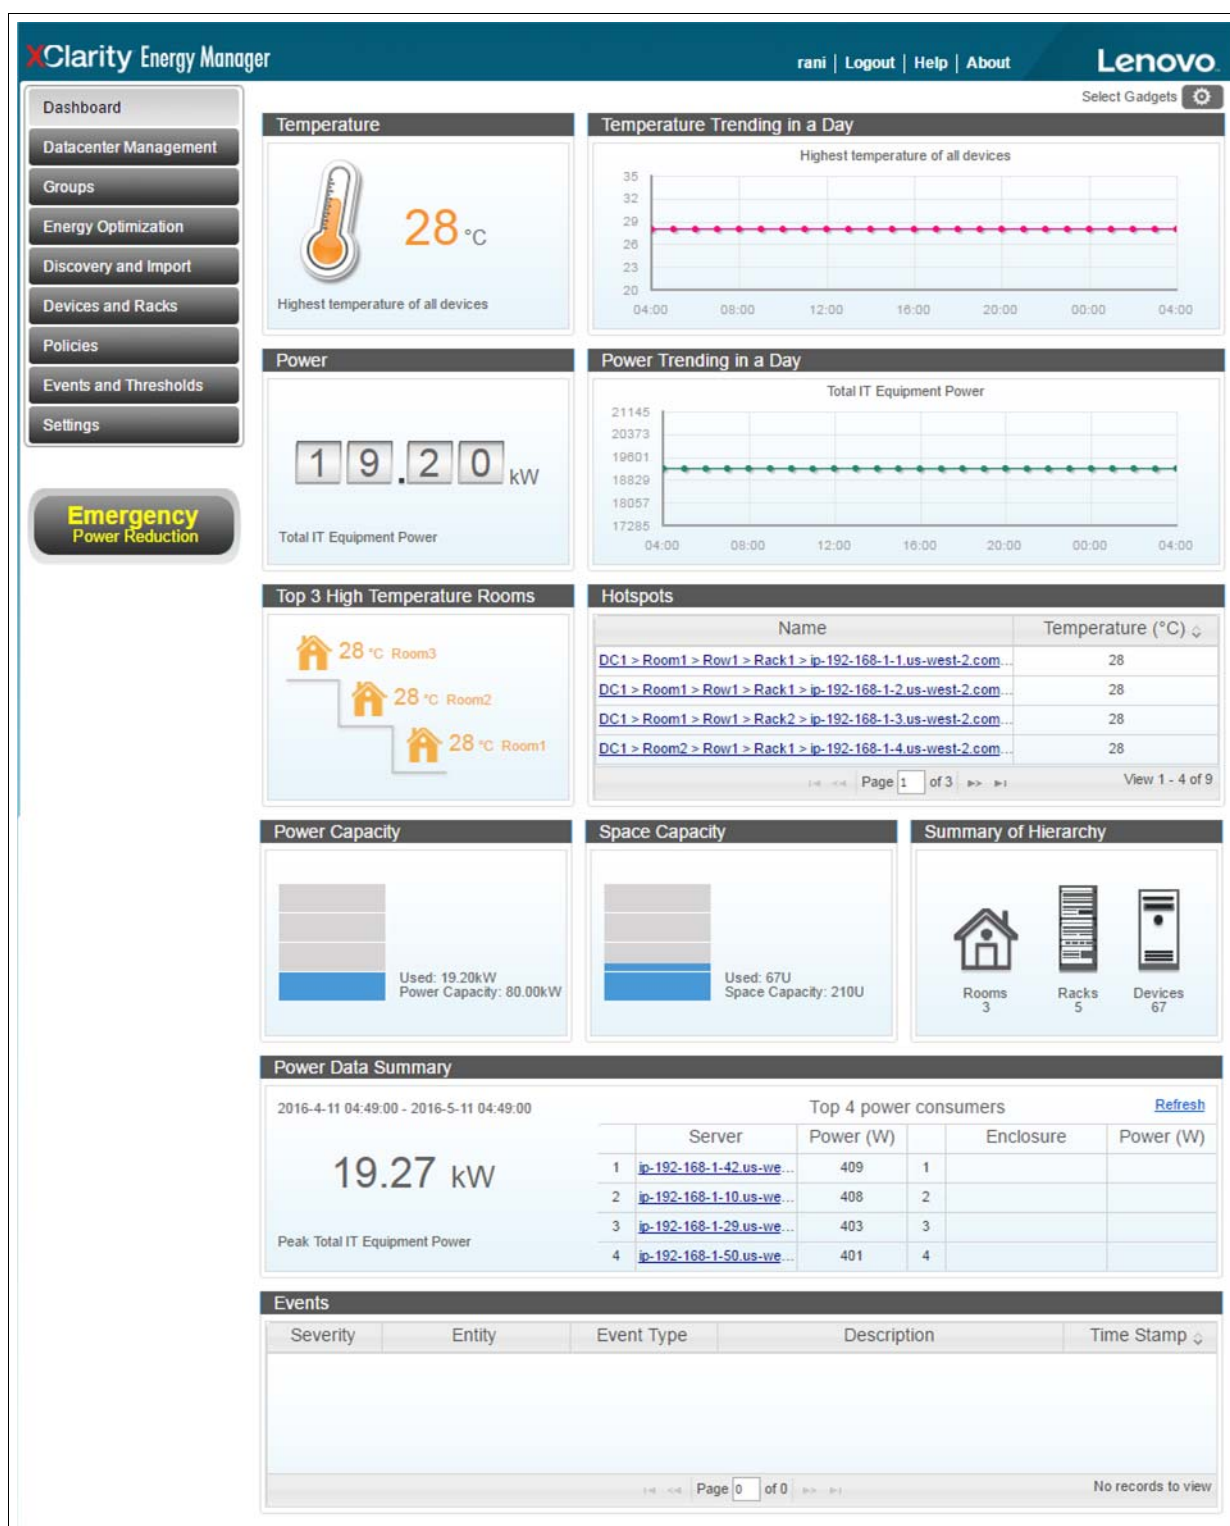


Figure 7-42 Energy Manager Dashboard

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 shown in Figure 7-43. You can view statistics as well as how many racks are in the room, and how many devices exist.

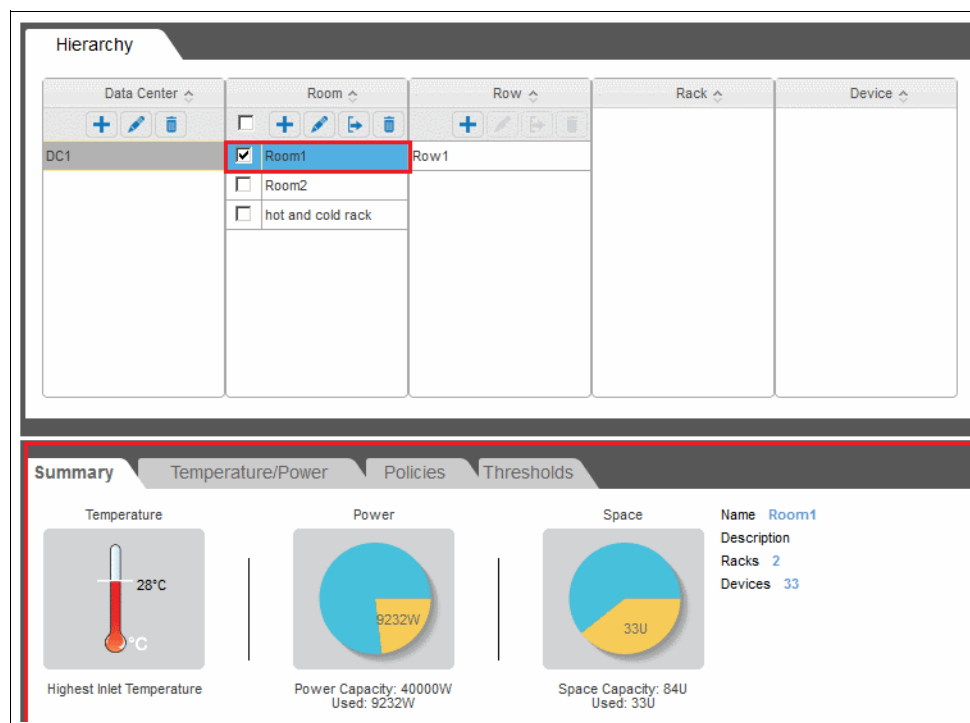


Figure 7-43 Hierarchy room selection displays room statistics

Likewise, when a rack is selected, information for that particular rack is displayed. 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 256 and “Setting thresholds” on page 257.

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 shown in Figure 7-44 on page 255.

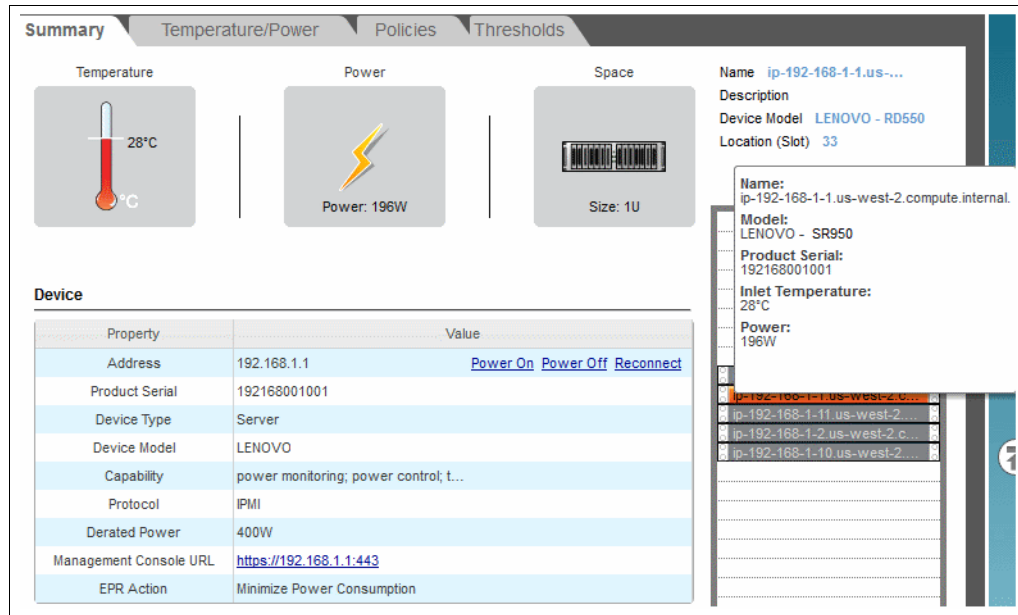


Figure 7-44 Device selected in rack

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. Then select the Data center, Room, Row, Rack and then the Device as shown in Figure 7-45.

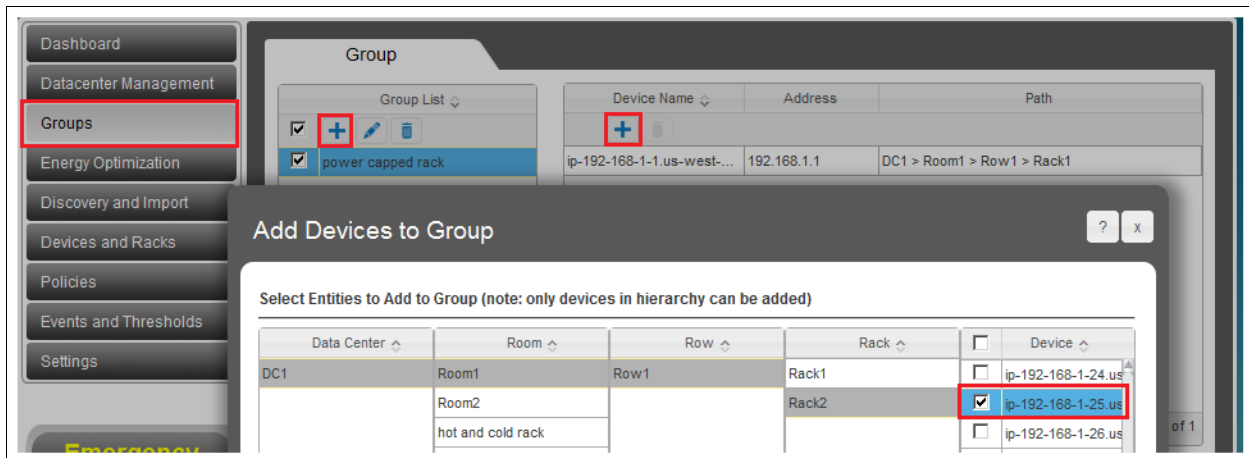


Figure 7-45 Creating a group

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:

1. Go to the Settings page and click **Add Subscriber**.
2. Fill in the email server configuration.
3. If you are only after threshold based events, check the **Subscribe threshold-based events only**.
4. 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:

1. Click the **Datacenter Management** page.
2. Select a device from the hierarchy by selecting the Data center, Room, Row, Rack and then the Device.
3. Click the **Policies** tab.
4. Click the **Add** button.
5. In the popup dialog, Figure 7-46, 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).
6. Set the schedule for the policy as either a permanent policy, for a specific time or as recurrent.

The screenshot shows the 'Add Policy' dialog box. On the left is a sidebar with navigation tabs: Dashboard, Datacenter Management, Groups, Energy Optimization, Discovery and Inventory, Devices and Rack, and Policies. The 'Policies' tab is selected. The main 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 Type:** Custom Power Limit (selected from a dropdown)
- Power Cap (W):** 230
- Policy Enabled:** ☒

Figure 7-46 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:

1. Click the **Datacenter Management** page.
2. Select a device from the hierarchy by selecting the Data center, Room, Row, Rack and then the Device.
3. Click the **Thresholds** tab.
4. Click the **Edit** option to set the threshold. Figure 7-47 displays the Thresholds page.

The screenshot shows the 'Thresholds' tab for a device with IP address 'ip-192-168-1-1.us-west-2.compute.internal'. It contains two sections: 'Power Thresholds' and 'Temperature Thresholds'.

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	Edit Clear
Lowest Inlet Temperature	Less	Edit Clear
Average Inlet Temperature	Greater	Edit Clear
Average Inlet Temperature	Less	Edit Clear

Figure 7-47 Setting a threshold

7.4.6 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 7-48 shows the Energy Optimization page.

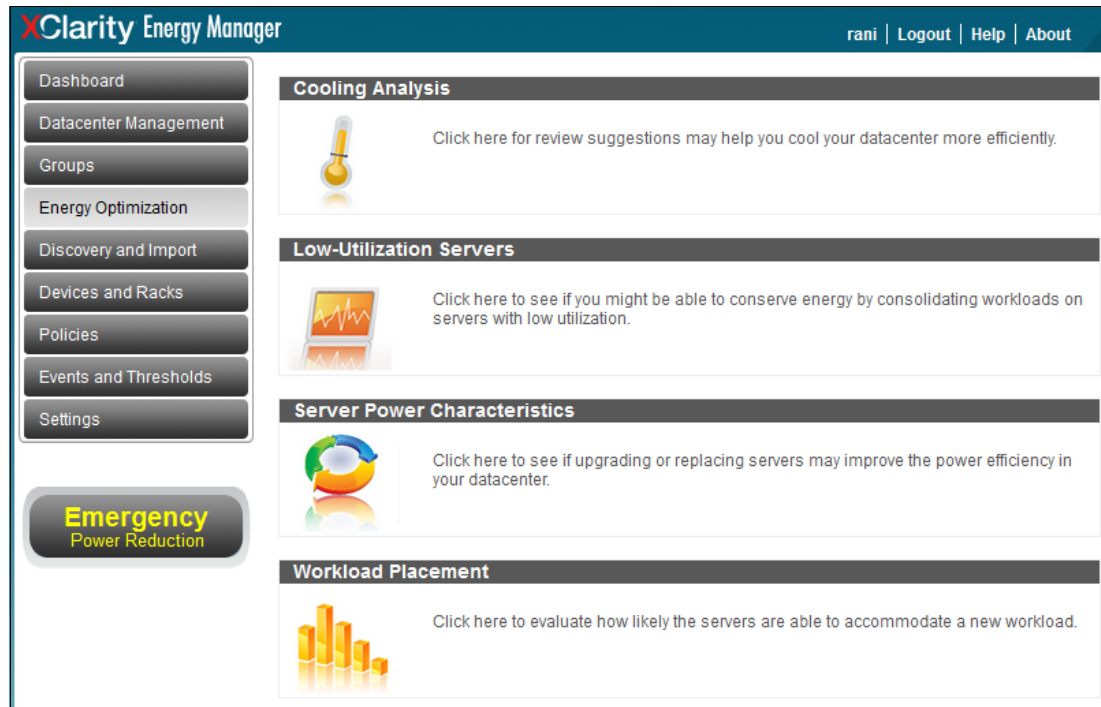


Figure 7-48 Energy Optimization page

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 shown in Figure 7-49.

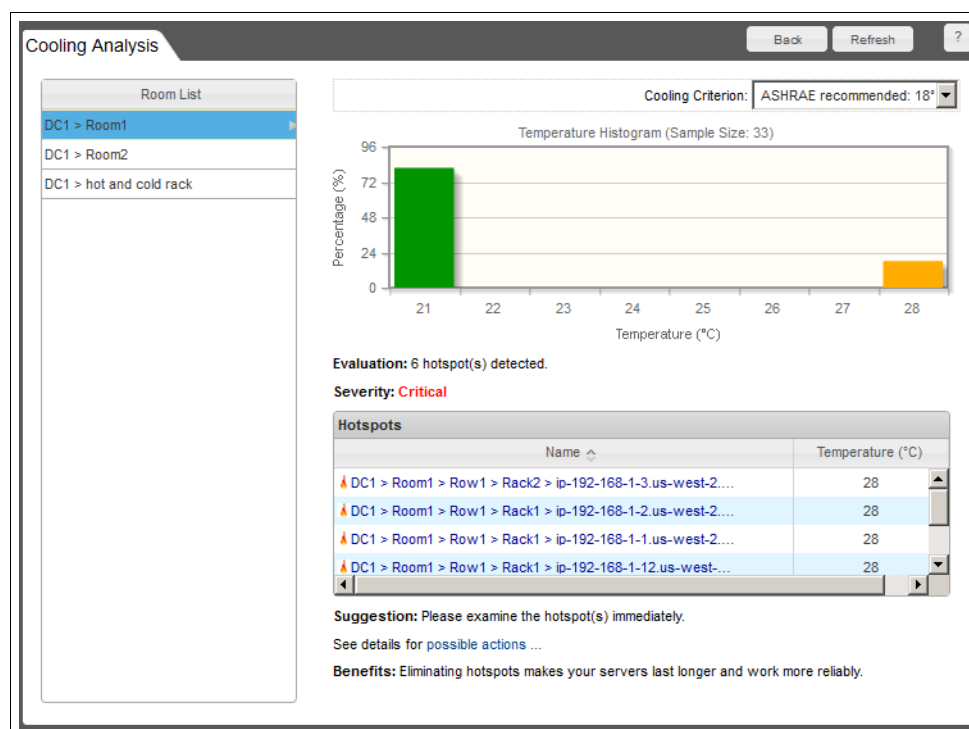


Figure 7-49 Cooling analysis page

Energy Manager will identify the servers that are causing hotspots in the data center and suggest an action to eliminate the hotspot, as shown in Figure 7-50. In this example, there are 6 servers with inlet temperatures that are higher than 27 °C.

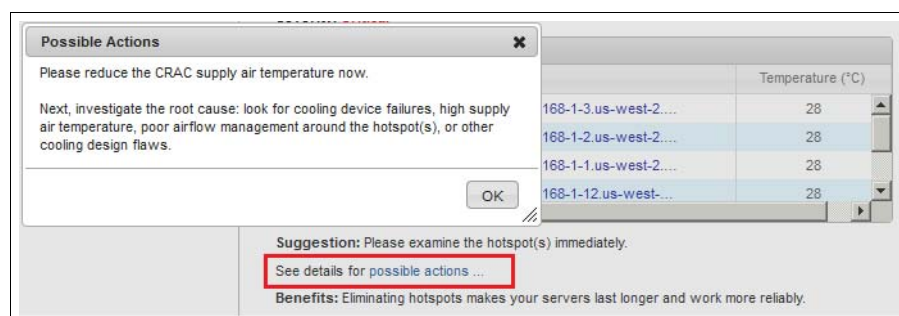


Figure 7-50 Cooling analysis page - possible action

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 as shown in Figure 7-51.



Figure 7-51 Low utilization servers detected

Server Power Characteristics

The Server Power Characteristics page represents the power usage of a server, as shown in Figure 7-52.




Figure 7-52 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 shown in Figure 7-53 which shows a break down of how peak power and idle power are distributed.



Figure 7-53 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 7-54.

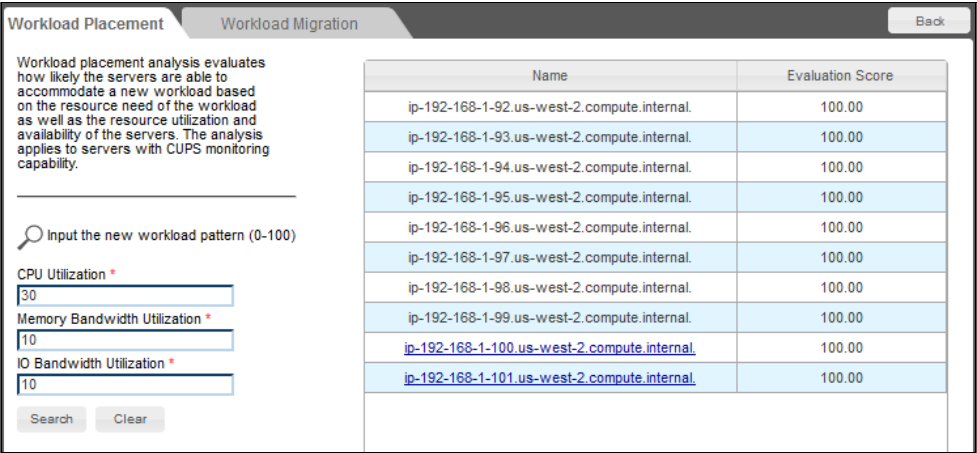


Figure 7-54 Workload analysis

7.5 Lenovo XClarity Controller management functions

Lenovo XClarity Controller (XCC) is the management processor for the server. XCC consolidates the service processor functionality, super I/O, video controller, and remote presence capabilities into a single chip on the server system board.

There are two ways to access XCC:

- ▶ Web-based interface. To access the web-based interface, point your browser to the IP address for the management processor.
- ▶ Command-line interface. To access the CLI interface, use SSH or Telnet to log in to the management processor.

For additional information on how to access and set up the networking for XCC, refer to 5.2.1, “Accessing XCC” on page 142.

XCC is operational whenever power is applied to a server, even if the server is not turned on. From the XCC interface, you can perform the following management functions:

- ▶ Diagnosing and management: health summary and system event logs
- ▶ Remote management function to log in to the server itself.
- ▶ Power the server on and off.

This section discusses the following management functions with XCC:

- ▶ 7.5.1, “XCC functions to diagnose and manage the server” on page 263
- ▶ 7.5.2, “Remote control” on page 266
- ▶ 7.5.3, “Powering the server on and off with XCC” on page 269

7.5.1 XCC functions to diagnose and manage the server

XCC provides hardware monitoring of the server and its installed components. This section discusses the diagnostic and management features of XCC:

- ▶ “Health summary”
- ▶ “Hardware event logs and alerting” on page 264
- ▶ “Remote control” on page 266

Health summary

The first screen that you see after completing the login credentials is the home page, Figure 7-55 on page 264. This page includes a summary of the health of the hardware. The home page also provides system information, quick actions, remote console and power utilization information.

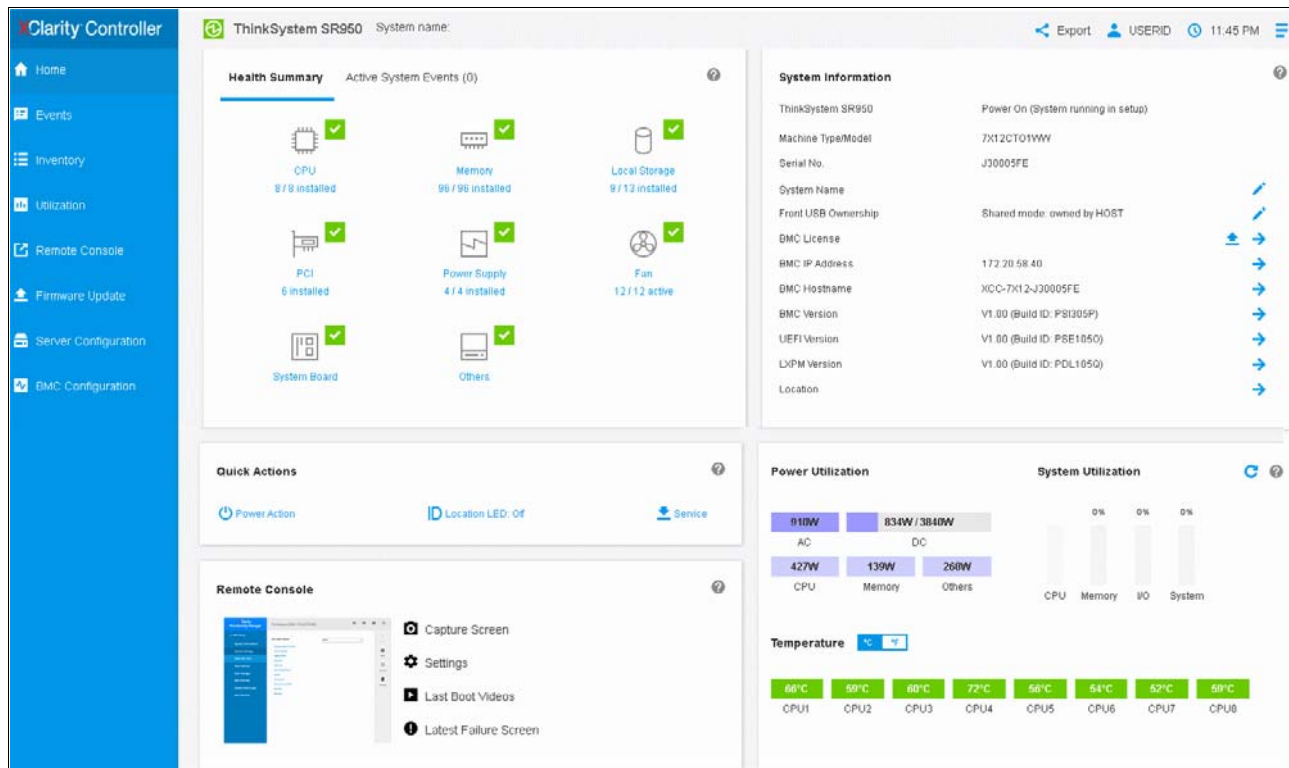


Figure 7-55 XCC home page

A green tick beside each hardware icon indicates that there are no problems from a hardware perspective.

XCC can check on the status of the following server components:

- ▶ CPUs
- ▶ Memory
- ▶ Local storage
- ▶ PCI adapters, including ML2 adapters and the LOM adapter
- ▶ Power supplies
- ▶ Fans
- ▶ System board
- ▶ Riser cards
- ▶ Compute boards

When a hardware error is detected in the server, either a yellow warning symbol or a red X replaces the green tick.

XCC does not check the functional status of most PCIe adapters regarding their hardware connections to external devices. You should refer to the system event log from within the operating system or the switch logs of the network and Fibre Channel switches to resolve any connectivity issues.

Hardware event logs and alerting

For information on the server 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 conditions that led up to a failure. The event log can be saved to a text file that can then be sent to Lenovo support.

The hardware event log is shown in Figure 7-56.

Severity	Source	Common ID	Message	Date
Information	System	FXSPSD00031	Hot Spare enabled for ThinkSystem SR950.	August 30, 2017 4:54:21 AM
Information	Power	FXSPPW20081	Host Power has been turned on.	August 23, 2017 11:23:47 PM
Information	System	FXSPNM40111	ENET[CM:ep1] DHCP-HSTName=XCC-7X12-J30005FE, DN=mgmt.bts.lab, IP@=172.20.50.40, SN=255.255.255.0, ...	August 23, 2017 11:23:43 PM
Information	System	FXSPNM40111	ENET[CM:ep1] DHCP-HSTName=XCC-7X12, DN=mgmt.bts.lab, IP@=172.20.50.40, SN=255.255.255.0, GW@=17...	August 23, 2017 11:16:34 PM
Information	System	FXSPNM40111	ENET[CM:ep1] DHCP-HSTName=XCC-7X12, DN=mgmt.bts.lab, IP@=172.20.50.40, SN=255.255.255.0, GW@=17...	August 11, 2017 9:07:05 AM
Information	System	FXSPNM40201	ENET[CM:ep2] IPv6-LinkLocalHstName=XCC-7X12, IP@=fe80:0a94:eff:fe40:79e5_Pref=64.	August 11, 2017 9:07:05 AM
Information	System	FXSPNM40141	LAN: Ethernet[MM: ep2] interface is now active.	August 11, 2017 9:07:05 AM
Information	System	FXSPNM40121	ENET[CM:ep2] IP: Cfg HstName=XCC-7X12, IP@=169.254.95.118, NetMask=255.255.0.0, GW@=0.0.0.0.	August 11, 2017 9:07:02 AM
Information	System	FXSPNM40281	ENET[CM:ep1] IPv6-LinkLocalHstName=XCC-7X12, IP@=fe80:0a94:eff:fe40:79e5_Pref=64.	August 11, 2017 9:07:02 AM
Information	System	FXSPNM40141	LAN: Ethernet[MM: ep1] interface is now active.	August 11, 2017 9:07:02 AM

Figure 7-56 XCC hardware event logs

Setting alerts for events

You can set alerts to be triggered for certain events that are recorded in the event log. This is configured under the **Events** → **Alert Recipients** in XCC. The alerts are sent via SMTP or SNMP v3.


The recipient of the alerts can receive events and event logs under three different categories: critical, attention and system. Figure 7-57 shows the Alert Receipts screen under the Hardware event log for creating an email alert.

Figure 7-57 Alert receipt screen for creating email alerts

Note: You will need to enable SNMPv3 traps in XCC via the **BMC Configuration** → **Network** → **SNMPv3** menu.

7.5.2 Remote control

Certain problems require that you interact with operating system or the UEFI setup panels or to the OS itself to detect or fix problems. For remotely managed servers, you can use the Remote Control feature of the XCC, which is accessed via the home page or left hand menu by selecting **Remote Console** → **Remote Control**.

Figure 7-58 shows the remote control launch screen. To launch the remote console session, hover over the thumbnail graphic under the Remote Console heading; the arrow icon appears . Click the icon to launch the remote session.

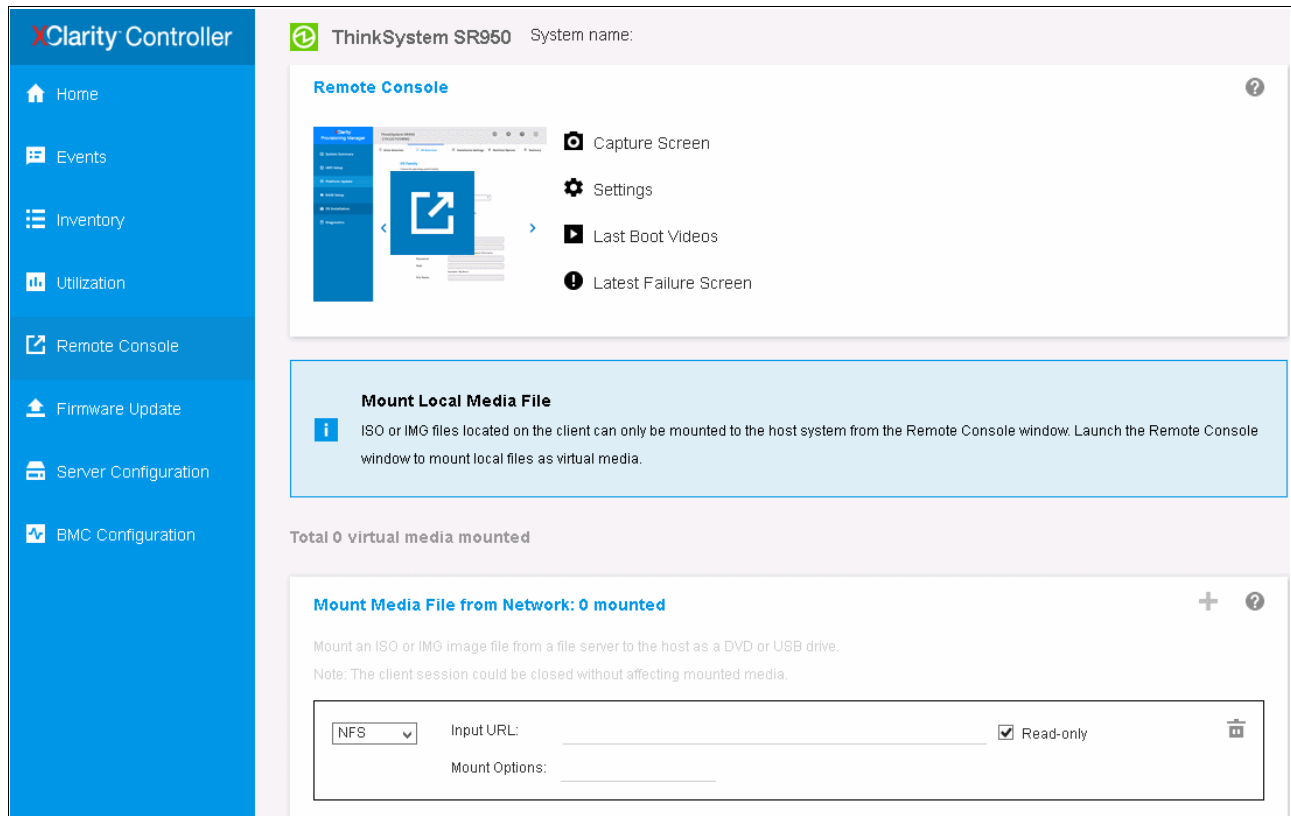


Figure 7-58 Remote console screen

Figure 7-59 on page 267 displays the remote console options when launching a session.

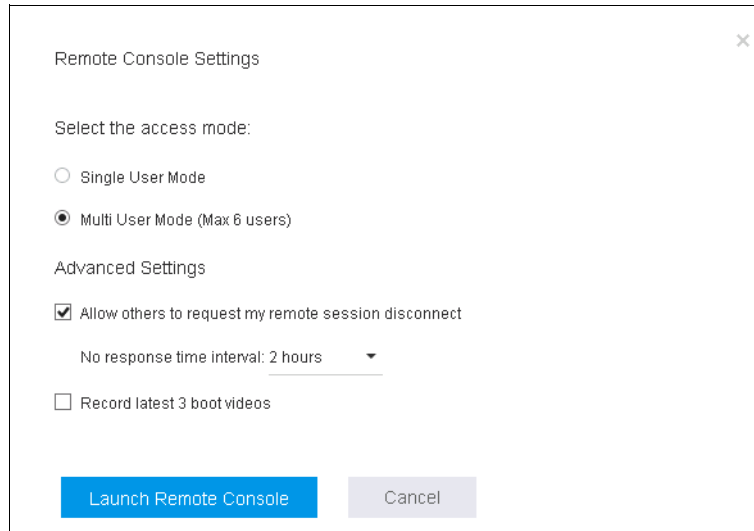


Figure 7-59 Remote console options

XCC 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 (CD/DVD-ROMs, ISO images, and USB devices) 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 XCC, you can power down, restart, or power on the server while maintaining the same remote control session.

Figure 7-60 on page 268 shows a remote session where the server has been booted into XClarity Provisioning Manager.

Popup blockers: Ensure your browser does not have a pop-up blocker enabled. If pop-ups are not allowed, the remote console session will not open.

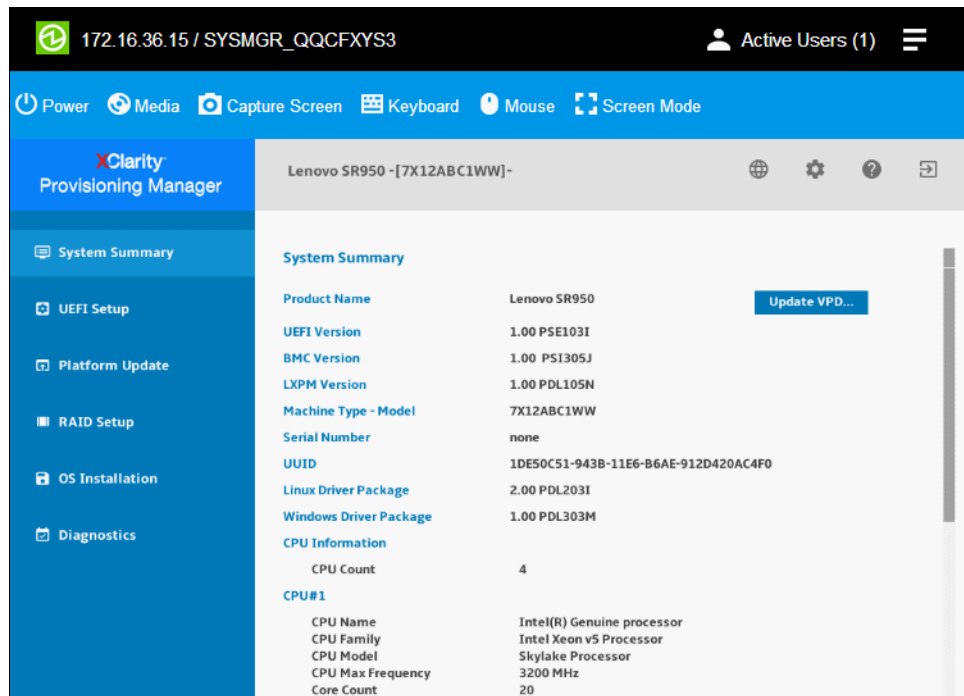


Figure 7-60 Remote console session

The remote console session also allows you to create your own customized keyboard key sequences by using the soft key programmer as shown in Figure 7-61.

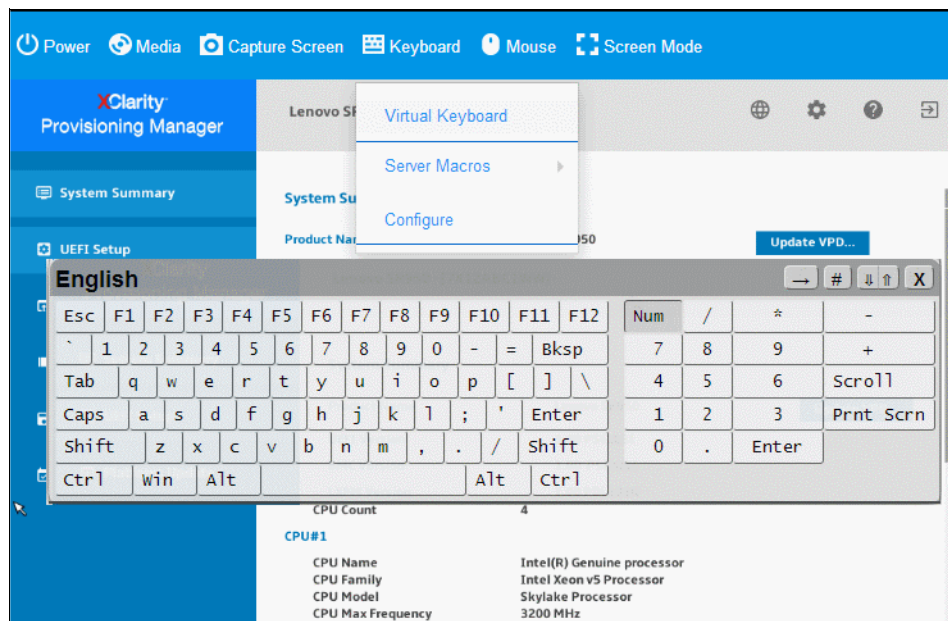


Figure 7-61 Virtual customizable keyboard

7.5.3 Powering the server on and off with XCC

Using the XCC interfaces, you can power the remote server on and off and back on again as long as the system is connected to a power source.

To power the system on and off from the XCC, follow these steps:

1. Log into XCC as explained in 5.2.1, “Accessing XCC” on page 142.
2. From the home page, scroll to the bottom of the page and click **Power Actions**. The available actions include:
 - Power on server
 - Power off server normally
 - Power off server immediately
 - Restart server normally
 - Restart server immediately
 - Boot server to System Setup
 - Trigger NMI
 - Schedule power actions
 - Restart management controller
 - AC power cycle server

These actions are shown in Figure 7-62.

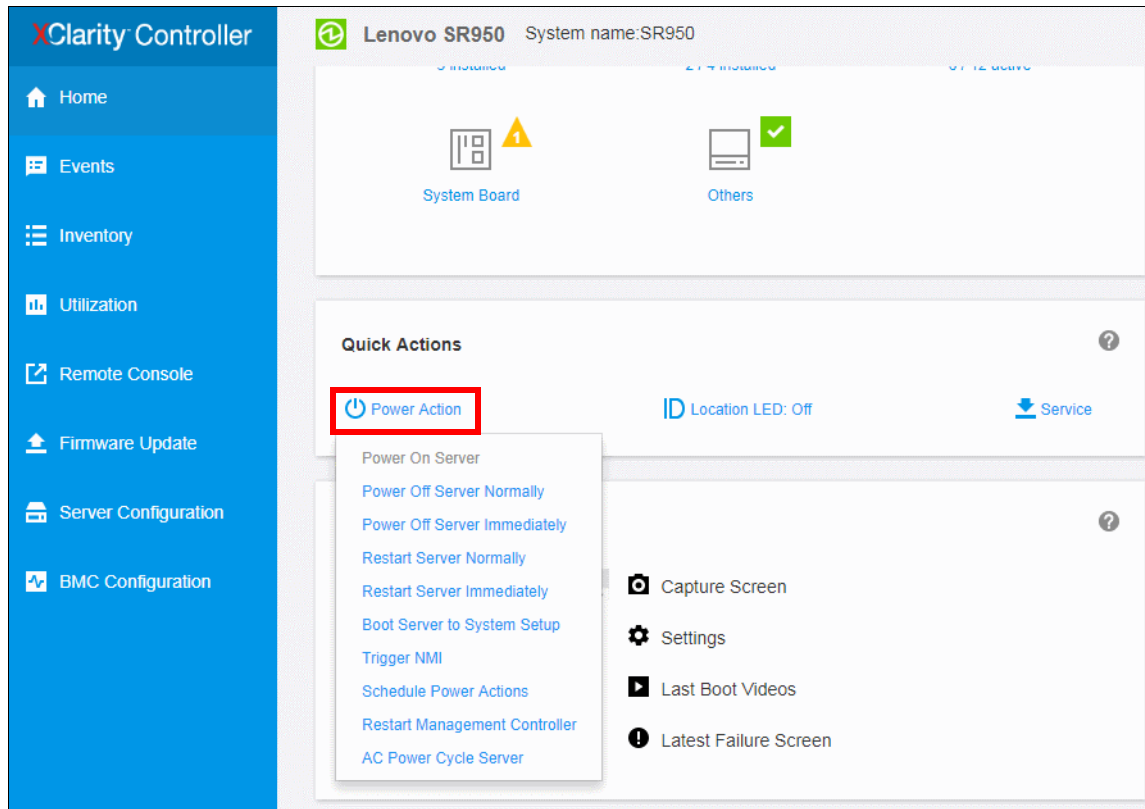


Figure 7-62 Power actions

3. If you wish to boot to a certain device or setup such as CD/DVD or F1 setup, use the **Server Configuration** → **Boot Options** page, as shown in Figure 7-63 on page 270.




Figure 7-63 Boot options

4. Select the device that the server will boot in to.
If you do not want to boot into a device then select the **No one-time boot** option and reboot the server from the **Home** page as explained in step 2.
5. When a boot option has been selected, click a radio button to select how and when the server will be rebooted. Click **Apply**.

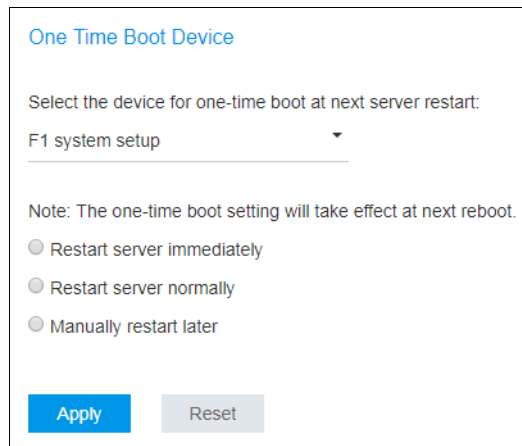


Figure 7-64 One-time boot options

For additional information on the XClarity Controller, see Systems Management Information Center at the following link:

http://sysmgmt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.systems.management.xcc.doc%2Fproduct_page.html

Additional XCC information in this book:

- For initial setup of the SR950 with XCC refer to 5.2, “Preparing the hardware with XClarity Controller” on page 142.
- For operating system installation with XCC refer to 6.2.1, “Installing an OS with XClarity Controller (XCC)” on page 185.

- For power and cooling monitoring with XCC, refer to 4.8.1, “XClarity Controller power and cooling” on page 109.

7.6 Lenovo XClarity Essentials

Lenovo XClarity Essentials (LXCE) is a collection of server management utilities that provides a less complicated method to enable management of the SR950 server more efficiently and cost-effectively.

XClarity Essentials includes the following utilities:

- OneCLI is a collection of several command line applications, which can be used to:
 - Configure the server.
 - Collect service data for the server. If you run Lenovo XClarity Essentials OneCLI from the server operating system (in-band), you can collect operating system logs as well. You can also choose to view the service data that has been collected or to send the service data to Lenovo Support.
 - Update firmware and device drivers for the server. Lenovo XClarity Essentials OneCLI can help to download UpdateXpress System Packs (UXSPs) for your server and update all the firmware and device drivers.
 - Perform miscellaneous functions, such as rebooting the server or reboot.

For more information, see 5.3.3, “XClarity Essentials OneCLI” on page 171.

- Bootable Media Creator (BoMC) is a software application that applies UpdateXpress System Packs and individual updates to your system.

Using Lenovo XClarity Essentials Bootable Media Creator, you can:

- Update the server using an ISO image or CD.
- Update the server using a USB key.
- Update the server using the Preboot Execution Environment (PXE) interface.
- Update the server in unattended mode.
- Update the server in Serial Over LAN (SOL) mode.

For more information, see 5.3.2, “XClarity Essentials Bootable Media Creator (BoMC)” on page 165.

- UpdateXpress is a software application that applies UpdateXpress System Packs and individual updates to your system.

Using Lenovo XClarity Essentials UpdateXpress, you can:

- Update the local server.
- Update a remote server.
- Create a repository of updates.

For more information, see 5.3.1, “XClarity Essentials UpdateXpress” on page 159.

7.7 LCD system information display panel

If you are physically in front of the server, it is easy to track hardware problems by using the LCD system information display panel on the front of the server. You can scroll through the display panel by using the up and down arrows that as shown in Figure 7-65.



Figure 7-65 LCD system information display panel

The LCD display panel allows quick access to system status, firmware, network, and health information. The menu structure of the display panel is shown in Figure 7-66.

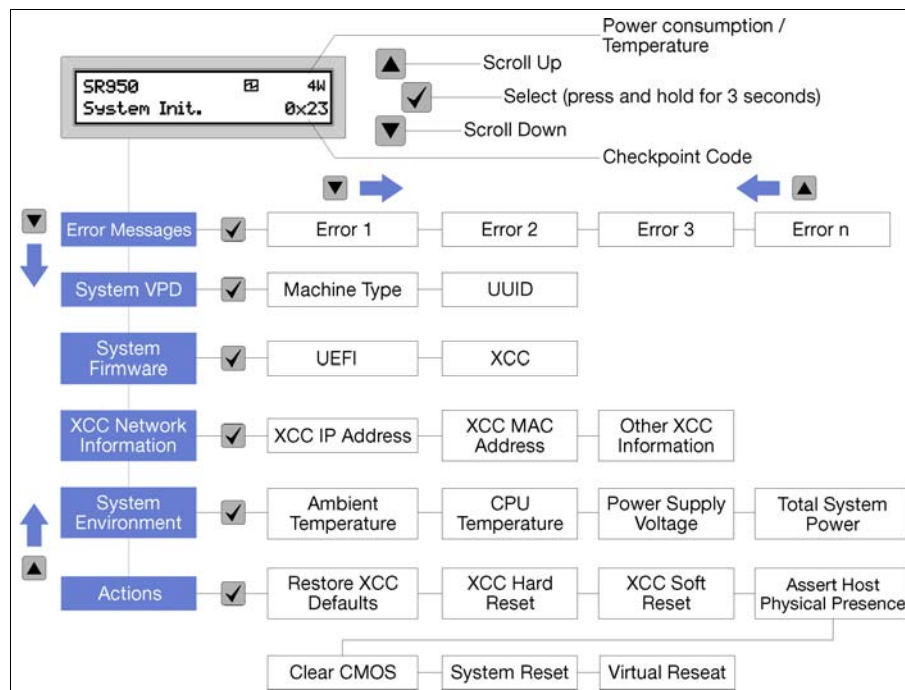


Figure 7-66 LCD display panel menu structure

If you are not physically near the server, you can see all events from a remote location with the XCC, including the XCC event logs. To access the event logs in XCC, refer to “Hardware event logs and alerting” on page 264.

Related publications

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

Lenovo Press publications

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

- ▶ SR950 Product Guide
<https://lenovopress.com/lp0647-thinksystem-sr950-server>
- ▶ SR950 Server Video Walkthrough
<https://lenovopress.com/lp0702-thinksystem-sr950-server-video-walkthrough>
- ▶ 3D Tour of the SR950
<https://lenovopress.com/lp0674-3d-tour-thinksystem-sr950>
- ▶ Memory Population Requirements for Lenovo ThinkSystem Servers
<https://lenovopress.com/lp0774-memory-population-requirements-for-thinksystem>
- ▶ Optimizing Memory Performance of Lenovo ThinkSystem Servers
<https://lenovopress.com/lp0697-optimizing-memory-performance-thinksystem>
- ▶ Articles on the SR950
 - Five Highlights of the ThinkSystem SR950
<https://lenovopress.com/lp0734-five-highlights-of-the-thinksystem-sr950>
 - Choosing between Lenovo ThinkSystem SR850 and SR950
<https://lenovopress.com/lp0737-choosing-between-thinksystem-sr850-and-sr950>
 - Workloads for 4-Socket and 8-Socket Servers
<https://lenovopress.com/lp0748-workloads-for-4-socket-and-8-socket-servers>
 - Usability in the Design of the ThinkSystem SR950
<https://lenovopress.com/lp0741-usability-in-the-design-of-the-sr950>
 - The Value of Refreshing Your 4-Socket Servers with the ThinkSystem SR950
<https://lenovopress.com/lp0753-the-value-of-refreshing-your-4-socket-servers-with-the-thinksystem-sr950>
 - ThinkSystem SR950 Memory Decisions
<https://lenovopress.com/lp0763-thinksystem-sr950-memory-decisions>
 - ThinkSystem SR950 Server Configurations
<https://lenovopress.com/lp0775-thinksystem-sr950-server-configurations>
 - The Value of Refreshing Your 8-Socket Servers with the ThinkSystem SR950
<https://lenovopress.com/lp0779-the-value-of-refreshing-your-8-socket-servers-with-the-thinksystem-sr950>
 - RAS Features of the Lenovo ThinkSystem SR950 and SR850

- <https://lenovopress.com/lp0777-ras-features-of-the-lenovo-thinksystem-sr950-and-sr850>
- Lenovo ThinkSystem SR950 New Options and Features - December 2017
<https://lenovopress.com/lp0827-thinksystem-sr950-new-options-and-features>
- ThinkSystem SR950 Performance Leadership
<https://lenovopress.com/lp0834-thinksystem-sr950-performance-leadership>
- Lenovo Servers for Mission Critical Workloads
<https://lenovopress.com/lp0841-lenovo-servers-for-mission-critical-workloads>
- Accelerate Your 4- and 8-Socket Server Refresh Cycle
<https://lenovopress.com/lp0849-accelerate-your-4-and-8-socket-server-refresh-cycle>
- Microsoft and Lenovo ThinkSystem SR950 – A Perfect Match
<https://lenovopress.com/lp0843-microsoft-and-thinksystem-sr950-a-perfect-match>
- SAP Business Process Applications and Lenovo ThinkSystem SR950 – A Perfect Match
<https://lenovopress.com/lp0853-sap-business-process-applications-and-thinksystem-sr950-a-perfect-match>
- ThinkSystem SR950 New Options - March 2018
<https://lenovopress.com/lp0856-thinksystem-sr950-new-options-march-2018>
- SAP HANA and Lenovo ThinkSystem SR950 – A Perfect Match
<https://lenovopress.com/lp0876-sap-hana-lenovo-thinksystem-sr950-perfect-match>

Other publications and online resources

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

- ▶ Lenovo ThinkSystem SR950 product publications:
<http://thinksystem.lenovofiles.com/help/index.jsp>
 - Quick Start
 - Rack Installation Guide
 - Setup Guide
 - Hardware Maintenance Manual
 - Messages and Codes Reference
 - Memory Population Reference
- ▶ ThinkSystem SR950 drivers and support
<http://datacentersupport.lenovo.com/products/servers/thinksystem/sr950/7x11/downloads>
- ▶ Lenovo ThinkSystem SR950 product page:
<http://www3.lenovo.com/us/en/p/77XX7HSSR95>

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™	Flex System™	System x®
AnyBay™	Lenovo®	ThinkServer®
BladeCenter®	Lenovo XClarity™	ThinkSystem™
Bootable Media Creator™	Lenovo(logo)®	TruDDR4™
eX5™	ServerGuide™	UpdateXpress System Packs™

The following terms are trademarks of other companies:

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

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

ActiveX, Excel, Hyper-V, Internet Explorer, Microsoft, Microsoft Edge, 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.

Lenovo[™]

Lenovo ThinkSystem SR950 Planning and Implementation Guide



(0.5" spine)
0.475" <-> 0.873"
250 <-> 459 pages



Lenovo ThinkSystem SR950 Planning and Implementation Guide

**Covers the new
flagship
high-performance
ThinkSystem server**

**Provides technical
information about all
server features**

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

**Describes how to
manage the server
with various Lenovo
management tools**

Lenovo® ThinkSystem™ SR950 is a high-performance server that is designed for your most demanding, mission-critical workloads. Suitable workloads include in-memory databases, large transactional databases, batch and real-time analytics, ERP, CRM, and virtualized server applications.

The powerful 4U ThinkSystem SR950 can grow from two to eight Intel Xeon Scalable Family processors, and with 96 DIMM sockets, supporting up to 12 TB of high-speed memory. The modular design of the SR950 aids upgrades and servicing with easy front or rear access to all major subsystems to maximize server availability.

The SR950 packs numerous fault-tolerant and high-availability features into a high-density, 4U rack-optimized design that reduces the space needed to support massive network computing operations and simplify servicing.

This book provides a general understanding of SR950 technology, components and management features. It also describes preparing and implementing the server, including power and cooling considerations, operating system installations, systems management, and firmware update tools.

This book is aimed at customers, Lenovo Business Partners, and Lenovo employees who want to understand the features and capabilities of the ThinkSystem SR950 and want to learn how to install and configure the server 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.**