

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

# SAP Business Suite on Lenovo X6 Systems

---

**Provides an overview of SAP  
Business Suite components**

---

**Discusses solution architecture and  
sizing considerations**

---

**Describes Lenovo X6 rack and  
blade server offerings**

---

**Explores Lenovo storage and  
networking options for SAP solutions**

**Rani Doughty**

**Al Kalafian**

**Ilya Krutov**



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

**Last update on June 2016**

This edition applies to Lenovo X6 rack servers and Flex System X6 compute nodes.

**© Copyright Lenovo 2016. All rights reserved.**

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

# Contents

<b>Notices</b> .....	iv
Trademarks .....	v
<b>Preface</b> .....	vi
Authors .....	vi
<b>Chapter 1. Introduction to SAP</b> .....	1
1.1 Client/server concept .....	2
1.2 Introduction to SAP .....	2
1.3 SAP architecture .....	3
1.3.1 SAP environment architecture .....	4
1.3.2 SAP Business Suite .....	8
1.3.3 SAP S/4HANA .....	12
1.3.4 SAP analytics solutions .....	13
1.3.5 Mobile solutions .....	14
1.3.6 SAP database and technology .....	15
1.4 SAP and Lenovo System x: A key relationship .....	18
<b>Chapter 2. Sizing</b> .....	21
2.1 Sizing methodology .....	22
2.1.1 SAP Quick Sizer .....	22
2.1.2 Lenovo sizing and planning questionnaire for SAP solutions .....	23
2.1.3 SAP Application Performance Standard .....	24
2.2 Sizing the Lenovo solution .....	24
2.2.1 Target CPU utilization .....	24
2.2.2 Resource categories .....	25
2.2.3 Steps in the sizing process .....	25
2.3 Summary .....	26
<b>Chapter 3. Hardware and software considerations</b> .....	27
3.1 SAP Business Suite hardware implementation strategy .....	28
3.2 Virtualization .....	28
3.2.1 Advantages of virtualization .....	28
3.2.2 Virtualization technologies .....	29
3.2.3 Lenovo X6 for virtualized SAP systems .....	32
3.3 Solution architecture considerations .....	33
3.3.1 Scaling up or scaling out .....	33
3.3.2 High availability .....	34
3.3.3 When to use virtualization .....	36
3.4 Choosing the correct hardware .....	37
3.4.1 Servers supported by SAP .....	37
<b>Chapter 4. Lenovo X6 systems, storage, and networking offerings</b> .....	43
4.1 Introduction to Lenovo X6 .....	44
4.1.1 Fast application performance .....	44
4.1.2 Agile design characteristics .....	44
4.1.3 Resilient enterprise platforms .....	45
4.2 Lenovo X6 systems overview .....	45
4.3 Lenovo X6 rack server design and architecture .....	49

4.3.1	Modular components	49
4.3.2	Processor subsystem	52
4.3.3	Intel Xeon E7 features	54
4.3.4	Memory subsystem	60
4.3.5	Storage subsystem	63
4.3.6	Networking and I/O	64
4.3.7	Scalability	66
4.4	Lenovo Flex System X6 compute node architecture	66
4.4.1	Lenovo Flex System Enterprise Chassis	67
4.4.2	Lenovo Flex System X6 compute node architecture	67
4.4.3	Processor subsystem	68
4.4.4	Memory subsystem	69
4.4.5	Storage subsystem	69
4.4.6	I/O subsystem	69
4.4.7	Scalability	72
4.4.8	Lenovo Flex System X6 compute nodes comparison	74
4.5	Lenovo Flash and SSD internal storage	74
4.5.1	SAS/SATA SSD technology	75
4.5.2	NVMe SSD technology	75
4.5.3	Flash Storage Adapters	77
4.5.4	Lenovo X6 RAS features	78
4.5.5	Hardware fault tolerance	78
4.5.6	Intel Xeon processor RAS features	79
4.5.7	Machine Check Architecture recovery	80
4.5.8	Memory RAS features	81
4.5.9	Lenovo XClarity Administrator	83
4.5.10	Lenovo XClarity Energy Manager	83
4.5.11	Lenovo XClarity Integrators	84
4.5.12	Upward Integration Module for Zenoss	84
4.6	Lenovo Storage offerings	85
4.6.1	SAP customer requirements	85
4.6.2	Lenovo Entry SAN Storage	86
4.6.3	IBM Storwize for Lenovo	88
4.7	Lenovo Networking offerings	92
4.7.1	Lenovo RackSwitch G7028	92
4.7.2	Lenovo RackSwitch G7052	94
4.7.3	Lenovo RackSwitch G8052	95
4.7.4	Lenovo RackSwitch G8124E	97
4.7.5	Lenovo RackSwitch G8264	100
4.7.6	Lenovo RackSwitch G8264CS	102
4.7.7	Lenovo RackSwitch G8272	105
4.7.8	Lenovo RackSwitch G8296	107
4.7.9	Lenovo RackSwitch G8332	110
	<b>Abbreviations and acronyms</b>	<b>113</b>
	<b>Related publications</b>	<b>115</b>
	Lenovo Press publications	115
	Online resources	115

# 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:

BladeCenter®	RackSwitch™	TruDDR4™
eX5™	Lenovo(logo)®	VMready®
Flex System™	ServeRAID™	vNIC™
Lenovo®	ServerProven®	X-Architecture®
Omni Ports™	System x®	

The following terms are trademarks of other companies:

Basis, Intel, Intel Core, 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.

Hyper-V, Microsoft, SQL Server, Windows, Windows Server, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

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

# Preface

SAP is the maker of one of the world's most popular inter-enterprise software suites. It provides collaborative business solutions for all types of industries. The company's flagship offering is the SAP Business Suite, which includes solutions for enterprise resource planning (ERP), customer relationship management (CRM), product lifecycle management (PLM), supplier relationship management (SRM), and supply chain management (SCM).

This Lenovo® Press publication describes the solution architecture of running SAP on Lenovo X6 systems, including x3850 X6 and x3950 X6 rack servers and Flex System™ x280/x480/x880 X6 compute nodes. It introduces the offerings and explains the processes that are involved in correctly sizing an SAP platform. The paper also describes tools and resources that are available to help you make the best decisions, and reviews the hardware offerings from Lenovo storage and networking.

This paper is intended for SAP administrators and technical solution architects. It is also for Lenovo Business Partners and Lenovo employees who want to know more about SAP offerings and available Lenovo solutions for SAP customers.

## Authors

This paper was produced in collaboration with Lenovo Press by a team of subject matter experts from around the world.



**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 world-wide 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 systems management, pre- and post-sale tools, and infrastructure planning.



**Al Kalafian** is an SAP Infrastructure Architect and the Team Lead for Lenovo SAP Center of Competence in North America. Al started at IBM in 1981, and he has over 20 years of experience with SAP. He has deep technical experience with all SAP applications that are running on the Intel platform and SAP appliance offerings, such as SAP BWA and SAP HANA. Al has several certifications, including SAP Certified Basis Consultant, SAP Certified Installer, IBM Certified Technical Professional, System x hardware, Microsoft Windows, and Red Hat Linux.



**Ilya Krutov** is a Project Leader at Lenovo Press. He manages and produces pre-sale and post-sale technical publications on various IT topics, including x86 rack and blade servers, server operating systems, virtualization and cloud, networking, and storage. Ilya has more than 15 years of experience in the IT industry, backed by professional certifications from Cisco Systems, IBM, and Microsoft. He has written more than 200 books, papers, and other technical documents. He has a Specialist's honors degree in Computer Engineering from the Moscow Engineering and Physics Institute.

Thanks to the authors of the previous editions of this paper:

- ▶ Odilon Goulart

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

From Lenovo:

- ▶ Mark T. Chapman
- ▶ David Watts
- ▶ Tag Robertson



# Introduction to SAP

This chapter describes SAP as an organization, its core business offerings, and its areas of strength. It includes the following topics:

- ▶ 1.1, “Client/server concept” on page 2
- ▶ 1.2, “Introduction to SAP” on page 2
- ▶ 1.3, “SAP architecture” on page 3
- ▶ 1.4, “SAP and Lenovo System x: A key relationship” on page 18

**Tip:** All URLs starting with `service.sap.com` need an S-User ID for access. If you do not have this kind of user account, request one from the following web page:

<https://websmp105.sap-ag.de/~sapidp/002006825000100141192000E/>

SAP notes can be accessed by using the following URL format, where *number* is the SAP note number that you want to access:

`https://service.sap.com/sap/support/notes/number`

For example, to access SAP note #171356, use the following address:

<https://service.sap.com/sap/support/notes/171356>

## 1.1 Client/server concept

The concept of the *client/server model* has several uses but all are based around the use of a client system to access a centralized server to use a service. When a network application is described as being *client/server*, the application runs a service on a server and a client version of the application is used to access the service.

The term *client/server network* is one where a centralized server provides network services to other clients in the network. Such a network is described as a domain in Linux, Macintosh, UNIX, and Windows operating systems.

Sometimes a server is used to centralize resources that are shared by many users. For example, file servers centralize file storage, database servers centralize data storage, and web servers centralize the distribution of web content.

## 1.2 Introduction to SAP

SAP was started in 1972 by five former IBM employees in Mannheim, Germany. They started SAP AG to turn the real business processing vision into reality. SAP AG is one of the world's largest inter-enterprise software organizations and a leading independent software supplier. SAP stands for *Systeme, Anwendungen, Produkte in der Datenverarbeitung* which, translated to English, means *Systems, Applications, Products in Data Processing*.

The first software that was produced by SAP AG was called the R/1 system. The "R" stands for *real-time data processing*. At the end of the 1970s, R/1 was replaced by R/2, a mainframe-based software that was designed to handle many languages and currencies. This product was especially successful in multinational European companies.

In the 1990s, SAP R/3 had success by using the client/server concept, uniform appearance of graphical interfaces, consistent use of relational databases, and the ability to run on systems from multiple vendors. The client/server concept that was introduced at this time meant that the processing of an application was split between the server and the client workstation, where the server handled the centralized functions and the client workstation handled the user interface functions. The data management was separate from the server.

In 1992, SAP AG introduced R/3, a client/server enterprise resource planning (ERP) software that runs on multiple platforms, including UNIX and Microsoft Windows operating system. SAP helped customers structure their business operations by integrating business processes by using the ERP software offering. The ERP system allows the exchange of data and information between various business units, such as finance, human resources, and the logistics of an organization.

SAP R/3 is an integrated tool to replace an obsolete and inefficient IT infrastructure. Initially, companies had departmental IT infrastructures to meet their automation requirements. Each department had its own policies, interfaces, and related IT infrastructure. Over time and through experience, a consolidated approach to link interfaces was developed. SAP is an integrated approach to such interfaces.

SAP enables business process change through its process preferred practices. It drives the organization towards process orientation to achieve higher productivity standards and to reduce production and labor costs.

During the last decade, SAP AG was updated to include the Java stack. This feature enables many other features, including web and mobile systems integration.

SAP offers the following characteristics, which makes it a worldwide leading ERP application:

- ▶ Hardware platform independence
- ▶ Database and operational system abstraction
- ▶ Multiple languages, time zones, and currency at same time
- ▶ Integration and support of a wide range of business solutions
- ▶ Support of multiple and different clients
- ▶ Flexibility of customization to meet customer requirements
- ▶ Scalability to grow to unlimited users
- ▶ Resources to extend preferred practices to an entire value chain
- ▶ IT investment protection with programs, such as Safe Passage for customers that are running other ERP solutions

### 1.3 SAP architecture

SAP offers a complete range of enterprise software applications and business solutions to drive businesses that range from enterprise to small. It offers these business applications and solutions to handle every aspect of the business. The Windows and Linux operating system markets are important to SAP, and System x is an ideal platform on which to run SAP solutions.

The SAP NetWeaver solution map (as shown in Figure 1-1) highlights various IT processes that are supported by solutions from SAP and its partners.

User Productivity Enablement	Running an Enterprise Portal	Enabling User Collaboration	Business Task Management	Mobilizing Business Processes	Enterprise Knowledge Management	Enterprise Search
Data Unification	Master-Data Harmonization		Master-Data Consolidation	Central Master-Data Management		Enterprise Data Warehousing
Business Information Management	Enterprise Reporting, Query, and Analysis	Business Planning and Analytical Services	Enterprise Data Warehousing	Enterprise Knowledge Management	Enterprise Search	
Business Event Management	Business Activity Monitoring			Business Task Management		
End-to-End Process Integration	Enabling Application-to-Application Processes	Enabling Business-to-Business Processes	Business Process Management	Enabling Platform Interoperability	Business Task Management	
Custom Development	Developing, Configuring, and Adapting Applications			Enabling Platform Interoperability		
Unified Life-Cycle Management	Software Life-Cycle Management			SAP NetWeaver® Operations		
Application Governance and Security Management	Authentication and Single Sign-On			Integrated User and Access Management		
Consolidation	Enabling Platform Interoperability	SAP NetWeaver Operations	Master-Data Consolidation	Enterprise Knowledge Mgmt.	Enterprise Data Warehousing	
Enterprise SOA Design and Deployment	Enabling Enterprise Services					

Figure 1-1 SAP NetWeaver solution map

The SAP IT practices are identified by SAP as ways to map business requirements to specific IT projects as a means of bringing order to the infrastructure, as described in Table 1-1.

Table 1-1 IT practices with SAP NetWeaver

Activity	Benefit
User productivity enablement	Help users and groups improve productivity through enhanced collaboration, optimized knowledge management, and personalized access to critical applications and data.
Data unification	Consolidate, rationalize, synchronize, and manage master data for improved business processes.
Business information management	Increase the viability, reach, and usefulness of structured and unstructured enterprise data.
Business event management	Ensure that business events from multiple systems are distributed to the appropriate decision makers in the context of relevant business processes.
End-to-end process integration	Make disparate applications and systems work together consistently to perform business processes.
Custom development	Rapidly create enterprise-scale applications that drive the company's differential advantage.
Unified lifecycle management	Automate application management and processes to optimize an application's lifecycle.
Application governance and security management	Maintain an appropriate level of security and quality in intellectual property and information assets.
Consolidation	Deploy a consolidated technology platform with the ability to allocate computing power according to changing business needs.
Enterprise service-oriented architecture (SOA) design and deployment	Consolidate and standardize basic processes, and use investments to compose distinctive business processes.

### 1.3.1 SAP environment architecture

An SAP *landscape* is a group of two or more SAP systems. These SAP systems often include *development*, *quality and test*, and *production* systems. One SAP system consists of one or more SAP instances. An SAP *instance* is a group of processes that are started and stopped at the same time.

Each instance can run different processes with different purposes on the same SAP system. Some systems, such as Process Orchestration (formerly *Program Integration*), have Advanced Business Application Programming (ABAP) and Java stacks, meaning a process by which to run each stack. However, other SAP systems have only one stack (ABAP or Java). The stacks that your system has depends on the required features and system type.

The SAP environment includes the following types of instances:

- ▶ Primary application server instance

Each SAP system has at least one primary application server instance that has the following components, which are divided by stack (AS ABAP and AS JAVA):

- Usage type AS ABAP:
  - Work processes (dialog, batch, spool, and update)

- Dispatcher
  - Gateway
  - Internet Communication Manager (ICM)
  - Internet Graphics Service (IGS)
- Usage type AS Java:
- Java dispatcher
  - Server processes
  - Software Deployment Manager (SDM)
  - IGS
- ▶ Central services instance (SCS and ASCS)
- This instance is installed with the primary application server instance and can be on the same server or inside another server (for cluster purposes, for example). There are two SAP central services instances, one instance for the Java stack, called *Java central services instance* (SCS) and another instance for the ABAP stack, called *ABAP central services instance* (ASCS). These instances form the basis of communication, synchronization, and database update for Java or ABAP clusters. This instance consists of the message server and the enqueue server.
- ▶ Database (DB) instance
- This instance is a mandatory component of SAP system based on the SAP NetWeaver platform. It has the RDBMS software and database. This component can be installed on the same server instance of the primary instance and the central services instance, but you can also install it on a separated instance for workload distribution and high availability purposes.
- ▶ Additional application server instance (AAS), optional
- This optional instance is used to add more resources to an SAP system and must be installed in an extra server. When you install more instances on your SAP system, this instance requires more resources (processor and memory) from the database instance. Therefore, resize the database instance when new applications are added (as described in Chapter 2, “Sizing” on page 21).
- ▶ Enqueue Replication Server instance (ERS), optional
- This optional instance is used to replicate the enqueue table between the instances and to ensure the same database locks and updates in high availability systems. This instance must be installed on different servers (hardware) to remove a single point of failure (SPOF) for this component.

In an SAP environment, you can use a 2-tier (primary application server instance) or a 3-tier landscape (primary application server instance with another application server instances).

Most SAP customers implement a 2-tier client/server architecture. One tier is the presentation layer that provides an interface to the user. The other tier is the application and database layer that operates and hosts the business processes.

Bigger SAP implementations use 3-tier or multitier landscapes that allow easier implementation of high availability for each component or tier.

SAP flexibility enables you to decide the best tier approach; that is, a two-, three-, or multitier architecture, as shown in Figure 1-2 on page 6.

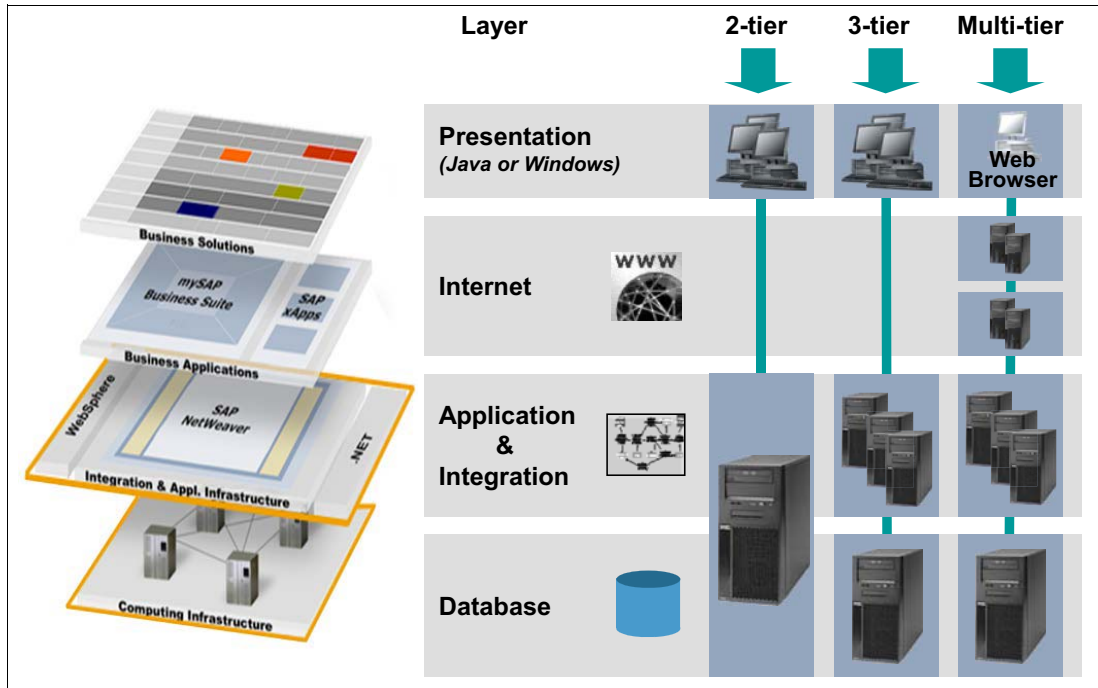


Figure 1-2 Various software layers of SAP ERP (© SAP AG 2007, 2010. All rights reserved.)

These layers feature the following classifications:

- ▶ **Presentation layer**

This layer provides an SAP graphical user interface, which is commonly referred to as *SAPGUI*. Because SAP offers an Internet-based interface as well, *SAPGUI* can achieve minimum bandwidth consumption by using the TCP/IP protocol.

- ▶ **Application layer**

The application layer processes user transactions and performs other functions (such as application interface, print server, and database access).

- ▶ **Internet layer**

In a multitier architecture, this layer provides a web interface for browser access to business applications.

- ▶ **Database layer**

The database layer is used for storing the SAP data that is generated by businesses and the SAP application programs, which are used by the database from the SAP application servers at run time. An organization has the flexibility to choose the database of its choice (including SQL Server, Oracle, DB2, SAPDB/MaxDB, Sybase, and HANA).

As shown in Figure 1-2, 2-tier architecture can have a presentation layer and a consolidated hardware server for application and database layers. A typical 3-tier architecture consists of presentation hardware, single or multiple application servers for load sharing and scalability, and often a single database or production server. The multitier architecture deploys a web-based presentation layer, an Internet server, multiple application servers, and a centralized database server.

An SAP implementation project is implemented in a phased manner. It requires different types of systems during different phases to run the project. It also often requires training users and understanding your business process, customizing the SAP application, testing the customized application, and porting it to the production server.

The following servers generally are required to run a project:

- ▶ Development system (DEV)

This system is required to customize the application to meet the business process requirements of the customer. Because the workload on such servers is minimal, a low-end hardware server meets the requirement.

- ▶ Test and Quality Assurance (QA) systems

SAP applications that are customized on the development server often are tested on this server before moving them to the production environment. A copy of the production database often is used to test the performance and functionality of SAP before moving this configuration or changing to a production environment. The testing and QA involves the upgrades that are related to the operating system, database software, hardware, firmware updates, SAP patches upgrades, and so on. Because the QA or test systems often are used for performance test, copy the database from production periodically to ensure that the test is in the same environment as the real production system. Because of this performance requirement, the QA or TST hardware must be as similar as possible to production solution in size and in architecture.

- ▶ Training system (TRN)

This system is used to provide training to users and is used for testing any specific requirements. Because the system is used only to provide training to users without any load factor, the hardware requirements often are not significant.

- ▶ Application server that is used for the production system (PRD)

This server handles all user activity inside an SAP system. To perform this task, all programs that are related to the SAP application are hosted on the SAP application servers. Depending on the workload and scalability factor, there can be multiple application servers or a single application server. Because the application server often handles all user connectivity and processes, its hardware requirement can vary based on the number of users, the number of batch jobs that are required, or both.

- ▶ Database server that is used for the PRD

This server handles the highest single load factor in a system environment because all queries and updates go to the database server. The hardware requirement is a high configuration, which can be scaled further. Because these servers are critical, these servers must be backed up regularly and have a high level of built-in redundancy. For a large enterprise, a scalable external storage appliance to store the data to meet ever growing database requirements is ideal. Because the server often is installed with the application server, you must add the application server if this configuration is used.

- ▶ SAP Solution Manager Product

This product provides facilities in technical support for distributed systems with functionality that covers all key aspects of solution deployment, operation, and continuous improvement. It is also an SAP requirement for SAP Support and Performance Analysis.

Organizations have differing requirements. Therefore, considering these requirements, they can plan the architecture that suits them best. SAP sizing tools, as described in Chapter 2, “Sizing” on page 21, can help you determine the architecture to implement.

An ideal scenario is to have a multitier infrastructure or at a minimum, a 3-tier architecture. However, combinations can be selected that consider various factors, such as future growth, workload factor, and budget. It is common to find a single server that is used for development and testing. Similarly, many environments use a common server for application and database, which makes it a 2-tier environment.

## 1.3.2 SAP Business Suite

The SAP Business Suite offers organizations a choice of integrating business units that use the complete business suite or various modules separately. SAP Business Suite consists of a set of core applications, industry applications, and supplementary applications that are built on the SAP NetWeaver technology platform.

This section describes the core applications of SAP Business Suite.

### SAP Enterprise Resource Planning Solution

The SAP ERP Solution addresses the core business software requirements of the most demanding organizations, in all industries and sectors. Built on R/3, the SAP ERP solution offers complete integration for the core business applications of any size organizations. It is an analytical tool that helps with business decisions. It also provides the capability to manage financial databases, assets, cost accounting, production operations, and corporate services of an organization.

The SAP ERP solution runs on several platforms, from Windows to Linux systems that run on the most-used databases of the market, or any possible combination between them. For more information about supported systems, see Chapter 3, “Hardware and software considerations” on page 27.

SAP ERP uses the client/server model and can store, retrieve, analyze, and process corporate data in many ways for financial analysis, production operation, human resource management, and most other business processes. SAP ERP makes it possible to access the ERP database and applications through Internet access and web browsers. A sales representative can start the workflow for a sales order by completing an electronic form on a notebook computer that is translated into input for the R/3 system. Other interfaces (such as Microsoft Office with Duet) can also be used. The web implementation adheres to the Workflow Client API standard of the Workflow Management Coalition (WfMC).

SAP ERP is delivered with selected standard processes enabled and many other optional processes and features disabled. At the heart of SAP ERP are thousands of tables that control the way that the processes are run. This feature allows an organization to upgrade to any of the full range of SAP solutions, which integrates with Customer Relationship Management (CRM), Product Lifecycle Management (PLM), and Supply Chain Management (SCM) applications.

SAP ERP offers the following solutions that support the key functional areas of any organization:

- ▶ **SAP ERP Financials**

This application provides a complete financial management solution for a broad range of industries. It integrates accounting, financial reporting, performance management, and corporate governance requirements of an organization.

- ▶ **SAP ERP Operations**

This solution helps organizations achieve operational excellence in key areas of procurement and logistics execution, product development and manufacturing, sales, and services.

- ▶ **SAP ERP Human Capital Management**

This solution delivers leading-edge human management capabilities to an organization that enables it to maximize workforce potential. It automates talent management, workforce process management, and payroll and workforce deployment, which increases the efficiency of the available human resources.



- ▶ **SAP ERP Corporate Services**

This solution allows organizations to manage space, enterprise assets, project portfolios, corporate travel, environment, health, and global trade services more effectively.

## **SAP Customer Relationship Management**

The SAP CRM application provides a platform for marketing, sales, and service professionals to obtain complete customer intelligence that they can use to effectively manage customers processes. It enables multichannel customer interactions, including mobile (tablets and smartphones), Internet, and social media, and offers a dedicated communications infrastructure that helps to connect all users anytime, anywhere.

## **SAP Product Lifecycle Management**

SAP PLM application helps companies manage, track, and control all product-related information over the complete product and asset lifecycle and throughout the extended supply chain. SAP PLM facilitates creativity and frees the process of product innovation from organizational constraints.

## **SAP Supply Chain Management**

SAP SCM enables adaptive supply chain networks by providing planning and execution capabilities to manage enterprise operations and visibility, collaboration, and radio frequency identification (RFID) technology to streamline and extend those operations beyond corporate boundaries. SAP SCM transforms traditional supply chains from linear, sequential processes into open, configurable, responsive supply networks across a globally distributed environment.

## **SAP Supplier Relationship Management**

SAP Supplier Relationship Management (SRM) integrates strategic practices for supplier qualification, negotiation, and contract management. It performs this task tightly and cost-effectively with other enterprise functions and their suppliers' processes through a single analytical framework and support for multichannel supplier enablement. SAP SRM accelerates and optimizes the entire end-to-end procure-to-pay process by supporting integrated processes and enforcing contract compliance, which results in realizable savings.

To complement the core applications, SAP AG developed specific software that is based on the requirements of each industry sector.

## **SAP Industry applications**

The SAP Industry applications address the requirements of specific business processes for many industries by complementing the basic business processes that are common to all large enterprises. The architecture and business functionality of the SAP Industry applications are a result of SAP in-depth understanding of industry-specific business requirements and the resulting business processes.

SAP Industry applications include the following most common applications:

- ▶ **SAP Aerospace and Defense:**
  - SAP Bulk Fuel Management
  - SAP Complex Assembly Manufacturing
- ▶ **SAP for Automotive Solutions: SAP Dealer Business Management**
- ▶ **SAP for High Tech: SAP Solution Sales Configuration**

- ▶ SAP for Consumer Products:
  - SAP Apparel and Footwear
  - SAP NetWeaver Master Data Management, Global data synchronization option
  - SAP Trade Promotion Optimization
  - SAP Accelerated Trade Promotion Planning
  - SAP Extended Warehouse Management for Fashion
  - SAP Allocation and Rescheduling Optimization
- ▶ SAP for Oil and Gas: SAP for Oil and Gas Secondary Distribution
- ▶ SAP for Utilities:
  - SAP Utility Customer E-Services
  - SAP Pricing and Costing for Utilities
  - SAP Asset Lifecycle Accounting
  - SAP Business Process Tracking for Utilities
  - SAP Customer Energy Management
  - SAP waste and recycling solutions by PROLOGA
  - SAP Smart Meter Analytics
  - SAP Energy Portfolio Management
- ▶ SAP for Retail:
  - SAP Customer Analytics for Retail
  - SAP Demand Management offerings
  - SAP Forecasting and Replenishment for Retail and Wholesale Distribution
  - SAP Promotion Management for Retail
  - SAP Point-of-Sale Base
  - SAP Point-of-Sale
  - SAP POS Data Management
  - SAP Workforce Management
  - SAP Fraud Watch
- ▶ SAP for Defense and Security:
  - SAP Military Data Exchange
  - SAP Defense MILSTRIP Purchasing
  - SAP Defense Command and Control by Systematic
- ▶ SAP for Public Sector:
  - SAP Document Builder
  - SAP Taxpayer online services
  - SAP Public Budget Formulation
  - SAP Intelligence Analysis for Public Sector by Palantir
- ▶ SAP for Banking:
  - SAP Bank Analyzer
  - SAP Capital Yield Tax Management
  - SAP Deposits Management
  - SAP Funding Management
  - SAP Price Optimization for Banking
  - SAP Payment Engine
  - SAP Connector to eBAgent
  - SAP Process Objects Builder
  - SAP Tax Classification and Reporting for FATCA
  - SAP Liquidity Risk Management

- ▶ SAP for Insurance:
  - SAP Policy Management
  - SAP Reinsurance Management
  - SAP Product and Quotation Management for Insurance
- ▶ SAP for Engineering, Construction, and Operations: SAP Common Area Maintenance Expense Recovery
- ▶ SAP for Media: SAP Classified Advertising Management
- ▶ SAP for Professional Services:
  - SAP Multiresource Scheduling
  - SAP Commercial Project Management
- ▶ SAP for Telecommunications:
  - SAP for Telecommunications
  - SAP Convergent Mediation by DigitalRoute
- ▶ SAP for Healthcare: E-Care Social Connect

### **SAP Supplementary Applications and Service Extensions**

The Supplementary Applications and Service Extensions support specialized business processes that are common to almost all industries. They deliver a short time to value, appeal to specialized business users, and offer a high degree of process flexibility. These applications can be plugged in to your SAP system as one add-on or can be installed on a new infrastructure.

There are more than 100 Supplementary Applications and Service Extensions. They can be divided by core applications (such as ERP and SCM) that are related to an industry segment (such as manufacturing) or they can work together with any industry or application, such as Duet Enterprise software (Microsoft Office integration) and SAP Business One (SAP for Small Business).

The following sections describe the most common Supplementary Applications and Service Extensions.

#### ***Duet Enterprise***

Duet software enables seamless access to SAP business processes and data through Microsoft Office that interacts with enterprise applications. Duet is the result of a collaboration between SAP and Microsoft and is the first joint product that was created and supported by these two companies.

#### ***SAP Business One***

To help the new businesses stay profitable and remain competitive, SAP offers the Business One solution. This solution helps emerging and dynamic businesses manage their operations efficiently and cost-effectively while overcoming the unique challenges they face. It is an affordable business tool that was developed specifically for small and mid-sized businesses.

Consider the following points about this solution:

- ▶ For companies or subsidiaries of 10 to several hundred employees
- ▶ Simple to use and learn
- ▶ Powerful report generators
- ▶ Easy data navigation
- ▶ Implementation in a few days or weeks
- ▶ Offered for Windows and Linux
- ▶ Adaptable by implementation partners to customer or industry-specific requirements

- ▶ Perfect solution for branches or subsidiaries of SAP accounts
- ▶ Data can be exchanged with SAP by using SAP NetWeaver Process Integration

### ***SAP Business All-In-One partner solutions***

SAP Business All-In-One solutions are designed for the small and midsize enterprises. These solutions are pre-configured to meet the requirements of this industry segment and provide the following advantages:

- ▶ Targeted for industries with 100 - 2500 employees
- ▶ Designed for ease of use
- ▶ Offer the power of SAP world class business solutions
- ▶ Pre-configured and can be deployed in a short time
- ▶ Provide an easy to grow path and can be implemented in a phased approach
- ▶ Available on all platforms that run SAP NetWeaver
- ▶ Easily integrated with non-SAP applications

In collaboration with SAP, Lenovo offers small and mid-sized businesses a cost-effective way to get ERP applications up and running quickly.

## **1.3.3 SAP S/4HANA**

In 2015 SAP released SAP Business Suite 4 SAP HANA (or S/4HANA) which is a Business Suite built natively on the in-memory database SAP HANA. This fourth generation of a business suite software offering from SAP is only available with SAP HANA database underneath. The tight integration of the application components with the database allows it to integrate all mission-critical processes of an enterprise within one system and provides instant insight into a business.

SAP uses SAP Fiori as the user interface platform to ensure a simple and consistent experience for SAP users across products and end-user devices including tablets and mobile phones. SAP Fiori UX is designed to replace SAP UI and supports HTML5 and JavaScript.

SAP S/4HANA can be installed on-premises or consumed as a cloud-based offering. As of May 2016, the following components are available with the SAP S/4HANA on-premises edition:

- ▶ SAP S/4HANA Enterprise Management
- ▶ SAP S/4HANA LoB Products (to enhance core functions of SAP S/4HANA Enterprise Management for specific lines of business)
- ▶ SAP S/4HANA LoB Products for specific industries (to enhance SAP S/4HANA Enterprise Management to provide industry specific benefits for certain lines of business)
- ▶ SAP S/4HANA Compatibility Packs

In addition, SAP S/4HANA Finance is available, which is an on-premises edition consisting of SAP Accounting powered by SAP HANA. This edition was part of the initial release of SAP S/4HANA in March 2015.

The following SAP S/4HANA cloud offerings are available as of May 2016:

- ▶ SAP S/4HANA Marketing Cloud (for the marketing line of business)
- ▶ SAP S/4HANA Marketing Cloud, base option (core version of SAP S/4HANA Marketing Cloud)
- ▶ SAP S/4HANA Professional Services Cloud (for the professional services industry)
- ▶ SAP S/4HANA Enterprise Management Cloud (for a full ERP business scope. This is the equivalent to the on-premises Enterprise Management offering).

The on-premises edition follows a yearly release cycle while the cloud edition follows a quarterly release cycle.

For more information about SAP S/4HANA on Lenovo X6 systems, refer to *In-memory Computing with SAP HANA on Lenovo X6 Systems*, SG24-8086.

### 1.3.4 SAP analytics solutions

SAP analytics solutions can help IT managers and CIOs to make better-informed decisions more easily with intelligent data. Analytics software is built for industries and line-of-business processes and are enriched with content. These analytics solutions are divided in the following major groups:

- ▶ Business intelligence
- ▶ Enterprise performance management
- ▶ Governance, risk, and compliance
- ▶ Applied analytics

#### **Business intelligence**

*Business intelligence* is used for business reports, dashboards, and predictive analysis and includes the following basic software:

- ▶ SAP BusinessObjects BI platform
- ▶ SAP BusinessObjects Mobile
- ▶ BI rapid-configuration package for SAP HANA
- ▶ SAP Crystal Reports
- ▶ SAP BusinessObjects Dashboards
- ▶ SAP BusinessObjects Explorer
- ▶ SAP BusinessObjects Analysis, edition for Microsoft Office
- ▶ SAP Business Planning and Consolidation for Public Sector
- ▶ SAP Planning for Public Sector
- ▶ SAP BusinessObjects Design Studio
- ▶ SAP Predictive Analysis
- ▶ SAP Visual Intelligence

#### **Enterprise performance management**

The main objectives of enterprise performance management (EPM) are business planning, consolidation, and performance. These solutions can be integrated with Microsoft Office (Duet) for reports.

EPM includes the following standard packages:

- ▶ SAP Financial Consolidation
- ▶ SAP Intercompany
- ▶ SAP Profitability and Cost Management
- ▶ SAP Business Planning and Consolidation, version for SAP NetWeaver
- ▶ SAP Business Planning and Consolidation, version for the Microsoft platform
- ▶ SAP Strategy Management
- ▶ SAP Notes Management
- ▶ SAP Disclosure Management
- ▶ SAP Financial Information Management
- ▶ SAP ERP client for E-Bilanz
- ▶ SAP Spend Performance Management
- ▶ SAP Supply Chain Performance Management

## **Governance, risk, and compliance**

You can use the governance, risk, and compliance (GRC) solution to help manage access and process control and business risk management. It is helpful to have a system with access, process, and risk management on the same infrastructure.

The GRC solution includes the following standard packages:

- ▶ SAP Access Control
- ▶ SAP Process Control
- ▶ SAP Risk Management
- ▶ SAP Global Trade Services
- ▶ SAP Electronic Invoicing for Brazil
- ▶ SAP Sustainability Performance Management

## **Applied analytics**

Risk Reporting, Sales Analysis, and Net margin analysis are the major objectives of the Applied analysis solutions, which include the enterprise analytics solution for banking.

The following packages are included:

- ▶ SAP Enterprise Risk Reporting for Banking
- ▶ SAP Trade Promotion Effectiveness Analysis
- ▶ SAP Sales Analysis for Retail
- ▶ SAP Business Planning and Consolidation for Banking
- ▶ SAP Net Margin Analysis
- ▶ SAP BusinessObjects Business Intelligence, Edge edition, standard package
- ▶ SAP BusinessObjects Business Intelligence, Edge edition with data integration
- ▶ SAP BusinessObjects Business Intelligence, Edge edition with data management
- ▶ SAP Business Planning and Consolidation, Edge edition, version for the Microsoft

## **1.3.5 Mobile solutions**

SAP made a few acquisitions in the mobile solutions industry, such as Sybase and Syclo, to extend leadership in enterprise mobility. SAP now has a mobile product portfolio for new mobile asset management and field service solutions that runs on the SAP mobile platform. The SAP mobile portfolio was segregated into mobility management, mobile platform, and mobile apps.

### **SAP enterprise mobility management**

The SAP mobility management tools help companies to manage mobile devices, such as tablets and smartphones, and all the tasks that are associated with buying, securing, deploying, and maintaining these devices. It also helps to organize company documents between these devices.

SAP mobility management primarily uses the following software:

- ▶ SAP Afaria
- ▶ SAP Mobile Documents (desktop and server)

### **SAP mobile platform**

By using the SAP mobile platform, cost-effective and innovative apps can be designed that use an open, standards-based mobile application. This platform helps reduce the development time and cost during development phase.

The SAP mobile platform primarily uses the following software:

- ▶ Sybase Unwired Platform
- ▶ SAP Mobile Platform (Enterprise Edition)
- ▶ Agentry
- ▶ Sybase 365 Mobiliser Platform
- ▶ Sybase 365 Brand Mobiliser Platform

### **SAP mobile apps**

SAP has more than 200 mobile apps that are ready for deployment. For more information about SAP mobile apps, see this website:

<http://www.sap.com/pc/tech/mobile/software/solutions/mobile-services-overview.html>

## **1.3.6 SAP database and technology**

This section describes the software that is necessary to orchestrate business applications that close the gap between insight and action on-premises, on-demand, or mobile. SAP uses the SAP platform to deliver key functions to support business application orchestration, enhance team productivity, business processes, and application integration. SAP also features database solutions that are integrated with their technology platform.

### **SAP database**

The database category includes all SAP Database-related software, including the following in-memory database solutions and mobile databases:

- ▶ SAP HANA (Enterprise and Platform Edition)
- ▶ SAP Sybase Adaptive Server Enterprise
- ▶ SAP Sybase IQ
- ▶ SAP Sybase Advantage Database Server
- ▶ SAP Sybase SQL Anywhere
- ▶ SAP Operational Process Intelligence

### **SAP application development and integration**

The application development and integration category incorporates business functionality, which is made available as ready-to-use enterprise services and process components through its enterprise services repository. It also provides an integrated platform of composition technologies for business processes, composing applications, and deploying solutions.

SAP NetWeaver is an open technology platform that unifies technology components in a single platform, which reduces the need for custom integration and ensures that mission-critical business processes are reliable, secure, and scalable. SAP introduced NetWeaver to enable organizations to integrate their SAP and non-SAP solutions. Because NetWeaver supports present and future SAP and non-SAP solutions, the platform is the center of a growing system of applications and services.

SAP is evolving its solutions into service-oriented business applications that are based on NetWeaver, which allows SAP applications to integrate with non-SAP applications. For the first time, SAP is offering its platform for third-party development. Many independent software vendors already are building their applications on SAP NetWeaver.

Initially started as an integration and application platform, SAP NetWeaver evolved to become a composition platform that allows model-based development to enhance open enterprise services that are delivered by SAP Business Suite applications. In 2006, SAP NetWeaver was enhanced to extend its use as a business process platform.

SAP NetWeaver reduces total cost of ownership (TCO) across an organization's IT landscape, which frees up resources and enables an organization to refocus on growth. Its components are integrated in a single platform and are preconfigured with applications, which eliminates the need for many integration projects, reduces complexity, and speeds up implementation. The platform supports IT standardization and consolidation so that companies can use IT investments, including SAP and non-SAP systems.

NetWeaver is a web-based, cross-application platform that can be used to develop SAP applications and other applications. By using NetWeaver, a developer can integrate information and processes from geographically dispersed locations by using diverse technologies, including Microsoft .NET, IBM WebSphere, and Oracle Java technologies.

NetWeaver was tagged as a product that can help spur industry adoption of web services. Although web services are often seen as the development model of the future, the implementation rate is not high, often because of competition and incompatibility among enabling products.

NetWeaver can extend SAP solutions with your components in the following ways:

- ▶ If you are starting from scratch, the Composite Application Framework can support the development of components that run in the NetWeaver environment.
- ▶ If the custom components are deployed in a .NET, Java Platform, Enterprise Edition, or WebSphere environment, NetWeaver supports their migration to run in the Web Application Server.

### ***Enterprise components***

The SAP NetWeaver component supports platform-independent web services, business applications, and standards-based development that enable you to use technology assets for web services-oriented solutions. All Enterprise Components that are described in this section run on top of SAP NetWeaver.

NetWeaver features the following components:

- ▶ **SAP NetWeaver Application Server**  
This component supports platform-independent web services, business applications, and standards-based development with which you can use technology assets for web services-oriented solutions.
- ▶ **SAP NetWeaver Composition Environment**  
By using this component, you can develop Java and composite applications from scratch and on top of services.
- ▶ **SAP NetWeaver Business Warehouse Accelerator**  
This component enables you to integrate data from across the enterprise and transform it into practical, timely business information to drive sound decision-making.
- ▶ **SAP NetWeaver Process Integration**  
This component delivers open integration technologies that support process-centric collaboration across the extended value chain.
- ▶ **SAP NetWeaver Mobile**  
This component provides a runtime environment that can be used now and in the future that is based on open and flexible technology standards. It is a powerful development environment for building integrated mobile solutions with native or browser-based user interfaces.



- ▶ **SAP NetWeaver Master Data Management**  
This component ensures cross-system data consistency and helps integrate business processes across the extended value chain.
- ▶ **SAP NetWeaver Enterprise Search**  
This component provides unified real-time access to enterprise data and information.
- ▶ **SAP NetWeaver Identity Management**  
This component unifies identity management and provides user access according to current business roles. It is also used to manage passwords with self-service capabilities and approval workflow.
- ▶ **SAP NetWeaver Visual Business**  
This component provides a user interface that visualizes data from SAP and external data sources in a single window.
- ▶ **SAP NetWeaver Portal**  
This component unifies critical information and applications to give users role-based views that span the enterprise, which enables you to make full use of your information resources.
- ▶ **SAP Gateway (formerly known as *Project Gateway*)**  
This component provides connectivity between devices, environments, and platforms to SAP software by using market standards.
- ▶ **SAP Single Sign On**  
This component integrates into authentication processes to work as single sign-on to enterprise applications, network security, and digital signatures.
- ▶ **SAP Decision Service Management**  
This component aims to simplify and accelerate the management of decision logic that is driving your applications processes.

## **SAP IT management**

SAP includes software to help the IT better manage their assets. One of the key elements of this segment is the SAP Solution Manager. SAP Solution Manager is a central support and system management suite that is provided to SAP customers as part of their license agreement.

The idea is to reduce and centralize the management of these systems and business processes. SAP Solution Manager includes features, such as Template Management, Solution Documentation, Test Management, Upgrade Management, and Change Control System.

SAP also developed an data extraction tool that helps the IT team during data migration between SAP and systems. SAP Test Data Migration Server (TDMS) is a high-speed data extraction tool that can transfer your data from production to quality and test systems to assure accurate data during tests and training.

SAP provides the following IT management software:

- ▶ SAP Solution Manager
- ▶ SAP Test Data Migration Server
- ▶ SAP Test Acceleration and Optimization
- ▶ SAP NetWeaver Adaptive Computing Controller
- ▶ SAP Landscape Transformation
- ▶ SAP IT Infrastructure Management

The following Solution Extensions for SAP Solution Manager are also available:

- ▶ SAP Extended Diagnostics by CA Wily
- ▶ SAP Central Process Scheduling by Redwood
- ▶ SAP IT Process Automation by Cisco
- ▶ SAP Enterprise Modeling by Software AG
- ▶ SAP Productivity Pak by ANCILE
- ▶ SAP User Experience Management by Knoa

### **SAP information management**

SAP also includes data management processes that uses BusinessObjects, Sybase, and OpenText software solutions. The following software is available:

- ▶ SAP Data Services
- ▶ SAP Data Federator
- ▶ SAP Metadata Management
- ▶ SAP Data Quality Management
- ▶ SAP Information Steward
- ▶ SAP Event Insight
- ▶ SAP BusinessObjects Accelerator
- ▶ SAP Sybase Event Stream Processor
- ▶ SAP Sybase PowerDesigner
- ▶ SAP Sybase RAP
- ▶ SAP Sybase Replication Server
- ▶ SAP Archiving and Document Access by OpenText
- ▶ SAP Extended Enterprise Content Management by OpenText
- ▶ SAP Portal Content Management by OpenText
- ▶ SAP Portal Site Management by OpenText

## **1.4 SAP and Lenovo System x: A key relationship**

The Lenovo X-Architecture® family of products meet the growing computing demands of the SAP system landscape with unequalled price and performance. The simplification of deployment and maintenance allows SAP customers to focus on their implementation and better manage their infrastructure costs. These considerations are critical when any complex software solution is deployed, especially SAP.

The Lenovo X6 family is the next generation of powerful, easy-to-use, and cost-effective industry standard IT systems for SAP solutions. It is an ideal platform for SAP customers that need reliability, manageability, and scalability with the flexibility to run Windows or Linux. The Lenovo X6 family is the foundation of a cost-effective SAP solution landscape. This family of products provides a cost-efficient, scalable platform with superior performance and unmatched reliability.

Lenovo X6 offerings feature the following benefits for Lenovo and SAP customers:

- ▶ Unmatched reliability, manageability, and scalability in cost-effective systems
- ▶ Choice of Linux or Windows operating systems
- ▶ Enterprise-class robustness with hot-swap components and robust mainframe that are inspired technologies of the Lenovo X6 platform
- ▶ Designed with high-availability enterprise workloads (such as SAP)
- ▶ Meets the SAP landscape needs with pay-as-you-grow capabilities, which delivers exceptional, flexible, and cost-effective scalability for future growth

- ▶ Server consolidation on to x3850/x3950 X6 servers or Lenovo Flex System x280/x480/x880 X6 compute nodes, which results in fewer physical SAP systems to manage, which lowers costs
- ▶ Up to eight Intel Xeon E7 family processors (24-core processors of 3.2 Ghz, total of 192 cores, and up to 60 MB L3 cache)
- ▶ Up to 12 TB of memory with ECC, Chipkill, Redundant Bit Steering, memory mirroring, and memory rank sparing features
- ▶ Hot-swap parts (hard drives, power supplies, fans, and optional I/O books) of the rack servers (x3850 X6/x3950 X6)
- ▶ Infrastructure integration with Lenovo Flex System solutions

Based on the 64-bit Intel Xeon processor, the Lenovo X6 family of products provides advantages for the following types of customers:

- ▶ Any customer who is evaluating SAP solutions
- ▶ Customers who have reached the limit of their current server capacity and need more power for their SAP application
- ▶ Customers who are running back-level SAP releases and who must upgrade to avoid the end-of-maintenance trap
- ▶ Customers who are looking for new levels of reliability, manageability, and scalability in a cost-effective 64-bit platform
- ▶ SAP customers who is considering the addition of SAP components
- ▶ SAP customers who are considering consolidating and simplifying their application server layer
- ▶ Customers who want to use virtualization technology to increase the use of their infrastructure to save on administration costs

For more information about Lenovo System x servers and Lenovo Flex System compute nodes for SAP solutions, see this website:

<http://www.lenovo.com/solutions/sap/>



# Sizing

Sizing is the process by which the performance requirements of the system (for example, response time and throughput) determine the right hardware configuration. In this context, *right* means able to deliver the expected performance at a minimum cost.

This chapter includes the following topics:

- ▶ 2.1, “Sizing methodology” on page 22
- ▶ 2.2, “Sizing the Lenovo solution” on page 24
- ▶ 2.3, “Summary” on page 26

## 2.1 Sizing methodology

The Lenovo and SAP sizing methodology is based on SAP benchmarks, information from SAP, and actual customer experiences. Lenovo uses sizing tools and customer input to approximate the system resource requirements; however, actual customer results can vary.

This section explains the Lenovo and SAP sizing methodology, terminology, and basic information that can help you to understand the sizing process and sizing estimate results.

### 2.1.1 SAP Quick Sizer

SAP provides all partners and customers with a web-based sizing tool that is called the SAP Quick Sizer. Customers can provide input directly through this tool or provide Lenovo with information so that a sizing can be performed on their behalf.

The SAP Quick Sizer guides you through a structured sizing questionnaire. The Quick Sizer gathers information about your organization's business requirements and translates this data into generic system requirements; that is, platform-independent specifications for processor, memory, and disk. The Quick Sizer offers the following sizing models:

- ▶ User-based sizings

By using this user-based model, you count the number of active users by the SAP functional module. SAP considers this model to be limited in its ability to estimate the SAP resource requirements because it does not consider important sizing factors (such as user behavior, peak versus average workload, the amount of batch processing, reporting, and user customization). SAP recommends the user-based sizing model for small businesses only.

- ▶ Quantity-structure-based sizings

This model is more thorough than the user-based model because it considers actual or expected SAP workload throughput. In addition to the number of SAP users, this model gathers detailed information about the business processes and objects that are used, including the number of dialog transactions, workload profiles, peak usage times, retention periods for business objects, and background and reporting processes. SAP recommends the quantity-structure-based model for medium and large businesses.

You can complete the user-based sizing questions, quantity-structure-based sizing questions, or both. When both models are used, the Quick Sizer provides SAP workload estimates for both models; however, the Lenovo sizing specialist develops the Lenovo hardware recommendation from the two workload-estimates, in concurrence with you.

For more information about SAP references about sizing, see this website:

<http://service.sap.com/sizing>

To run your own sizing, you can use the SAP Quick Sizer tool, which is available at this website:

<http://service.sap.com/quicksizer>

## 2.1.2 Lenovo sizing and planning questionnaire for SAP solutions

The sizing questionnaire is used to gather the requirements that are used to estimate the hardware resources to run SAP Business Suite on Lenovo servers. The questionnaire was designed so that it can be answered without detailed knowledge of SAP solutions. Lenovo established the Lenovo Expert Technical Sales (LETS) team to assist with queries and help size and configure a target SAP solution. The LETS team was established on a geographical basis to address them locally. Contact your local LETS team for information and questionnaires. The LETS team can be contacted at the addresses that are listed in Table 2-1.

Table 2-1 *Lenovo Expert Technical Sales teams*

<b>LETS email address</b>	<b>Geographical location</b>
LETSanz@lenovo.com	Australia/New Zealand
LETSasean@lenovo.com	ASEAN
LETSemea@lenovo.com	EMEA
LETSgcg@lenovo.com	GCG
LETSisa@lenovo.com	ISA
LETSjapan@lenovo.com	Japan
LETSkorea@lenovo.com	Korea
LETSlap@lenovo.com	LA Portuguese
LETSlas@lenovo.com	Spain
LETSna@lenovo.com	North America

The purpose of the questionnaire is to collect information to estimate the Lenovo hardware resources that are required to run the SAP application suite on Lenovo System x. The sizing estimate results include recommendations for processor, memory, and disk for the server infrastructure. In addition, information from the sizing questionnaire can be used by a Lenovo sales representative or Lenovo Business Partner to develop a Lenovo infrastructure proposal, including other hardware and software (for example, for systems management to support the SAP applications).

A sizing estimate is an approximation of the hardware resources that are required to support an SAP solution or component implementation. It is a pre-sales effort that is based on information that is available at a point in time, which provides an entry into understanding the customer's hardware requirements. Your actual experiences vary from the sizing estimate for many reasons, including batch and reporting workloads and custom code. The degree of variability can range from small to significant.

Sizing the hardware requirements for each SAP implementation is an iterative process, which can be refined and repeated several times. Check the sizing input data and estimations during the implementation project. It is important to understand that the sizing estimate is a pre-installation effort, which is based mainly on standard assumptions and benchmark performance data. It cannot replace capacity planning for installed systems.

### 2.1.3 SAP Application Performance Standard

The basic SAP workload unit is SAP Application Performance Standard (SAPS). SAPS is a definition of throughput that was coined by SAP capacity planning and performance testing personnel. A total of 100 SAPS are defined as 2,000 fully business-processed order line items per hour in the standard SAP Sales and Distribution (SD) application benchmark. This amount is equivalent to 2400 SAP transactions per hour.

The capability of processors is measured during the standard SD benchmark test, which is certified by SAP and rates the processor on its capability to process work at 100% CPU loading.

The SAP Quick Sizer calculates the workload (in SAPS) and adjusts it to allow for a suitable processor use. Therefore, if a workload of 4,800 SAP SD benchmark transactions per hour was required, the Quick Sizer calculates this requirement as 200 SAPS; however, allowing for a target processor load of 33%, it adjusts this amount to find a processor that is capable of 600 SAPS @ 100% (=200 at 33%).

In an SAP environment, a processor can perform the following workloads:

- ▶ Database Server Function
- ▶ Central Server Function
- ▶ Application Server Function

Lenovo rates each of its processors on its ability to perform each of these functions. These calculations are based on SAP-certified benchmarks. Lenovo uses its judgment to calculate combinations that are not benchmarked. SAP-certified benchmarks are performed for database servers and central servers only. Therefore, application server numbers are calculated on the results of considerable field experience that are expressed as the number of application servers per database server.

Most customers run database servers with central instance functions (for example, messaging, lock management, and spooling) that are running on the database server. As a result, the SAP benchmarks are run with these functions offloaded.

SAPS ratings relate to an absolute business-related workload. Therefore, the changing complexity of SAP applications causes the ratings to vary from one release of SAP to the next.

## 2.2 Sizing the Lenovo solution

Lenovo uses information from the SAP Quick Sizer and the questionnaire as input to the hardware mapping process. The Quick Sizer provides processor, memory, and disk requirements. This section explains how Lenovo uses information from the Quick Sizer tool to develop the Lenovo hardware recommendations.

### 2.2.1 Target CPU utilization

Based on the sizing inputs, the Quick Sizer approximates processor, memory, and disk consumption. For user-based sizings, SAP sizing results are calculated to meet an average CPU use of 65% for the online interactive workload, including a buffer for unspecified workloads (such as batch jobs, reporting, printing, backup, and software interfaces) and for peak periods. The Quick Sizer quantity-structure-based model also uses a target CPU use of 65%. Different target CPU utilizations can be calculated with the mutual concurrence and responsibility of the Lenovo sales team and the customer.



## 2.2.2 Resource categories

For the processor and the disk, the Quick Sizer results are expressed as resource categories, which are SAP's hardware independent resource specifications.

For CPU sizing, each resource category represents a range of SAPS and disk space. Table 2-2 shows an example of the CPU and disk resource categories and the SAPS and disk space requirements.

Table 2-2 CPU and disk sizing resource categories

Category	Up to ... SAPS	Up to ... GB disk	Up to ... IOPS
XS	8,000	400	3,200
S	16,000	500	6,000
M	32,000	1,000	12,000
L	48,000	2,000	20,000
XL	72,000	2,500	28,000
XXL	Contact Lenovo (LETS) or SAP for detailed sizing		

If your CPU sizing exceeds 72,000 SAPS, your disk sizing exceeds 2.5 TB, or the I/O value is more than 28,000 I/Os per second, contact your local Lenovo LETS team or SAP for a detailed sizing.

Figure 2-1 shows SAP Quick Sizer that refers to these resource categories.

The screenshot shows the 'Active Users Results for SAP Business Solutions' table. The 'CPU cat.' column contains 'M' and the 'DB Disk cat.' column contains 'L'. Both cells are circled in red. The table also shows 'SAPS (total)' as 7.800 and 'DB Disk (MB, total)' as 420.000.

Solution	Release	CPU cat.	SAPS (total)	Memory (total, MB)	DB Disk cat.	DB Disk (MB, total)
NW-APPSRV	640	M	7.800	46.080	L	420.000

Figure 2-1 Resource categories in SAP Quick Sizer

## 2.2.3 Steps in the sizing process

Lenovo performs the following steps for every Quick Sizer sizing request:

1. Determines the potential SAP workload requirements.

The SAP Quick Sizer analyzes your input and calculates a generic sizing result. The Quick Sizer provides the storage requirement and resource categories for CPU and disk.

**Important:** If you completed the user and quantity-structure based sizings with reasonable inputs, discuss the two workloads with your local LETS team to determine the Lenovo hardware requirements, which can vary subject to the analysis of the two results.

2. Selects a 2-tier or 3-tier hardware configuration.

Based on the potential workload, select a 2-tier or 3-tier hardware configuration:

- In a 2-tier configuration, one server provides the database and application server functions.
- In a 3-tier configuration, there is one database server and one or more separate application servers.

If one server can handle the database and application server workloads, select the 2-tier configuration. This configuration often is appropriate for smaller SAP workloads. If your workload does not fit in a single server, select a 3-tier configuration (consider special system landscape issues).

3. Determines the Lenovo server requirements.

Based on the SAPS capacity rating for each Lenovo processor, find a database or central server that can support at least the number of SAPS that are indicated by the Quick Sizer. For application servers, the work can be split across several servers whose joint capability must exceed the target SAPS number.

## 2.3 Summary

The Lenovo and SAP sizing methodology is continually reviewed and revised to provide the best possible estimate of the Lenovo hardware resources that are required to run SAP. Guidelines for sizing SAP solutions come from a number of sources, including SAP, SAP benchmarks, and customer feedback. Based on this information and your completed sizing questionnaire, the LETS team analyzes your requirements and recommends a Lenovo hardware configuration.

For more information about the SAP database benchmark for two- and 3-tier configurations, see this website:

<http://global.sap.com/campaigns/benchmark/index.epx>

# Hardware and software considerations

This chapter describes major software and hardware factors for future SAP-based solution installations and includes the following topics:

- ▶ 3.1, “SAP Business Suite hardware implementation strategy” on page 28
- ▶ 3.2, “Virtualization” on page 28
- ▶ 3.3, “Solution architecture considerations” on page 33
- ▶ 3.4, “Choosing the correct hardware” on page 37

## 3.1 SAP Business Suite hardware implementation strategy

During production operation, most customers maintain non-production and production systems that are running on separate Lenovo servers. This strategy provides the following advantages:

- ▶ Test new releases and system modifications without affecting users.
- ▶ Minimize initial hardware investment by installing the hardware in phases.
- ▶ Increase security and segregate physical access to production information.

During the application development phase, customers typically install one or more non-production systems for development and testing of the SAP system. In many cases, each non-production system is a 2-tier configuration that is running on its own server. However, it is a good idea to have at least one system that is an exact replica of production, which often is the QA system. The QA system then can be used for performance testing, feature or functionality upgrades, and other changes to a real replica of the production system and data. This process is another way to ensure that the productive system remains stable and performs as needed.

In addition to the development and quality assurance systems, customers install other servers for other non-production systems (for example, training system, and sandbox system).

When the SAP system is ready to run in production, customers install the production system hardware. For 3-tier configurations, more application servers can be installed and the database server upgraded, as needed as new users are added to the production system. For 2-tier configurations, the central database/application server can be upgraded to handle the extra workload as users are added to the system. Also, if a 2-tier installation grows beyond the capability of a single-server, it can be configured as a 3-tier configuration by adding application servers, which is when customers often split the SAP Database and Central Services from each other.

## 3.2 Virtualization

This section describes the reasons for virtualization of SAP systems on the x86 platform and the options that are available.

### 3.2.1 Advantages of virtualization

There are several motivations for the use of virtualization. The primary focus for all of the motivations is increased utilization and reduction of the total cost of ownership (TCO) of the IT landscape.

#### **Landscape consolidation**

Today's SAP landscapes are growing more complex. Starting with the classical development, test, and production system landscape, extending these landscapes with sandbox or educational systems, deploying these systems into a 3-tier configuration with multiple servers per system, and adding infrastructure-related systems (such as the SAP Solution Manager, SLD, and SAP Router), your SAP system landscape can easily grow to several dozens of physical SAP system servers.

SAP systems are normally sized for 65% of the total capacity that is required. This sizing ensures that there is enough computing resources to handle the critical peak workload, which often represents only a small part of the system's day, month, or even year. By consolidating several dedicated servers (each running at an average low use with periodic peaks) into a shared virtualized environment, you can achieve a reduction of physical servers, which leads to less cost and reduced complexity.

### **Load shifting**

The ability to move virtualized SAP systems seamlessly between physical servers enables you to move applications that temporarily demand more resources than the physical server on which they are currently running to a more powerful server easily when needed. This way, an application that needs increased resources only once a year, once a quarter, or once a week (for example, for a high-load batch run) can be moved to a powerful server for that task and can be on a less powerful server for the other times. In general, virtualization gives you much more flexibility in SAP system planning, which enables you to use the physical resources much more effectively.

### **Optimized application availability**

Another useful scenario is to free physical servers for scheduled maintenance. Before taking down the server for the maintenance, all virtualized SAP systems are moved to other available physical servers in the datacenter. This way, you can shut down the server without affecting the application's availability. Therefore, you can maximize the application uptime while servicing your servers independently from application availability requirements. If a physical server fails, you can start a virtualized SAP system again on a new server quickly (represented as files on a shared storage system) without having to wait for hardware service.

### **Rapid provisioning**

Virtual machines can be duplicated easily because they are represented by a set of files. With virtualization, it takes only minutes for a new virtual server to be up and running. This feature is an important advantage compared to the time that it takes to put a physical server in place. The physical server must be mechanically mounted and connected to cables. Then, the operating systems must be installed, patches must be applied, parameters must be modified, and the wanted application must be installed.

Rapid provisioning helps in quickly setting up of more SAP application servers to increase the performance of an SAP system by scaling out and enables providing sandbox SAP systems for testing or education, on-demand within minutes instead of days. You can use resources that might otherwise remain unused and sit idle waiting for higher workload times.

These possibilities, which are provided by virtualization, help to increase server use and application availability and to provide a high degree of flexibility.

## **3.2.2 Virtualization technologies**

A number of virtualization technologies are available for SAP on Lenovo System x. This section provides an overview of these technologies and what is supported by SAP for productive SAP systems.

## Supported guests systems overview

Table 3-1 lists the support status, depending on the virtualization technology and operating system (OS) that are used for the virtualized SAP system.

Table 3-1 Virtualization support status overview

Virtualization technology	Windows Server guest OS	SUSE Linux Enterprise Server guest OS	Red Hat Enterprise Linux guest OS
VMware ESX	Yes	Yes	Yes
Hyper-V	Yes	No	No
Kernel-based virtual machine (KVM)	No	Yes	Yes

The following sections provide more information about support and restrictions for each combination that is listed in Table 3-1. For more information about supported operating systems versus virtualization technology, see SAP note #1492000 – *General Support Statement for Virtual Environments*.

## VMware ESX

SAP supports VMware virtualization in production and non-production environments. Table 3-2 lists Windows guest systems that are supported by VMware running SAP.

Table 3-2 Supported Windows 64-bit guest operating systems with SAP on VMware

Guest OS	VMware ESX Server 3.x and VMware vSphere 4.x, 5.x, 6.x
Windows Server 2003	Supported
Windows Server 2008	Supported
Windows Server 2008 R2	Supported
Windows Server 2012	Supported
Windows Server 2012 R2	Supported

Table 3-3 lists Linux guest operating systems that are supported by VMware that is running SAP.

Table 3-3 Supported Linux guest operating systems with SAP on VMware

Guest OS	VMware ESX Server 3.x and VMware vSphere 4.x, 5.x, 6.x
Red Hat Enterprise Linux 5	Supported
Red Hat Enterprise Linux 6	Supported
Red Hat Enterprise Linux 7	Supported
SUSE Linux Enterprise Server 8	Supported
SUSE Linux Enterprise Server 9	Supported
SUSE Linux Enterprise Server 10	Supported
SUSE Linux Enterprise Server 11	Supported

For more information about the support status and restrictions that apply, see the following notes:

- ▶ #1409608 – *Virtualization on Windows*
- ▶ #1122387 – *Linux: SAP Support in Virtualized Environment*

For more information about support for Windows and Linux that run on VMware and the databases that are supported, see this website:

<http://scn.sap.com/docs/DOC-27384>

For servers that have Lenovo ServerProven® status for VMware ESX<sup>1</sup> and that are certified for SAP on Windows systems or for SAP on Linux, VMware ESX is supported to run SAP software in a virtualized Windows or Linux environment.

## Hyper-V

Hyper-V is included with Microsoft Windows Server 2008, 2008 R2, 2012, and 2012 R2. Table 3-4 lists the supported guest systems that run on Hyper-v with SAP.

Table 3-4 Supported Linux guest operating systems with SAP on Hyper-v

Guest OS	Hyper-V Windows Server 2008 (R2)	Hyper-V Windows Server 2012	Hyper-V Windows Server 2012 (R2)
Windows Server 2003	Supported	Not Supported	Not Supported
Windows Server 2008	Supported	Supported	Supported
Windows Server 2008 R2	Supported	Supported	Supported
Windows Server 2012	Supported	Supported	Supported
Windows Server 2012 R2	Not Supported	Supported	Supported

Desktop versions of these operating systems are not relevant for SAP.

**Linux support:** Red Hat Enterprise Linux (RHEL) and SUSE Linux Enterprise Server (SLES) are not supported to run on Hyper-V.

For more information about the supported Windows systems that run on Hyper-V and the databases that are supported, see this website:

<http://scn.sap.com/docs/DOC-40885>

For more information about the support status and restrictions that apply, see SAP note #1246467 – *Hyper-V Configuration Guideline*.

More information about virtualization of SAP applications on Windows, see this website:

<http://www.sdn.sap.com/irj/sdn/windows-virtualization>

<sup>1</sup> For more information about Lenovo System x servers, see:  
<http://www.ibm.com/systems/info/x86servers/serverproven/compat/us/nos/vmware.html>

## Kernel-based virtual machine

The newest virtualization technology that is approved for productive use with SAP is kernel-based virtual machine (KVM). It is included as part of RHEL and has approval from SAP to run the following guest releases:

- ▶ Red Hat Enterprise Virtualization (RHEV) 3.0 and 3.1
- ▶ RHEL 5.4 to 5.8 and 6.1 to 6.3, 7
- ▶ SLES 11 SP1 or later
- ▶ SLES 12 SP1 and later

Para-virtualization is not possible with KVM, and all supported versions must be 64-bits. You can run your host system only on hardware that is certified for Linux KVM environments. For more information about the support status and restrictions that apply, see the following notes:

- ▶ *#1522993 – Linux: SAP on SUSE KVM – Kernel-based Virtual Machine*
- ▶ *#1400911 – Linux: SAP on Red Hat KVM – Kernel-based Virtual Machine*

For more information about ServerProven hardware, see the following resources:

- ▶ SAP note *#171380 – Linux Released IBM hardware (Intel and AMD processors)*
- ▶ The following website and click **Supported Platforms** under the Key Topics section:

<http://www.sap.com/linux>

## 3.2.3 Lenovo X6 for virtualized SAP systems

Lenovo X6 enterprise servers are the ideal platform for business-critical and complex SAP applications, such as database processing, customer relationship management (CRM), and enterprise resource planning (ERP) applications, and highly consolidated, virtual server environments. With multiple workloads that are running on the same server, performance remains important, but reliability and availability become more critical than ever.

The Lenovo X6 family of scalable rack servers consists of the Lenovo System x3850 X6 server (a 4U 4-socket server) and the x3950 X6 server (an 8U 8-socket server). By using proven technologies of the previous generations of Enterprise X-Architecture, these servers introduce the following benefits:

- ▶ New levels of fault tolerance and resiliency with advanced reliability, availability, and serviceability (RAS) features that are implemented in hardware and software
- ▶ Agility and scalability with fit-for-purpose modular “bookshelf” that is ready to support multiple technology upgrades
- ▶ Significant improvements in response time with ultra-low latency, stretched memory speeds that exceed Intel specifications and innovative flash memory-channel storage offerings

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

Lenovo X6 systems that use the Intel Xeon processor E7 v2, v3, and v4 families deliver an extensive and robust set of integrated advanced RAS features that prevent hardware faults from causing an outage. Part selection for reliability, redundancy, recovery, and self-healing techniques and degraded operational modes are used in a RAS strategy to avoid application outages. By using this strategy, Lenovo X6 systems can help increase application availability and reduce downtime by enabling 24x7 mission-critical capabilities to run your core business services.



## 3.3 Solution architecture considerations

The implementation of SAP systems is a complex and expensive effort that can affect company life. Because the introduction of an ERP system entails the re-engineering of internal business processes, the introduction of SAP concerns more than only the IT department. Entire departments (and sometimes the entire company) often are asked to participate in the definition of new, optimized SAP-controlled business processes. In addition to dedicating employees to the support of SAP customizing and Advanced Business Application Programming (ABAP) or Java development, the company must invest in IT resources to ensure that system operational requirements are met, particularly performance, availability, and serviceability.

With the focus on functional requirements for a solution (in most cases regarding non-functional requirements, such as performance, reliability requirements, or the ability of an SAP installation to handle increasing loads) your organization must consider the options that are described in the following sections when an SAP solution is designed.

### 3.3.1 Scaling up or scaling out

Scalability is key to the successful implementation of an ERP software. Your business grows, and you expect this ability to grow from your SAP system. The SAP software is designed for scalability, with the ability to perform the following tasks:

- ▶ Scale up (also known as *vertical scaling*) with a 2-tier system
- ▶ Scale out (also known as *horizontal scaling*) with a 3-tier system

#### 2-tier system

This type of topology presents a configuration where all components are on the same physical server. It is easy to install and maintain and, therefore, is appropriate for first-time installations where you want to test and familiarize yourself with the functions of the application server.

At this stage, it is easy to determine what kind of topology is needed and to plan the system landscape. Although this topology is easily configurable and inexpensive, it also has drawbacks that you must consider. Because all components are on a single machine, they compete for resources, which affects the performance processes.

A 2-tier system configuration has the following key advantages:

- ▶ Uses the power of 64-bit and scalability.
- ▶ Provides the best performance, with no overhead for database connection, no network traffic, and no shadow processes
- ▶ Makes it easy to administer

The 2-tier approach is suited for small installations or installations with performance demands not being dynamic. Large installations (for example, the use of the System x3950) can use the extra performance that you can achieve from this type of installation.

#### 3-tier system

In a 3-tier system, the application servers run separated from the database server on different machines. By separating the servers onto separate machines, the capacity of a single machine in terms of performance can be lower than in a 2-tier environment.

A 3-tier system configuration offers the following key advantages:

- ▶ Lower total cost of acquisition (TCA) for large configurations
- ▶ Scales flexibly (scale out)
- ▶ Affordable and flexible high availability (HA) concepts
- ▶ Application dependent, because some SAP applications require a 3-tier configuration, for example, Advanced Planner and Optimizer (APO)
- ▶ Cover short-time load peaks easily by adding application servers temporarily

In general, the 3-tier approach is the more flexible and often the only viable option, especially for large installations.

### 3.3.2 High availability

As more customers use SAP systems for global operations to support mission-critical business functions (such as sales and order-entry and continuous manufacturing) the need for maximized system availability becomes crucial. Many companies now require 24x7 reliability for their SAP systems.

*Clustering* is the use of multiple computers that are connected to act as a single entity to run common tasks. If one computer fails, the other computers pick up the workload automatically. By adding more computers, you can increase the power, and the design can support more users.

*Switchover clusters* ensure HA of the SAP system by switching critical services, which represent a single point of failure (SPOF), across multiple hosts in the cluster. When a primary node fails or is taken down for maintenance, switchover software automatically switches the failed software unit to another hardware node in the cluster. Applications or users that are accessing the failed service experience a short delay but resume normal processing after the switchover. This method is the longest supported and most common method of achieving HA in a traditional SAP system.

When an SAP system is designed to use a switchover cluster technology for HA, it is important to know the effect of the sizing effort. Best practices require sizing each cluster node for 100% of the total required capacity to ensure that the surviving cluster node can handle the increased workload in a failover condition.

Figure 3-1 shows an example of a switchover setup for SAP.

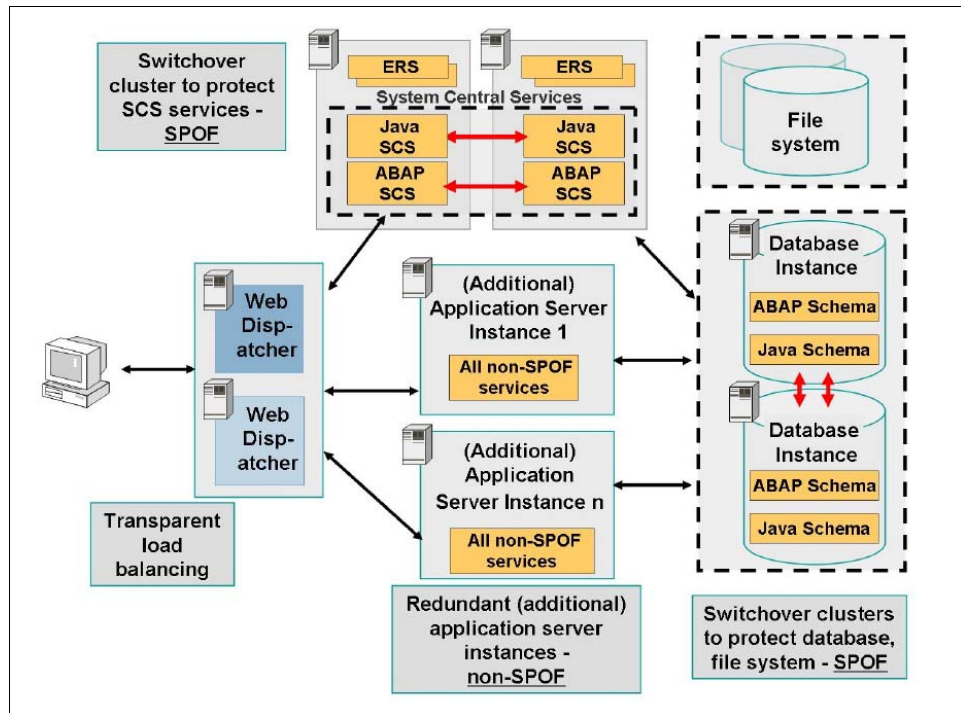


Figure 3-1 Example of a switchover setup for SAP (© SAP AG 2007, 2010. All rights reserved.)

For more information about how to achieve HA for SAP Solutions at this SAP Community, see this website:

<http://scn.sap.com/docs/DOC-7848>

For more information, see SAP note #803018 – NetWeaver High Availability capabilities.

The following section describes the use of Veritas software for an HA solution that is available to complement the Lenovo System x configuration.

### HA with Veritas software

Veritas Storage Foundation for Windows systems brings advanced volume management technology to Windows Server 2003, 2008 (R2), and 2012 (R2). The combination of Veritas Storage Foundation, Veritas Storage Foundation Cluster File System, Veritas Cluster Server HA/DR, Veritas Volume Replicator, and Cluster Server agents that are designed specifically for SAP applications enables you to automate cluster monitoring, and in the case of a detected failure, enables a graceful failover to your SAP system.

For more information about how to create HA and disaster recovery for your SAP application, see this website:

[http://www.symantec.com/connect/sites/default/files/13803401-1\\_GA\\_WP\\_HADR-SAP\\_02.08.pdf](http://www.symantec.com/connect/sites/default/files/13803401-1_GA_WP_HADR-SAP_02.08.pdf)

## HA with Microsoft Cluster Server

By using Microsoft Cluster Server (MSCS), a server's resources can fail over from one machine to another. MSCS is part of the operating system in Windows Server 2003, 2008, 2008 R2, 2012, and 2012 R2. This service provides HA and scalability for mission-critical applications (such as databases, messaging systems, and file and print services) and can be used to implement a highly available SAP system.

For more information about MSCS for all SAP software, see this website:

<http://scn.sap.com/docs/DOC-7848#section24>

### 3.3.3 When to use virtualization

Although virtualization has its advantages, it is not always the best decision to virtualize everything. Virtualization comes with a price in terms of performance and limitations. When virtualization is used, consider the following important items:

- ▶ Processor virtualization

Processor virtualization provides virtual processors to the virtualized guest systems. Each virtualized guest appears to run on a dedicated processor or set of processors, with each processor having its own registers and control structures. Virtualization solutions often limit the number of these virtual processors per virtualized guest system.

**Note:** A virtualized cpu (vCPU) does not equal the performance of a physical CPU core. This reduction occurs because with Hyper-Threading enabled (highly recommended from Intel processors formerly codenamed *Nehalem* and newer), VMWare sees a thread as a CPU. Therefore, there are two Virtual CPUs per core in Intel based systems.

- ▶ Memory virtualization

Every access to memory in a virtualized environment must be managed by the virtualization layer (the hypervisor), which translates virtual memory addresses into physical addresses. Memory virtualization is supported by the processors with features (such as page-table virtualization) that allow the hypervisor to efficiently handle memory access by the guest systems.

As with processors, the virtual memory that is available to virtualized guests might be limited, depending on the virtualization solution that is used. Another possible problem with memory virtualization is that some virtualization software does not ensure the memory to host. Therefore, some applications that run only in real memory might have problems during paging if there was no real memory any longer.

- ▶ I/O virtualization

In a virtualized environment, I/O subsystems (such as network, local disk, or Fibre Channel access) also must be virtualized. I/O virtualization was a problem in the early days of x86-virtualization and to an extent still is. Workloads that generate high volumes of I/O tend to see a performance degradation when virtualized. An example of such a workload is a database server. Although it might be a viable path to go for smaller database installations, consider running high-load database servers on physical hardware (that is, not virtualized), and virtualize only the application server parts of an SAP system.

Lenovo and several other SAP hardware partners conducted SAP SD benchmarks with virtualization enabled to show the performance difference between a virtualized and non-virtualized system. These benchmarks run a single guest on the physical server, by which the performance penalty for the introduction of the virtualization layer into the stack is shown. Benchmarks indicate that virtualization (depending on technology that is used) introduces a performance degradation of less than 10%. Considering the low accuracy that can be achieved with a usual sizing, this number is small enough not to affect how you size your SAP solution.

In addition to the performance penalties, each of the virtualization technologies has limits regarding the maximum number of virtual processors, the maximum amount of memory, and, possibly, other parameters of the virtualized system.

## 3.4 Choosing the correct hardware

Choosing the correct hardware for a complex solution such as SAP can be a challenge. You must consider many facets, as described in the following sections.

For more information about the hardware or combination of hardware that is supported by SAP, see the following website (an SAP S-user ID is required):

<http://service.sap.com/platforms>

### 3.4.1 Servers supported by SAP

SAP supports various platforms for their software, but SAP AG requires a special SAP certification for each hardware or software platforms that run SAP applications with a few exceptions, such as SAP Business One. To be supported by SAP, every SAP software component (except the presentation tier) must run on SAP-certified platforms.

SAP platform certification includes the benefits:

- ▶ Allows SAP to reference a list of reliable and standardized platforms when a customer requests that information
- ▶ Offers a service to SAP customers with greater market transparency and a way to compare offers
- ▶ Delivers a performance indicator that can be publicly available and quotable
- ▶ Ensures that the SAP-certified hardware platform is stable and scalable when implementing SAP solutions

When certified, an SAP-certified hardware platform is then supported for all SAP releases that are available, including all SAP approved databases.

#### Definitions

The following terms are used regarding the formal SAP hardware certification process:

- ▶ *Processors*, *cores*, and *threads* are defined as specified by the Transaction Processing Performance Council (TPC) CPU working group. For more information about these definitions, see this website:

<http://www.tpc.org>

Under the CPU working group report from 15.08.2005, these terms have the following definitions:

- Processor

A component that contains one or more cores. The number of processors that is claimed by the test sponsor must be consistent with the way the product is represented in the test sponsor's marketing collateral, including that which is not related to TPC benchmark results. The reported number of processors must be the number of processors that are enabled for the benchmark.

- Core

The core is the execution unit can run one or more processor threads. If a hardware unit can run more than one concurrent processor thread without requiring the processor threads to share execution resources, it is more than one core. It can rely on other assist units (such as cache and hardware accelerators), some or all of which can be shared by multiple cores. The reported number of cores must be the number of cores that are enabled for the benchmark.

- Execution unit

This unit is the electronic circuits that are necessary to implement the semantics of all possible instructions in computer architecture.

- Processor thread

This thread is hardware necessary to maintain the state of a software thread. The reported number of processor threads must be the number of concurrent processor threads that are enabled for the benchmark.

- Software thread

The thread is an instruction sequence that performs operations within an address space and is scheduled by software.

- ▶ A *hardware platform* includes the following definitions:

- By the processor family

- Determined by the processor manufacturer; for example, Intel, AMD instruction set. Mere differences in processor speed or cache size is considered as the same family.

- By the chipset or the on-chip controller architecture, which includes all components that enable the data transfer between processor, memory, and I/O.

- External I/O (such as the storage adapter or equivalent [Fibre Channel, RAID controller, and so forth]) and network access are not part of the platform definition.

- ▶ A *server family* is formed by those servers that share the hardware platform as defined. They are considered as one platform and do not depend on the type of physical construction (Flex Systems nodes, Towers, rack servers, and so forth) or on the number of processors.

## Server certification for SAP on Windows

All SAP technology partners that have a partner and support contract with SAP or technology partners that are accepted by any SAP regional or country branch can certify platforms for running SAP solutions on Microsoft Windows systems.

### **Server certification rules**

To achieve certification, a Windows server must adhere to the following rules:

- ▶ The certified hardware must be listed within the Microsoft Hardware Compatibility List (HCL) for all Windows versions that are supported by this platform.
- ▶ SAP allows only 64-bit processors to be certified.
- ▶ I/O controllers that are supported by the hardware vendor for SAP solutions must be declared within the certification. No separate certification is required for I/O adapters, storage, and other hardware components. Therefore, the vendor vouches for functionality and support of these components. Recommendations from hardware vendors are to be observed.
- ▶ SAP requires that the hardware vendors certify at least one system per server family. The certification of a server family must be completed on the largest system. The size of the system is defined by the largest number of processors, cores, and threads that are possible in the system. A certification includes all smaller systems. If questions arise, it must be clarified through the hardware certification workgroup.
- ▶ One benchmark must be completed to certify a server family. When the certification is granted, new larger family members can be added to the certification without publishing a benchmark immediately. A new benchmark must be completed for the new or larger family member within six months of the date the new member was added. When no benchmark is completed for the announced members of the server family, the certification is revoked for all servers larger than the family member with a published benchmark.

The certification includes a performance indicator for the certified platform. The following tasks must be completed successfully:

- ▶ The tests consist of an SAP SD 2-tier, SD 3-tier, or SD parallel standard application benchmark. The general SAP standard application benchmark<sup>2</sup> rules apply.
- ▶ Any Windows operating system and any database system that is released for the aforementioned SAP release (or is released in the next 180 days) can be used.
- ▶ The minimum requirement for the 2-tier benchmark is 80% average processor use in the high load phase.
- ▶ Only the database server is certified on a 3-tier benchmark.

### **Determining which servers are supported**

After successful certification, the certified systems are listed on the SAPonWin.com website. For more information about certified Lenovo System x servers, see the following resources:

- ▶ <http://saponwin.com/pub/hardware.asp?l=vendor&sl=12&la=en>
- ▶ See the following website and click **Vendor of Certified Hardware** → **Lenovo**:  
<http://www.SAPonWin.com>

The corresponding SD benchmark that is used for the certification is published at the following SAP website:

- ▶ <http://www.sap.com/benchmark>

### **Server validation for SAP on Linux**

The vendors of the operating systems must have a partner and support contract with SAP or at least have a support agreement with SAP. For Linux, SAP-supported vendors are Red Hat and Novell.

---

<sup>2</sup> See <http://www.sap.com/benchmark>

## **Server certification rules**

To achieve certification, a Linux server must adhere to the following rules:

- ▶ The certified hardware must be listed on the Red Hat Compatibility List (for Red Hat Enterprise Linux) or the Novell Hardware Compatibility List (for SUSE Linux Enterprise Server). For Lenovo servers, the hardware must be listed as supported for RHEL/SLES in the Lenovo ServerProven list.
- ▶ I/O controllers that are supported by the hardware vendor for SAP solutions must be declared within the certification. No separate certification is required for I/O adapters, storage, and other hardware components; therefore, the vendor vouches for functionality and support of these components. Recommendations from hardware vendors must be observed. The certification is valid for all Enterprise Linux versions and releases from the Linux distributors if they are supported by SAP and the hardware vendor.
- ▶ Hardware vendors must certify at least one system per server family. Server family certification must be completed on a representative system. The size of the system is defined by a reasonable number of processors (at least 4-way unless only 2-way is available by the vendor), and cores and threads that are possible in the system. If questions arise, the questions must be clarified through the hardware certification workgroup.
- ▶ The hardware vendor is fully responsible for providing support in case of any doubts that are related to the use of particular hardware with SAP.

The certification consists of a set of tests to verify the stability of the certified platform. One of the following specifications must be completed successfully:

- ▶ **Minimum**  
Tests must consist of functional tests that represent an SAP implementation onsite at the customer location.
- ▶ **Recommended**  
The recommended tests also can consist of an SAP SD 2-tier, SD 3-tier, or SD Parallel Standard Application Benchmark. The general SAP standard application benchmark rules apply. This benchmark does not have to be published. For more information, see the following website:  
<http://www.sap.com/benchmark>
- ▶ **Best of Breed**  
The best tests include one or both of the minimum and recommended tests and many other tests, such as I/O tests, ABAP sample reports and transactions, memory tests, and other tests as determined by the hardware partner to show the stability and reliability of the server family that is certified.

For an easier and faster certification process, SAP developed the SAP Certification Suite for Linux, which was tested and approved for use by SAP, the hardware partners, and the Linux distributors.

Specifically, the vendor can use a Linux operating system and any database system that is released and certified for use with SAP. When the hardware vendor chooses to test by using the SAP SD benchmark, the minimum requirement for the 2-tier benchmark is 80% average processor use in the high-load phase.



### ***Determining supported servers***

For more information about Lenovo Systems x servers that are certified for SAP on Linux, see the following resources:

- ▶ SAP note #171380
- ▶ SAP on Linux website:  
<http://www.sap.com/linux/>
- ▶ Red Hat (click **Vendor of Lenovo Certified Hardware** for each Linux distribution):  
<http://hardware.redhat.com/hcl/>
- ▶ SUSE (click **Vendor of Lenovo Certified Hardware** for each Linux distribution):  
<http://developer.novell.com/yesssearch/>



# Lenovo X6 systems, storage, and networking offerings

The Lenovo Systems portfolio is a valuable platform for implementing SAP applications. This chapter describes the Lenovo X6 rack server and Flex System offerings for SAP implementations. It also provides a brief review of the Lenovo Storage and Lenovo Networking offerings that can be used with the Lenovo servers for complete end-to-end infrastructure solutions.

This chapter includes the following topics:

- ▶ 4.1, “Introduction to Lenovo X6” on page 44
- ▶ 4.2, “Lenovo X6 systems overview” on page 45
- ▶ 4.3, “Lenovo X6 rack server design and architecture” on page 49
- ▶ 4.4, “Lenovo Flex System X6 compute node architecture” on page 66
- ▶ 4.5, “Lenovo Flash and SSD internal storage” on page 74
- ▶ 4.6, “Lenovo Storage offerings” on page 85
- ▶ 4.7, “Lenovo Networking offerings” on page 92

## 4.1 Introduction to Lenovo X6

The Lenovo X6 product family represents the sixth generation of the Enterprise X-Architecture. It delivers fast application performance, is based on an agile system design, and is a resilient platform that is needed for mission-critical databases, enterprise applications, and virtualized environments.

Lenovo X6 systems are scalable systems that can be expanded on demand and configured by using a building block approach. By using this approach, your server can grow in processing, input/output (I/O), and memory dimensions. Therefore, you can provision what you need now and expand the system to meet future requirements. System redundancy and availability technologies are more advanced than those technologies that were available in the x86 systems.

The X6 portfolio increases performance and virtualization density while decreasing infrastructure costs and complexity. By using this function, you can design faster analytics engines, reign in IT sprawl, and deliver information with high reliability. Lenovo X6 servers are fast, agile, and resilient.

### 4.1.1 Fast application performance

The x3850/3950 X6 rack servers and the Flex System x280/x480/x880 X6 compute nodes deliver fast application performance. With the new Intel Xeon processor E7 v4 family, the rack servers (x3850/x3950 X6) and the Flex System compute nodes (x280/x480/x880 X6) can deliver up to 12 TB of memory and 192 cores of processing power. Armed with these capabilities, you can host essential mission-critical applications, implement large virtual machines (VMs), many VMs, or run sizeable in-memory databases without compromises in performance, capacity, or scalability.

These business-critical, enterprise-class servers deliver fast performance with the Lenovo flash storage offerings for the X6 servers. Flash storage combines extreme IOPS performance and low response time for transactional database workloads. The flash technologies that are used in the X6 servers include PCIe NVMe drives and adapters (higher bandwidth and shorter latency) and Flash Storage Adapters (high performance NAND technology).

The X6 product family can also use high performing solid-state drives (SSDs). The use of SSDs instead of, or along with, traditional 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, which increases the performance of analytical workloads, transactional databases, and virtualized environments.

### 4.1.2 Agile design characteristics

The unique, adaptive modular design of the X6 rack server family and the flexibility of the Flex System compute nodes deliver agility that enables you to design a solution that meets your needs. At the same time, you can realize infrastructure cost savings by hosting multiple generations of technology in a single platform, without compromising performance or capacity.

The X6 platforms provide the following capabilities:

- ▶ You can configure the server or chassis to fit the unique requirements of your applications and workloads. You can also easily add, modify, or upgrade the X6 rack server or add extra nodes to a Flex System chassis by using modular components.

- ▶ You can scale capacity and performance by adding compute nodes to Flex System or installing extra processors and memory into X6 rack servers, which delivers the performance for growing applications without creating IT sprawl.
- ▶ You can use Lenovo XClarity software for automated provisioning of operating system deployments and firmware updates and realize time-to-value in minutes rather than days.
- ▶ You can use Lenovo XClarity Energy Manager, which provides energy management by monitoring power and temperature at the server level and the group level to improve business continuity and energy efficiency.
- ▶ You can realize an agile system design that is ready to host multiple generations of technology in a single server.

### 4.1.3 Resilient enterprise platforms

By using differentiated X6 self-healing technology, the X6 servers maximize uptime by proactively identifying potential failures and transparently taking necessary corrective actions. The X6 servers include the following unique features:

- ▶ Advanced Page Retire proactively protects applications from corrupted pages in memory, which is crucial for scaling memory to terabytes.
- ▶ Advanced Processor Recovery allows the system to automatically switch access and control of networking, management, and storage if there is a processor 1 failure, which provides higher availability and productivity.
- ▶ Lenovo XClarity Integrator for standard hypervisors enables the creation and management of policies to maintain high availability of VMs and concurrent updating of the system firmware, with no affect on application performance or availability.
- ▶ The X6 rack servers have a modular design that reduces service time by enabling quick and easy replacement of failed components.
- ▶ When virtualized, the X6 compute nodes can move workloads between nodes to avoid interruptions during system maintenance or system failure.

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

## 4.2 Lenovo X6 systems overview

The X6 systems portfolio consists of rack-mount servers that include x3850 X6 and x3950 X6 systems, and Flex System compute nodes that include x280, x480, and x880 X6 systems. In this section, we describe the main features and differences of each family.

### X6 rack servers

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

- ▶ Lenovo System x3850 X6
- ▶ Lenovo System x3950 X6

The x3850 X6 server is a 4U rack-optimized server that is scalable to four sockets. The x3950 X6 server is an 8U rack-optimized server that is scalable to eight sockets. These systems are designed for maximum usage, reliability, and performance for compute-intensive and memory-intensive workloads.

Figure 4-1 shows the Lenovo System x3850 X6 server.



Figure 4-1 Lenovo System x3850 X6 server

The x3850 X6 server has the following key characteristics:

- ▶ Up to four Intel Xeon E7-4800 v2, v3, or v4, or E7-8800 v2, v3, or v4 product family processors
- ▶ Up to 96 DIMM slots (24 DIMM slots per processor) for up to 6 TB of memory (by using 64 GB DIMMs)
- ▶ Up to 1600 MHz DDR3 and 1866 MHz DDR4 memory speeds and up to 3200 MHz SMI2 link speeds
- ▶ Up to eight 2.5-inch hot-swap drives or up to 16 1.8-inch hot-swap SSDs
- ▶ Support for 12 Gbps SAS connectivity for internal storage
- ▶ Support for Flash Storage Adapters for high IOPS performance
- ▶ Mezzanine LOM slot for the integrated NIC functionality (choice of dual-port 10 GbE or quad-port 1 GbE adapters)
- ▶ Up to 11 PCIe 3.0 I/O slots
- ▶ Internal USB 2.0 port for the embedded hypervisor

The x3950 X6 server resembles two x3850 X6 servers, where one is placed on top of the other; however, unlike eX5™ servers, x3950 X6 server uses a single-server design with a single backplane without any external connectors and cables.

Figure 4-2 shows the Lenovo System x3950 X6 server.



Figure 4-2 Lenovo System x3950 X6 server

The x3950 X6 server has the following key characteristics:

- ▶ Up to eight Intel Xeon E7-8800 v2, v3, or v4 product family processors
- ▶ Up to 192 DIMM slots (24 DIMM slots per processor) for up to 12 TB of memory (by using 64 GB DIMMs)
- ▶ Up to 1600 MHz DDR3 and 1866 MHz DDR4 memory speeds and up to 3200 MHz SMI2 link speeds
- ▶ Up to 16 2.5-inch hot-swap drives or up to 32 1.8-inch hot-swap SSDs
- ▶ Support for 12 Gbps SAS connectivity for internal storage
- ▶ Support for Flash Storage Adapters for high IOPS performance
- ▶ Two mezzanine LOM slots for the integrated NIC functionality (choice of dual-port 10 GbE or quad-port 1 GbE adapters)
- ▶ Up to 22 PCIe 3.0 I/O slots
- ▶ Two internal USB 2.0 ports for the embedded hypervisors

### **X6 compute nodes**

The Lenovo Flex System X6 portfolio consists of the following compute nodes:

- ▶ Lenovo Flex System x280 X6
- ▶ Lenovo Flex System x480 X6
- ▶ Lenovo Flex System x880 X6

The following systems are based on a two-socket double-wide compute node and differ only by the processor type that is installed and the degree to which the system can scale up:

- ▶ The Flex System x880 X6 Compute Node uses processors of Intel Xeon E7-8800 v2 or v3 family, and it supports the following configurations:
  - 8-socket (four double-wide compute nodes that are connected)
  - 4-socket (two double-wide compute nodes that are connected)
  - 2-socket (one double-wide compute node)
- ▶ The Flex System x480 X6 Compute Node uses processors of Intel Xeon E7-4800 v2 or v3 family, and supports the following configurations:
  - 4-socket (two double-wide computer nodes that are connected)
  - 2-socket (one double-wide compute node)
- ▶ The Flex System x280 X6 Compute Node uses processors of Intel Xeon E7-2800 v2 family and supports only 2-socket configurations as a single double-wide compute node.

Figure 4-3 shows Lenovo Flex System X6 nodes that are connected.



Figure 4-3 *Lenovo Flex Systems that are connected together*

The Flex System X6 compute nodes have the following key characteristics:

- ▶ Up to eight Intel Xeon processor E7-8800 v2 or v3, four E7-4800 v2 or v3, or two E7-2800 v2 product family processors
- ▶ Up to 48 DIMM slots (24 DIMM slots per processor) for up to 3 TB of memory (by using 64 GB DIMMs) per single-wide compute node
- ▶ Up to 1600 MHz DDR3 memory speeds and up to 2667 MHz SMI2 link speeds
- ▶ Up to eight enterprise class 2.5-inch hot-swap drives
- ▶ Up to 16 10 Gb Ethernet interfaces and 16 I/O adapter slots on the x880 node

The x480 is similar to an x880 externally, but only the x880 can scale to a four-node server. The x480 can scale up to 2 x480 nodes.



## 4.3 Lenovo X6 rack server design and architecture

The X6 systems offer the new “bookshelf” design concept that is based on a fixed chassis that is mounted in a standard rack cabinet. You do not need to pull the chassis in or out of the rack because you can access all chassis components from the front or rear, which is similar to pulling books from a bookshelf.

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

- ▶ Compute Book

A Compute Book contains one processor and 24 DIMM slots. It is accessible from the front of the server. For more information about Compute Books, see 4.3.2, “Processor subsystem” on page 52.

- ▶ Storage Book

A Storage Book contains standard 2.5-inch or 1.8-inch hot-swap drive bays. It also provides front USB and video ports, and it has two PCIe slots for internal storage adapters. It is accessible from the front of the server. For more information about Storage Books, see 4.3.5, “Storage subsystem” on page 63.

- ▶ I/O Book

An I/O Book is a container that provides PCIe expansion capabilities. I/O Books are accessible from the rear of the server. For more information about I/O Books, see 4.3.6, “Networking and I/O” on page 64.

The following types of I/O Books are available:

- The *Primary I/O Book* provides core I/O connectivity, including mezzanine LOM slot for an onboard network, three PCIe slots, Integrated Management Module II, and rear ports (USB, video, serial, and management).
- The hot-swap *Full-length I/O Book* provides three optional full-length PCIe slots, and two of the slots can host Graphics Processing Unit (GPU) adapters up to 300 W of total power per book.
- The hot-swap *Half-length I/O Book* provides three optional half-length PCIe slots.

### 4.3.1 Modular components

All modular components are hosted in a 4U (the x3850 X6 server) or 8U (the x3950 X6 server) rack drawer. Two passive backplanes (for 4U and 8U chassis) connect all modular components. Books are the same for 4U and 8U servers.

An x3850 X6 4U server can hosts up to four Compute Books, one Storage Book, one Primary I/O Book, and up to two optional I/O Books. In addition, 4U server supports up to four power supplies and up to 10 hot-swap dual-motor fans (eight fans on the front and two fans on the rear). Figure 4-4 on page 50 and Figure 4-5 on page 50 show front and rear of the x3850 X6 server.




Figure 4-4 x3850 X6 server front view




Figure 4-5 x3850 X6 server rear view

An x3950 X6 8U server can host up to eight Compute Books (four minimum), two Storage Books, two Primary I/O Books, and up to four optional I/O books. In addition, 8U server supports up to eight power supplies and up to 20 hot-swap dual-motor fans (up to 16 fans on the front and four fans on the rear).

Figure 4-6 and Figure 4-7 show front and rear of the x3950 X6 server.

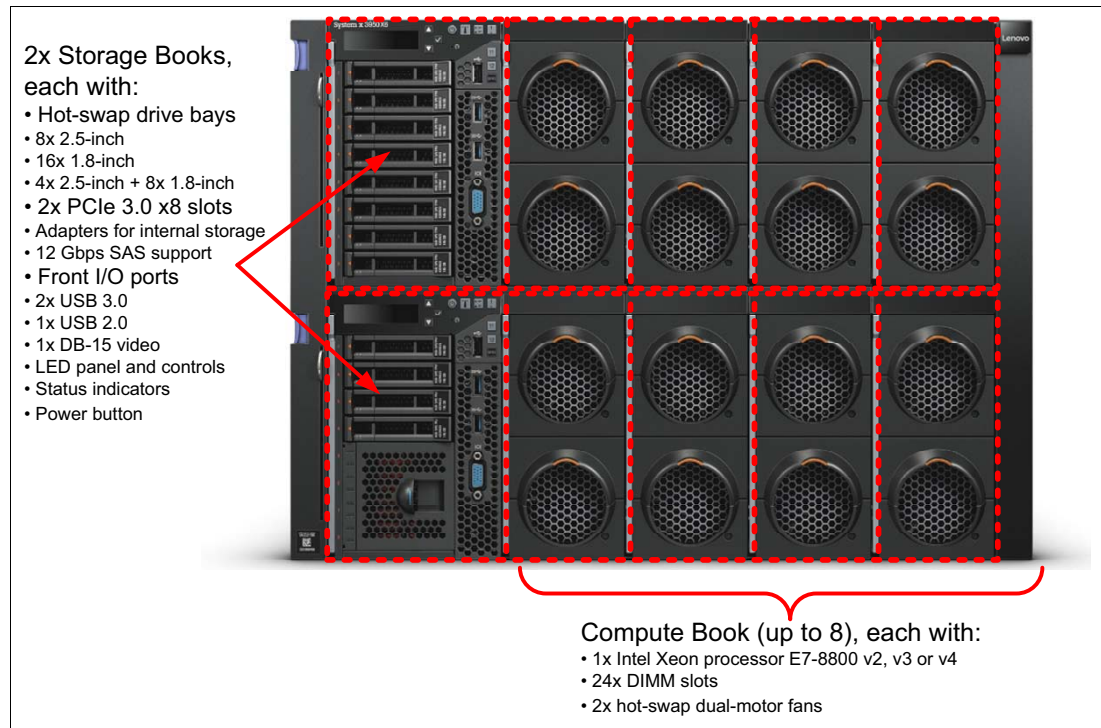


Figure 4-6 x3950 X6 server front view




Figure 4-7 x3950 X6 server rear view

## 4.3.2 Processor subsystem

The current X6 systems use the Intel Xeon processor E7-4800/8800 v2, v3, or v4 product family. The Intel Xeon processors used in the X6 systems are follow-ons to the previous-generation Intel Xeon processor E7-4800/8800 product family. The newest processors feature the Intel microarchitecture (formerly code named *Broadwell-EX*) and the 14 nm manufacturing process that provides higher core count, larger cache sizes, and DDR4 memory interface support.

**Note:** DDR4 memory support is model dependent and requires the Compute Book to have either v3 or v4 processors installed.

Intel Xeon E7 v2, v3 and v4 families support up to 24 DIMMs per processor and provide fast low-latency I/O with integrated PCIe 3.0 controllers.

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

The following groups of the Intel Xeon processor E7 v4 family are used in the X6 servers and support scaling to separate levels:

- ▶ The Intel Xeon processor E7-4800 v4 product family is supported only in the x3850 X6. This family supports four-processor configurations.
- ▶ The Intel Xeon processor E7-8800 v4 product family is supported in the x3950 X6 to scale to eight-socket configurations. It is also supported in the x3850 X6 for interoperability and exchange with the x3950 X6.

The Intel Xeon E7-4800 v4 and E7-8800 v4 product families offer the following key features:

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

- ▶ Resource director technology: Cache monitoring technology, cache allocation technology, memory bandwidth monitoring
- ▶ Security technologies: OS Guard, Secure Key, Intel TXT, Crypto performance (ADOX/ADCX), Malicious Software (SMAP), Key generation (RDSEED)

Table 4-1 compares the Intel Xeon E7-4800/8800 processors that are used in X6 systems.

Table 4-1 X6 processors comparisons

Feature	X6 family, Xeon E7 v2	X6 family, Xeon E7 v3	X6 family, Xeon E7 v4
Processor family	Intel Xeon E7-8800 v2 Intel Xeon E7-4800 v2	Intel Xeon E7-8800 v3 Intel Xeon E7-4800 v3	Intel Xeon E7-8800 v4 Intel Xeon E7-4800 v4
Processor codenames	Ivy Bridge EX	Haswell EX	Broadwell EX
Cores per CPU	Up to 15 cores	Up to 18 cores	Up to 24 cores
Last level cache	Up to 37.5 MB L3 cache	Up to 45 MB L3 cache	Up to 60 MB L3 cache
QPI	QPI 1.1 at 8.0 GT/s max	QPI 1.1 at 9.6 GT/s max	QPI 1.1 at 9.6 GT/s max
CPU TDP rating	Up to 155 W	Up to 165 W	Up to 165 W
DIMM sockets	24 DDR3 DIMMs per CPU	24 DDR3 DIMMs per CPU 24 DDR4 DIMMs per CPU	24 DDR4 DIMMs per CPU
Maximum memory speeds	2667 MHz SMI2	3200 MHz SMI2	3200 MHz SMI2
PCIe technology	PCIe 3.0 (8 GTps)	PCIe 3.0 (8 GTps)	PCIe 3.0 (8 GTps)

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

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

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

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

- ▶ Machine Check Architecture recovery (non-running and running paths)
- ▶ Enhanced Machine Check Architecture Gen2 (eMCA2)
- ▶ Machine Check Architecture I/O
- ▶ Security technologies: OS Guard, Secure Key, Intel TXT

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

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

The Intel Xeon E7-4800 v2 and E7-8800 v2 processors offer the following key features:

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

### **4.3.3 Intel Xeon E7 features**

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

#### **Intel Transactional Synchronization eXtensions**

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

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

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

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

## Intel Advanced Encryption Standard

Advanced Encryption Standard (AES) is an encryption standard that is widely used to protect network traffic and sensitive data. Advanced Encryption Standard - New Instructions (AES-NI), which is available with the E7 processors, implements certain complex and performance intensive steps of the AES algorithm by using processor hardware. AES-NI can accelerate the performance and improve the security of an implementation of AES over an implementation performed by software.

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

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

## Intel Virtualization Technology

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

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

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

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

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

## Hyper-Threading Technology

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

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

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

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

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

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



## Turbo Boost Technology 2.0

The Intel Xeon E7-8800/4800 family of processors brings enhanced capabilities to change processor speed with the new generation of 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 enables a dynamic increase in the clock speed of the active cores to gain a performance boost. For example, a 3.4 GHz 15-core processor can overclock the cores up to 3.7 GHz.

Turbo Boost Technology is available on a per-processor basis for the X6 systems. For ACPI-aware operating systems and hypervisors, such as Microsoft Windows 2008/2012, RHEL 5/6, SLES 11, VMware ESXi 4.1, and later, no changes are required to use it. Turbo Boost Technology can be used with any number of enabled and active cores, which results in increased performance for 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 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 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 temperature, power, and current, and reevaluates whether the current frequency is sustainable or if it must reduce the core frequency further. At any time, all active cores run at the same frequency.

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

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

## QuickPath Interconnect

The Intel Xeon E7 processors that are implemented in X6 servers each include two integrated memory controllers. Processor-to-processor communication is carried over shared-clock or coherent QPI links. Each processor has three QPI links to connect to other processors, as shown in Figure 4-8. On the left side is how the four sockets of the x3850 X6 are connected. On the right side is how all eight sockets of the x3950 X6 are connected.




Figure 4-8 QPI links between processors



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

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

For more information about QPI, see this website:

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

### **Intel Data Direct I/O**

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

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

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

### **RAS features**

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

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

## Machine Check Architecture recovery

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

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

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

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

## Security improvements

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

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

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

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

## Compute Book

The core modular element of the X6 design is a Compute Book, which includes the following components:

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

Figure 4-9 shows the Compute Book.



Figure 4-9 Compute Book

The system board of the Compute Book has two sides on which all components are installed. Figure 4-10 and Figure 4-11 on page 60 show the left and right sides of the Compute Book. The left side contains one processor and 12 DIMM slots. The right side contains 12 DIMM slots, for a total of 24 DIMMs per Compute Book.




Figure 4-10 Left side of Compute Book




Figure 4-11 Right side of Compute Book

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

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

All Compute Books in a server must be the same type; that is, they must have same processors and either all DDR3 or all DDR4 memory.

#### 4.3.4 Memory subsystem

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

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

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

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

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

The x3850 X6 supports up to a total of 96 DIMMs when all processors are installed (24 DIMMs per processor), and the x3950 X6 supports up to a total of 192 DIMMs. The processor and the corresponding memory DIMM slots are contained within Compute Books.

#### Memory operational modes

For X6 and Flex System, each processor has two integrated memory controllers, and each memory controller has two Scalable Memory Interconnect generation 2 (SMI2) links that are connected to two scalable memory buffers. Each memory buffer has two memory channels, and each channel supports three DIMMs, for a total of 24 DIMMs per processors.

Figure 4-12 shows the processor's memory architecture.

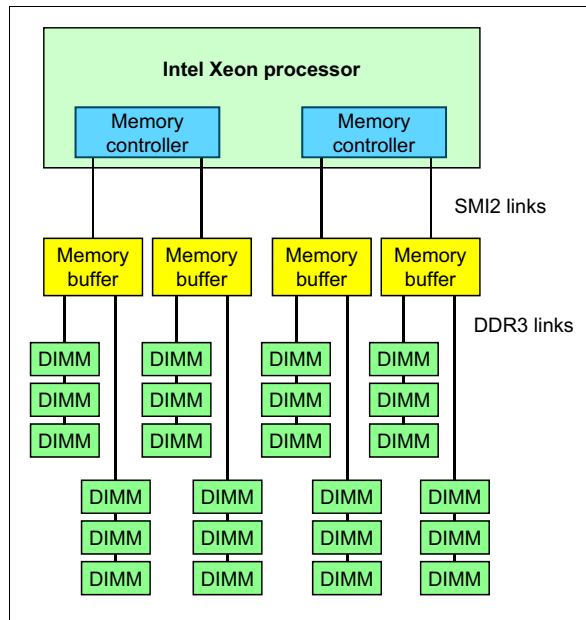


Figure 4-12 Intel Xeon processor E7-xxxx v2/v3 memory architecture

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

► Performance mode

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

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

Although in this mode DIMMs can be populated in any order, memory modules should be placed based on round robin algorithm between SMI2 channels and alternating between DDR channels for best performance.

► RAS (Lockstep) mode

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

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

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

- Chipkill, also known as SDDC
- RBS

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

Figure 4-13 shows the two modes. In RAS mode, both channels of one memory buffer are in lockstep with each other.

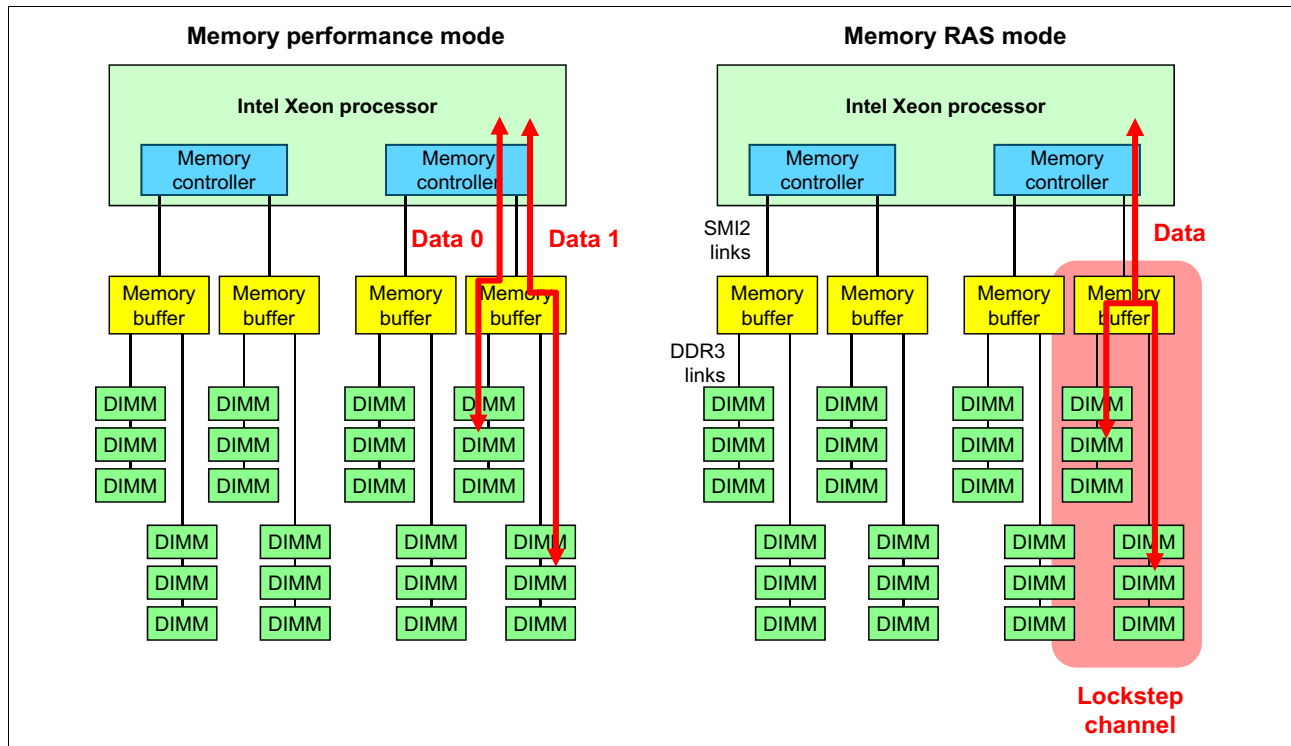


Figure 4-13 Memory modes: Performance mode (left) and RAS mode (right)

The following memory protection technologies also are supported:

- ▶ Error-correcting code memory (ECC)  
Used for detecting and correcting corrupted data.
- ▶ Single Device Data Correction (Chipkill)  
Advanced detection and correction of corrupted data. Single Device Data Correction (SDDC) is supported in performance mode.
- ▶ Double Device Data Correction (Chipkill + Redundant Bit Steering)  
If a memory chip fails, a spare memory chip is used to replace the failed chip. Double Device Data Correction (DDDC) is supported in RAS mode with paired DIMMs.
- ▶ Memory mirroring  
If a fault occurs in a memory channel, the memory controller shifts the data to another channel without disruption. This configuration is similar to a RAID 1 array in disk storage. If memory mirroring is used, DIMMs must be installed in pairs of the same type and size.
- ▶ Memory rank sparing  
Disables failed memory and activates a rank sparing DIMM to replace the failed active DIMM. If memory rank sparing is used, a minimum of two single-rank or dual-rank DIMMs must be installed per populated channel (the DIMMs do not need to be identical). In rank sparing mode, one rank of a DIMM in each populated channel is reserved as spare memory.

### 4.3.5 Storage subsystem

The x3850 X6 and x3950 X6 servers support the following types of internal storage:

- ▶ HDD storage: 2.5-inch SAS/SATA hot-swap HDDs
- ▶ SSD storage:
  - SATA 1.8-inch SSDs
  - 2.5-inch SAS/SATA SSDs
  - 2.5-inch NVMe Enterprise Performance PCIe SSD
- ▶ PCIe storage (Flash Storage Adapters):
  - P3700 NVMe Enterprise Performance Flash Adapter
  - io3 Enterprise Mainstream Flash Adapters
  - Enterprise io3 PCIe Flash Adapter

The following section describes the Storage Book with HDD and SSD storage for the x3850/x3950 X6 servers. For more information about flash and PCIe internal storage offerings, see 4.5, “Lenovo Flash and SSD internal storage” on page 74.

#### Storage Book

The 1.8-inch SSDs and 2.5-inch SSDs and HDDs are installed in Storage Book that is shown in Figure 4-14.




Figure 4-14 The X6 Storage Book

- ▶ The Storage Book supports up to eight HDDs (2.5-inch format), up to 16 1.8-inch SSDs, and up to 4 NVMe PCIe SSDs (2.5-inch format) in the following combinations:
  - ▶ 4x 2.5-inch HDD hot-swap drives
  - ▶ 8x 2.5-inch HDD hot-swap drives
  - ▶ 8x 1.8-inch hot-swap SSDs
  - ▶ 16x 1.8-inch hot-swap SSDs
  - ▶ 4x 2.5-inch HDD hot-swap drives + 8x 1.8-inch hot-swap SSD bays
  - ▶ 4x 2.5-inch PCIe NVMe SSDs
  - ▶ 2x 2.5-inch PCIe NVMe SSDs + 4x 2.5-inch HDD hot-swap drives
  - ▶ 2x 2.5-inch PCIe NVMe SSDs + 8x 2.5-inch HDD hot-swap drives

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



### **12 Gbps SAS internal storage**

The X6 servers support 12 Gb SAS connectivity for internal storage. The 12 Gb SAS doubles the data transfer rate that is compared to 6 Gb SAS solutions to fully unlock the potential of the PCIe 3.0 interface and to maximize performance for storage I/O-intensive applications, including databases, business analytics, and virtualization and cloud environment.

Lenovo 12 Gb SAS controllers provide performance benefits even for the 6 Gbps SAS drive infrastructure (especially for SSDs) by providing intelligent buffering capabilities. The 12 Gb SAS solution consists of Lenovo ServeRAID™ M5210 RAID controllers, performance optimized N2215 SAS/SATA HBAs, and specialized 12 Gb SAS drive backplanes.

### **1.8-inch SATA SSDs**

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 SSDs deliver up to 30,000 IOPS per single drive. These drives can deliver up to 240,000 IOPS and up to 2 GBps of sustained read throughput per SSD unit. In addition to its superior performance, the SSDs offers superior uptime with three times the reliability of mechanical HDDs, because SSDs have no moving parts to fail.

### **NVMe Enterprise Performance PCIe SSD**

Non-Volatile Memory Express (NVMe) is new PCIe 3.0 high performance SSD technology, which provides high I/O throughput and low latency. NVMe interface removes SAS/SATA bottlenecks and unleashes all of the capabilities of contemporary NAND flash memory. Each NVMe SSD drive has direct PCIe 3.0 x4 connections, which provide at least *twice* the bandwidth and *half* the latency of SATA/SAS-based SSD solutions. The NVMe SSD drives are optimized for heavy multithreaded workloads by using internal parallelism and many other improvements, such as enlarged I/O queues.

The NVMe SSDs are available in a 2.5-inch drive form-factor that is compatible with the X6 Storage Book; however, it requires a special NVMe backplane and NVMe PCIe extenders that are installed in Storage Book PCIe 3.0 x8 slots.

## **4.3.6 Networking and I/O**

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

PCIe 3.0 uses 128b/130b encoding scheme, which is more efficient than 8b/10b encoding that is used in PCIe 2.0 protocol. This approach reduces overhead to less than 2% comparing to 20% of PCIe 2.0, and doubles bandwidth at 8 GTps speed.

Each Intel Xeon E7-4800/8800 processor product family contains an Integrated I/O (IIO) module that provides 32 lanes of PCIe 3.0. These 32 lanes can be split into any combination of x4, x8, and x16.

The x3850 X6 server supports one mezzanine LOM slot (for an ML2 card) and up to 11 PCIe slots using up to three I/O Books, plus one Storage Book installed in the 4-socket chassis. The x3950 X6 server supports two mezzanine LOM slots and up to 22 PCIe slots with up to six I/O Books and two Storage Books installed in the 8-socket chassis.



The Primary I/O Book supplies three PCIe 3.0 slots and one mezzanine LOM slot. The following PCIe slots are available:

- ▶ 2x PCIe 3.0 x16 (x16-wired); full height, half length
- ▶ 1x PCIe 3.0 x16 (x8-wired); full height, half length

Mezzanine LOM slot (PCIe 3.0 x8 interface) is used for mezzanine network cards that provide the flexibility in choosing integrated NIC option, including:

- ▶ 2x 10 Gb Ethernet SFP+ ports
- ▶ 2x 10 Gb Ethernet RJ-45 ports
- ▶ 4x 1 Gb Ethernet RJ-45 ports

In addition, one port on the ML2 card can be configured as a shared management or data port.

Figure 4-15 shows the Primary I/O Book location and its components.




Figure 4-15 Primary I/O Book location and ports

In addition to PCIe expansion, Primary I/O Book provides I/O ports on the rear of the server. Primary I/O Book requires processors 1 and 2 in the x3850 X6 servers, and processors 1, 2, 5, and 6 in the x3950 X6 servers (two Primary I/O Books).

The following types of I/O Books are optional:

- ▶ Full-length
- ▶ Half-length

Full-length I/O Book supports high performance GPU adapters (up to 300 W) and other full-length cards, and Half-length I/O Book supports traditional PCIe adapters. Up to two optional I/O Books can be installed in the x3850 X6 server, and they require processors 3 and 4. Up to four optional I/O Books can be installed in the x3950 X6 server, and they require processors 3, 4, 7, and 8. Optional I/O Books are enabled for PCIe hot add and hot replace functionality.

Figure 4-16 shows the two optional I/O Books.



Figure 4-16 Optional I/O Books

The following I/O Book slots are available:

- ▶ Half-length I/O Book:
  - Two PCIe 3.0 x8 slots (x8 wired)
  - One PCIe 3.0 x16 slot (x16 wired)
- ▶ Full-length I/O Book:
  - Two PCIe 3.0 x16 (x16 wired)
  - One PCIe 2.0 x8 slot (x4 wired)

Both I/O Books accommodate full height adapters.

### 4.3.7 Scalability

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

## 4.4 Lenovo Flex System X6 compute node architecture

This section describes the Lenovo Flex System chassis and provides more information about the Lenovo Flex System X6 compute nodes (x280 X6/x480 X6/x880 X6).

## 4.4.1 Lenovo Flex System Enterprise Chassis

Figure 4-17 shows the front and rear views of the Lenovo Flex System Enterprise Chassis.



Figure 4-17 *Lenovo Flex System Enterprise Chassis: Front and rear*

The chassis provides 14 bays for half-wide nodes, four scalable I/O switch modules, and two Chassis Management Modules (CMMs). Current node configurations include half-wide and full-wide, and double-high and quad-high options. Power and cooling can be scaled up as needed in a modular fashion as more nodes are added.

For more information, see *Lenovo Flex System Enterprise Chassis Product Guide*, TIPS0863:  
<http://lenovopress.com/tips0863>

## 4.4.2 Lenovo Flex System X6 compute node architecture

The Lenovo Flex System X6 compute node has the following system architecture features:

- ▶ Two 2011-pin type R (LGA-2011) processor sockets.
- ▶ An Intel C602J PCIe Controller Hub (PCH).
- ▶ Two integrated memory controllers. Each memory controller provides two Intel SMI2 channels that interface with a Scalable Memory Buffers.
- ▶ Two memory channels per Memory Buffer.
- ▶ Up to three DIMMs per memory channel.
- ▶ A total of 48 DDR3 DIMM sockets.
- ▶ Support for LRDIMMs and RDIMMs.
- ▶ One LSI 3004 SAS controller with integrated RAID 0, 1, 10, and 1E to the two 2.5-inch internal drive bays.
- ▶ Integrated Management Module II (IMM2) for systems management.
- ▶ Four PCIe 3.0 I/O adapter connectors: two connectors have two x8 links (x16 total) interface, two other connectors have x8 interfaces.
- ▶ Two internal USB 2.0 ports for a hypervisor key and one external USB 3.0 connector for external USB devices.

Figure 4-18 shows the components on the system board of the X6 compute node.




Figure 4-18 Layout of the X6 compute node system board

The X6 compute node features the following key design aspects:

- ▶ The system uses scalable memory buffers to provide a total of eight memory channels per processor and three DIMMs per channel, for a total of 48 DIMMs per 2-socket compute node.
- ▶ The two processors in a compute node are connected via one QPI link. The other two QPI links are used to connect to other compute nodes in a scalable complex.
- ▶ I/O connectors 1 and 2 are connected to the first processor; I/O connectors 3 and 4 are connected to the second processor

### 4.4.3 Processor subsystem

The Lenovo Flex System X6 compute nodes feature one of the following Intel Xeon E7 v2 or v3 series processors:

- ▶ Intel Xeon Processor E7-2800 v2: Supports 2-socket configurations
- ▶ Intel Xeon Processor E7-4800 v2 or v3: Supports 2 and 4-socket configurations
- ▶ Intel Xeon Processor E7-8800 v2 or v3: Supports 2, 4, and 8-socket configurations

For more information about the Intel Xeon processor E7-4800 and E7-8800 V2 and V3 product family, see 4.3.2, “Processor subsystem” on page 52.



#### 4.4.4 Memory subsystem

The Lenovo Flex System X6 compute nodes support DDR3 memory with ECC protection. The X6 supports up to 48 DIMMs when two processors are installed (24 DIMMs per processor), and up to 192 DIMMs with 8-socket scaled configuration.

For more information about memory operational modes that are supported in the Flex system, see “Memory operational modes” on page 60.

#### 4.4.5 Storage subsystem

The X6 compute node has two 2.5-inch front-accessible hot-swap drive bays that are accessible from the front of the server. These bays are connected to the integrated ServeRAID M1210e controller.

The ServeRAID M1210e controller includes the following features:

- ▶ Based on the LSI SAS 3004 12 Gbps SAS/SATA RAID-on-Chip (ROC) controller
- ▶ Four-port controller with 12 Gbps throughput per port
- ▶ PCIe x4 Gen 2 host interface
- ▶ Supports RAID levels 0, 1, 10, and 1E; optionally supports RAID 5 and RAID 50.

The two 2.5-inch front-accessible drive bays can be replaced with four 1.8-inch drive bays, two 1.8-inch bays replacing each 2.5-inch bay, by using the ServeRAID M1200 Series Flex System Flash Kit for x880 X6.

#### 4.4.6 I/O subsystem

The X6 node has four I/O expansion connectors for attaching I/O adapters, as shown in Figure 4-19 on page 69. Installing I/O adapters allows the server to connect to switch modules in the Lenovo Flex System Enterprise Chassis.




Figure 4-19 Location of the I/O adapters in the Lenovo Flex System X6 compute node

A compatible switch or pass-through module must be installed in the corresponding I/O bays in the chassis, as listed in Table 4-2. Installing two switches means that all ports of the adapter are enabled, which improves performance and network availability.

Table 4-2 Adapter to I/O bay correspondence

I/O adapter slot in the X6 compute nodes	Port on the adapter	Corresponding I/O module bay in the chassis
Slot 1	Port 1	Module bay 1
	Port 2	Module bay 2
	Port 3 (for 4 & 8-port cards)	Module bay 1
	Port 4 (for 4 & 8-port cards)	Module bay 2
	Port 5 (for 8-port cards)	Module bay 1
	Port 6 (for 8-port cards)	Module bay 2
	Port 7 (for 8-port cards)	Module bay 1
	Port 8 (for 8-port cards)	Module bay 2
Slot 2	Port 1	Module bay 3
	Port 2	Module bay 4
	Port 3 (for 4 & 8-port cards)	Module bay 3
	Port 4 (for 4 & 8-port cards)	Module bay 4
	Port 5 (for 8-port cards)	Module bay 3
	Port 6 (for 8-port cards)	Module bay 4
	Port 7 (for 8-port cards)	Module bay 3
	Port 8 (for 8-port cards)	Module bay 4
Slot 3 (Dual-ASIC adapters not supported)	Port 1	Module bay 1
	Port 2	Module bay 2
Slot 4 (Dual-ASIC adapters not supported)	Port 1	Module bay 3
	Port 2	Module bay 4

Figure 4-20 shows the location of the switch bays in the rear of the Enterprise Chassis.




Figure 4-20 Locations of the I/O modules

Figure 4-21 shows how the two port adapters are connected to switches that are installed in the I/O Module bays in an Enterprise Chassis.

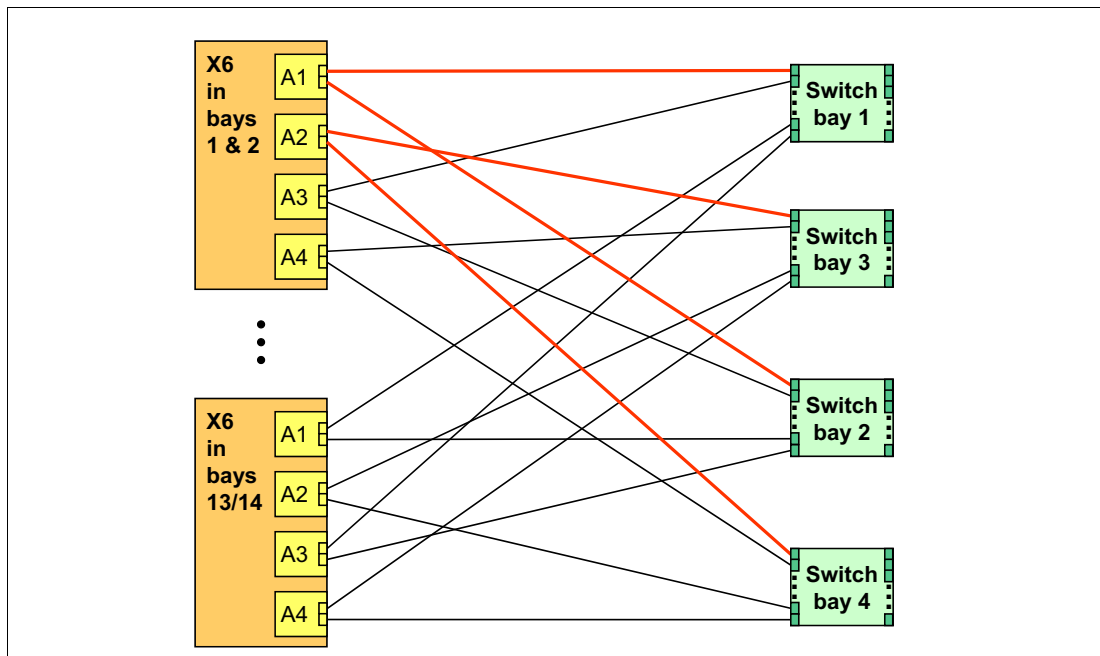


Figure 4-21 Logical layout of the interconnects between I/O adapters and I/O module

## 4.4.7 Scalability

The X6 compute nodes can scale to a 4-socket or 8-socket complex, depending on the processor that is installed into the node. Consider the following points:

- ▶ An x880 X6 node with Intel Xeon E7-8800 v3 family processors can scale up to an 8-socket configuration, and supports 2-socket, 4-socket, and 8-socket configurations. All processors in these configurations must be identical.
- ▶ An x480 X6 node with Intel Xeon E7-4800 v3 family processors can scale up to a 4-socket configuration, and supports 2-socket and 4-socket configurations. All processors in these configurations must be identical.
- ▶ An x280 X6 node with Intel Xeon E7-2800 v2 family processors supports only 2-socket configurations. The processors in these configurations must be identical.

The scaled X6 compute nodes are connected through a front interconnect system. The interconnection has a QPI bus and the sideband signals that are needed for correct operation.

The Lenovo Flex System x880 X6 4-Socket Scalability Connector is shown in Figure 4-22.



Figure 4-22 Lenovo Flex System x880 X6 4-Socket Scalability Connector

Lenovo Flex System x880 X6 4-Socket Scalability Connector allows for two x480 X6 or x880 X6 nodes to form a dual-node four-socket scaling complex. The scalability connection effectively connects all processors via their QPI links, as shown in Figure 4-23.

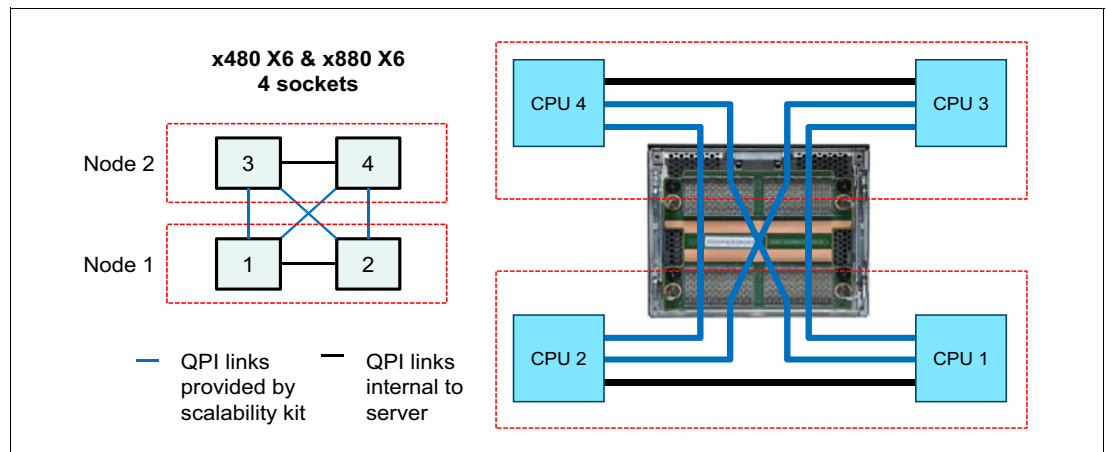


Figure 4-23 4-Socket Scalability Connector: QPI connections



The Lenovo Flex System x880 X6 8-Socket Scalability Connector is shown in Figure 4-24.




Figure 4-24 Lenovo Flex System x880 X6 8-Socket Scalability Connector

Lenovo Flex System x880 X6 8-Socket Scalability Connector allows for four x880 X6 nodes to form a 4-node, 8-socket scaling complex. Figure 4-25 shows the block diagram for this connector.

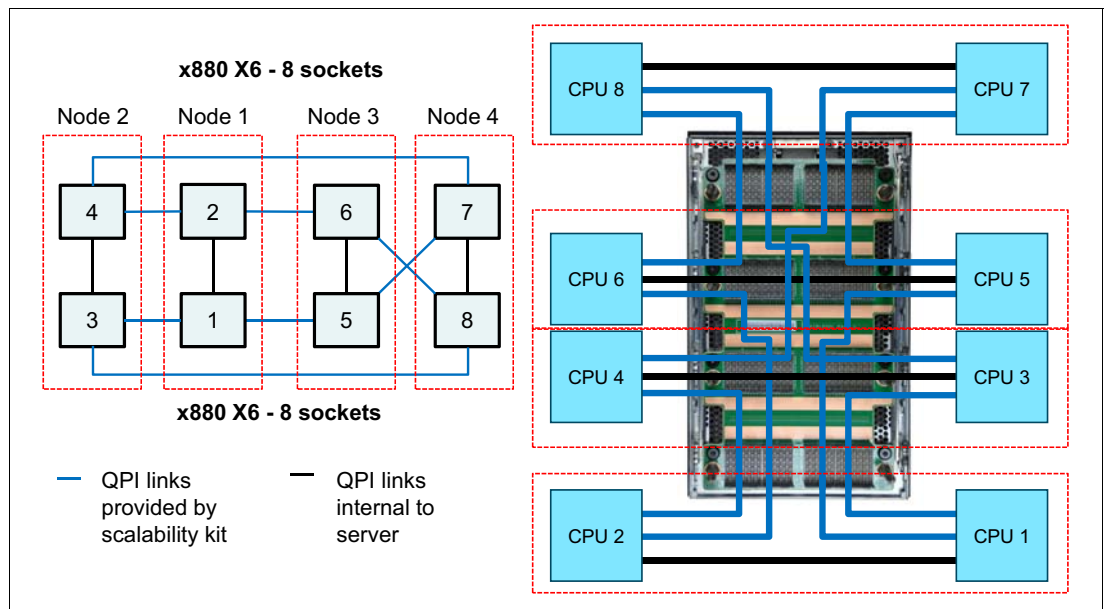


Figure 4-25 Block diagram for 8-socket Scalability Connector

## 4.4.8 Lenovo Flex System X6 compute nodes comparison

Table 4-3 shows a high-level comparison between the 2-socket x280 X6, the 4-socket x480 X6, and the 8-socket x880 X6, based on CPU, cores, memory, drives, interfaces, and I/O.

Table 4-3 High-Level Flex System X6 compute node comparison

Configurations		x280 X6	x480 X6	x880 X6
Form factor	2-socket	Full-wide (2 node bays)	Full-wide (2 node bays)	Full-wide (2 node bays)
	4-socket	Not available	Full-wide, double-high (4 node bays)	Full-wide, double-high (4 node bays)
	8-socket	Not available	Not available	
Maximum processors / cores	1-node	2 / 30	2 / 30	2 / 30
	2-node	Not available	4 / 60	4 / 60
	4-node	Not available	Not available	8 / 120
Maximum memory DIMM slots	1-node	48	48	48
	2-node	Not available	96	96
	3-node	Not available	Not available	192
2.5-inch drive bays	1-node	2	2	2
	2-node	Not available	4	4
	4-node	Not available	Not available	8
I/O adapter slots	1-node	4	4	4
	2-node	Not available	8	8
	4-node	Not available	Not available	16

## 4.5 Lenovo Flash and SSD internal storage

Currently, the processor, memory, and I/O subsystem are well-balanced and not considered as performance bottlenecks in most systems. The major source of performance issues is related to the storage I/O activity because of the speed of traditional HDD-based storage systems that still does not match the processing capabilities of the servers. This issue can lead to a situation when a powerful processor sits idle while waiting for the storage I/O requests to complete, which wastes its time and negatively affects user productivity, extends return on investments (ROI) time frame, and increases overall total cost of ownership (TCO).

Lenovo Flash storage offerings can help to address these issues by combining extreme IOPS performance with low response time. Lenovo X6 servers use flash storage offerings and SSDs to achieve fastest response time for analytical workloads, transactional databases, and virtualized environments.

X6 servers include the following Lenovo SSD and Flash Storage Adapter technologies, which are discussed in the following sections:

- ▶ 4.5.1, “SAS/SATA SSD technology” on page 75
- ▶ 4.5.2, “NVMe SSD technology” on page 75
- ▶ 4.5.3, “Flash Storage Adapters” on page 77

Table 4-4 compares the features of the SSDs and Flash Storage devices that are used in X6 servers.

Table 4-4 *Lenovo Flash storage and SSD devices*

Feature	io3 Enterprise Mainstream Flash Adapter	NVMe Enterprise Performance Flash Adapter	NVMe Enterprise Performance PCIe SSD	2.5-inch SSDs	Enterprise Value 1.8-inch SSDs
Form factor	PCIe adapter	PCIe adapter	2.5-inch drive	2.5-inch drive	1.8-inch drive
Interface	PCIe 2.0 x8	PCI 3.0 x4	PCIe 3.0 x4	6 or 12 Gbps SAS or 6 Gbps SATA	6 Gbps SATA
Capacity	Up to 6.4 TB	2.0 TB	2.0 TB	Up to 1.6 TB	Up to 800 GB
Max. random read IOPS per unit	Up to 340,000	450,000 (90% consistent)	Up to 450,000	Up to 100,000	Up to 65,000
Write latency	15 $\mu$ s	20 $\mu$ s	As low as 20 $\mu$ s	Less than 100 $\mu$ s	150 $\mu$ s
Hot-swappable	No	No	Yes <sup>a</sup>	Yes	Yes
RAID support	Chip-level redundancy; operating system software RAID	No	Operating system software RAID	Yes	Yes

a. Hot add and hot remove features are available with supported operating systems.

### 4.5.1 SAS/SATA SSD technology

SSD solutions can provide the following benefits:

- ▶ Significantly lower implementation cost of high performance I/O-intensive storage systems.
- ▶ Significantly higher performance and better response time for storage-intensive applications, with up to 10x lower response time compared to solutions based on mechanical 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, decision support, and other random I/O-intensive applications. Built on enterprise-grade MLC NAND flash memory, the SSD drives used in the X6 systems deliver up to 30,000 IOPS per single drive. Combined into an SSD unit, these drives can deliver up to 240,000 IOPS and up to 2 GBps of sustained read throughput per SSD unit. In addition to their superior performance, SSDs offer superior uptime with three times the reliability of mechanical disk drives, because SSDs have no moving parts to fail.

### 4.5.2 NVMe SSD technology

Non-Volatile Memory Express (NVMe) is new PCIe 3.0 high performance SSD technology that provides high I/O throughput and low latency for the x3850 X6 and x3950 X6 servers. Similar to the PCIe Flash Adapter (discussed in 4.5.3, “Flash Storage Adapters” on page 77) both of which share the same NVMe technology, NVMe interfaces remove SAS/SATA bottlenecks and enable all of the capabilities of contemporary NAND flash memory. Each

NVMe PCI SSD has direct PCIe 3.0 x4 connection, which provides at least *twice* the bandwidth and *half* the latency of SATA/SAS-based SSD solutions. NVMe drives are also optimized for heavy multithreaded workloads by using internal parallelism, along with many other improvements, such as enlarged I/O queues.

NVMe technology has the following key characteristics:

- ▶ Direct PCIe 3.0 connection. There is a PCIe 3.0 x4 connection with up to 4 GBps overall throughput for each NVMe drive.
- ▶ Low I/O latency. The average read/write latency for the P3700 drives is 20  $\mu$ s.
- ▶ Up to 2800 MBps sequential read speed with 128 KB blocks, and up to 2000 MBps sequential write speed with 128 KB blocks per drive.
- ▶ Up to 460,000 IOPS of random read with 4 KB blocks, and up to 175,000 IOPS of random writes with 4 KB blocks.
- ▶ A total of 65,536 I/O queues supported and 65,536 commands per queue supported, which provides great performance on heavily multithreaded workloads with combined sequential and random access.
- ▶ High endurance: The P3700 drives include High Endurance Technology (HET) which combines NAND silicon enhancements and SSD NAND management techniques to extend SSD write endurance up to 17 drive writes per day (DWPD) for 5 years.
- ▶ Available drive capacities of 400 GB, 800 GB, 1.6 TB, and 2.0 TB.
- ▶ Support for software RAID under operating system management.
- ▶ Hot add and hot remove features are available with supported operating systems.
- ▶ Most operating systems have native support of NVMe drives or provides support through software drivers, such as RHEL 6.5 and later, RHEL 7.0 and later, SLES 11 SP3, SLES 12, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, and VMware vSphere (ESXi) 5.5 and 6.0.
- ▶ NVMe drives can be used as boot drives.

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

Figure 4-26 shows the PCIe NVMe SSD.



Figure 4-26 NVMe PCIe SSD

For more information about SSD drives that are supported in the X6 servers, see the Lenovo Press Product Guides in the drives category:  
<https://lenovopress.com/servers/options/drives>

### 4.5.3 Flash Storage Adapters

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

The Flash Storage Adapter options include the following:

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

The Flash Storage Adapters combine high IOPS performance with low latency. As an example, with 4 KB block random reads, the 3.2 TB io3 Enterprise Mainstream Flash Adapter can read 345,000 IOPS, and the 2.0 TB P3700 NVMe Enterprise Performance Flash Adapter can read 400,000 IOPS, compared with 420 IOPS for a 15 K RPM 146 GB disk drive. The Enterprise Mainstream adapter can also write 385,000 IOPS with 4 KB random writes while the NVMe Enterprise Performance Flash Adapter can write 150,000 IOPS with 4 KB random writes. The read access latency is approximately 79 microseconds for the Enterprise Mainstream Flash Adapter, which is over 100th of the latency of a 15K RPM 146 GB disk drive. The write access latency is even less, with about 15 microseconds.

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

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

Figure 4-27 shows the 6.4TB io3 Enterprise Mainstream Flash Storage Adapter:




Figure 4-27 6.4 TB io3 Enterprise Mainstream Flash Storage Adapter

Figure 4-28 shows the P3700 NVMe Enterprise Performance Flash Adapter:



Figure 4-28 P3700 Enterprise Performance Flash Adapter

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

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

#### 4.5.4 Lenovo X6 RAS features

The ultimate design goal for Lenovo X6 systems as a part of the Enterprise X-Architecture strategy is to prevent hardware faults from causing an outage. Part selection for reliability, redundancy, recovery, self-healing techniques, and degraded operational modes are used in a fault resilient RAS strategy to avoid application outages.

#### 4.5.5 Hardware fault tolerance

The X6 systems offers several hardware-based fault tolerant features:

- ▶ Redundant hot-swap power supplies and fans
- ▶ Redundant hot-swap disk and solid-state drives with the RAID adapters
- ▶ I/O adapter fault-tolerance and hot-replace capabilities
- ▶ Integrated Management Module II (IMM2.1)
- ▶ Predictive Failure Analysis
- ▶ Light path diagnostics

IMM2 is a part of the RAS strategy. It supports error prevention with Predictive Failure Analysis (PFA), server self-healing, concurrent maintenance, and light path diagnostics for minimizing servicing downtime. In addition, IMM provides self-healing capabilities for the system by isolating faults and restarting a server in degraded mode (for example, if a processor failure occurred, the IMM disables it and restarts the server by using the remaining processors, if possible).

PFA for Lenovo System x servers is a collection of techniques that help you run your business with less unscheduled downtime. PFA allows the server to monitor the status of critical subsystems and notify the system administrator when components appear to be degrading. In most cases, replacement of failing parts can be performed as part of planned maintenance activity. As a result, unscheduled outages can be prevented.

Providing peace of mind, Lenovo stands behind its server products with a 3-year, onsite limited warranty. Labor and Lenovo parts are covered for the full duration of the warranty period, including parts that are identified during PFA.

PFA features on the X6 servers include:

- ▶ System memory
- ▶ System fans
- ▶ Processors
- ▶ Power supplies
- ▶ Voltage regulators on the system board
- ▶ Hard disk drives

Each feature provides local and remote warnings of impending failures that might cause unscheduled downtime. The potential benefit is higher server availability and reduced TCO.

Light path diagnostics allows systems engineers and administrators to easily and quickly diagnose hardware problems on the Lenovo System x servers. Hardware failures that in the past took hours to locate and diagnose can be detected in seconds, which avoids or reduces downtime.

Light path diagnostics constantly monitors selected components within the System x server. If a failure occurs, a light is illuminated on the front panel of the server to alert the systems administrator that there is a problem. An LCD panel on the front of the server indicates the failed component.

Light path diagnostics is performed at a hardware level, and it does not require an operating system to function. These characteristics make it reliable and allow you to diagnose problems before an operating system is even installed.

#### 4.5.6 Intel Xeon processor RAS features

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

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



## 4.5.7 Machine Check Architecture recovery

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

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

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

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

Figure 4-29 on page 80 shows an example of how an application can use the MCA recovery feature to handle system errors to prevent itself from being ended if there is a system error.

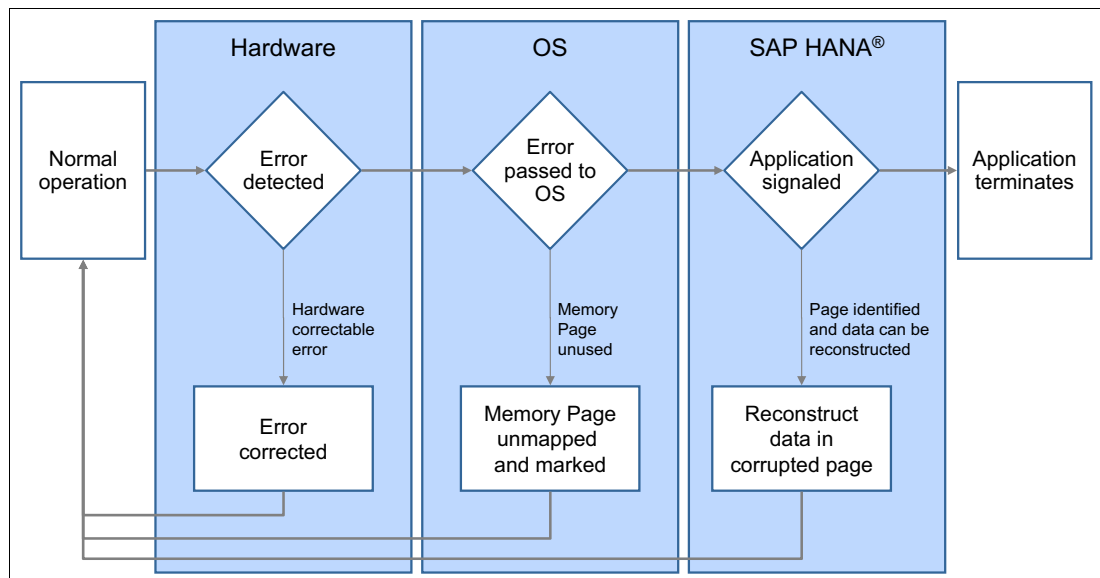


Figure 4-29 Intel Machine Check Architecture recovery example

If a memory error is encountered that cannot be corrected by the hardware, the processor sends an MCA recovery signal to the operating system. An operating system that supports MCA recovery now determines whether the affected memory page is in use by an application. If unused, it unmaps the memory page and marks it as in error. If the page is used by an



application, the operating system often holds that application or stops all processing and halts the system.

With the MCA-aware application, the operating system can signal the error situation to the application, which gives it the chance to try to repair the effects of the memory error. By using MCA recovery, the application can take an appropriate action at the level of its own data structures to ensure a smooth return to normal operation, which avoids a time-consuming restart or loss of information.

The Intel Xeon processor E7-2800/4800/8800 product families include the following new MCA recovery features:

- ▶ Execution path recovery: Ability to work with hardware and software to recognize and isolate the errors that were delivered to the execution engine (core).
- ▶ MCA I/O: Ability to report uncorrectable (fatal and non-fatal) errors to the software layers for further handling, such as determining the root cause of failure or preventive maintenance.
- ▶ Enhanced MCA (eMCA) Gen 1: Provides enhanced error log information to the operating system, hypervisor, or application that can be used to provide better diagnostic and predictive failure analysis for the system. This feature enables higher levels of uptime and reduced service costs.

## 4.5.8 Memory RAS features

You can enable various memory RAS features from the Unified Extensible Firmware Interface (UEFI) shell. These settings can increase the reliability of the system; however, there might be performance trade-offs when these features are enabled. The following sections provide a brief description of each memory RAS technology.

### Lenovo Chipkill

Chipkill memory technology, an advanced form of error correction code (ECC) from Lenovo, is available for the X6 servers. Chipkill (also known as SDDC) protects the memory in the system from any single memory chip failure. It also protects against multibit errors from any portion of a single memory chip. Chipkill on its own can provide 99.94% memory availability<sup>1</sup> to the applications without sacrificing performance and with standard ECC DIMMs.

### Redundant bit steering and double device data correction

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

Within the system, the models of the X6 servers that use the Intel Xeon processor E7 family support the Intel implementation of Chipkill plus redundant bit steering, which Intel refers to as *double device data correction* (DDDC). Redundant bit steering uses the ECC coding scheme that provides Chipkill coverage for x4 DRAMs. This coding scheme leaves the equivalent of one x4 DRAM spare in every pair of DIMMs. If a chip failure on the DIMM is detected, the memory controller can reconstruct data from the failed chip to the x4 spare.

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

<sup>1</sup> A White Paper on the Benefits of Chipkill-Correct ECC for PC Server Main Memory by Timothy Dell, IBM Microelectronics Division - Rev. 11/19/97

## Lenovo Advanced Page Retire

Advanced Page Retire is a Lenovo unique algorithm to handle memory errors. It is a built-in, sophisticated error handling firmware that uses and coordinates memory recovery features by balancing the goals of maximum up time and minimum repair actions. The algorithm uses short- and long-term thresholds per memory rank with leaky bucket and automatic sorting of memory pages with the highest correctable error counts.

It also uses hardware recovery features, followed by software recovery features, to optimize recovery results for newer and older operating systems and hypervisors. When recovery features are exhausted, the firmware issues a PFA. Memory that failed completely is held offline during restarts until repaired. Failed DIMMs are indicated by light path diagnostics LEDs that are physically at the socket location.

Lenovo thoroughly tests to verify the features and co-ordination between the firmware and the operating system or hypervisor.

## Memory mirroring

To improve memory reliability and availability beyond ECC and Chipkill (see “Lenovo Chipkill” on page 81), the memory controller can mirror memory data to two memory channels. To enable mirroring, you must have both memory channels populated with the same DIMM type and amount of memory.

Memory mirroring, or *full array memory mirroring (FAMM) redundancy*, provides a redundant copy of all code and data addressable in the configured memory map. Memory mirroring works within the chip set by writing data to two memory ports on every memory-write cycle. Two copies of the data are kept, similar to the way RAID 1 writes to disk. Reads are interleaved between memory ports. The system automatically uses the most reliable memory port as determined by error logging and monitoring.

If errors occur, only the alternative memory port is used until bad memory is replaced. Because a redundant copy is kept, mirroring results in only half the installed memory being available to the operating system. FAMM does not support asymmetrical memory configurations and requires that each port is populated in identical fashion. For example, you must install 2 GB of identical memory equally and symmetrically across the two memory ports to achieve 1 GB of mirrored memory.

Memory mirroring is independent of the operating system. There is a memory performance trade-off when memory mirroring is enabled.

## Memory rank sparing

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

In rank sparing mode, one rank per memory channel is configured as a spare. The ranks must be as large as the rank relative to the highest capacity DIMM for which sparing is set up. The size of the two unused ranks for sparing is subtracted from the usable capacity that is presented to the operating system.

Memory sparing is independent of the operating system. There is a memory performance trade-off when memory sparing is enabled.

## 4.5.9 Lenovo XClarity Administrator

Lenovo XClarity Administrator is a centralized resource management solution that is aimed at reducing complexity, speeding response, and enhancing the availability of Lenovo server systems and solutions.

XClarity Administrator provides agent-free hardware management for System x rack servers, including x3850 X6 and x3950 X6 servers, and Flex System compute nodes and components, including the CMM and Flex System I/O modules.

XClarity Administrator is a virtual appliance that is quickly imported into a Microsoft Hyper-V or VMware virtualized environment, which gives easy deployment and portability. The tool offers out-of-band agentless management to reduce complexity, which means that the endpoints that are managed do not need special software agents or drivers to be installed or maintained.

XClarity Administrator offers the following functions:

- ▶ Discovery
- ▶ Inventory
- ▶ Monitoring
- ▶ Firmware updates
- ▶ Firmware compliance
- ▶ Configuration management
- ▶ Deployment of operating systems and hypervisors to bare metal servers

The administration dashboard is based on HTML 5. Such a foundation allows fast location of resources so tasks can be run quickly.

For more information about XClarity, see *Lenovo XClarity Administrator*, TIPS1200, which is available at this website:

<http://lenovopress.com/tips1200>

## 4.5.10 Lenovo XClarity Energy Manager

Lenovo XClarity Energy Manager (LXEM) is a standalone software tool 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, LXEM helps data center administrator improve business continuity and energy efficiency.

Features and functions of the LXEM include the following:

- ▶ Monitoring energy consumption, estimating the power demand, and reallocating power to servers as needed.
- ▶ An 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).

For more information on Lenovo XClarity Energy Manager, refer to the following support website: <http://support.lenovo.com/us/en/downloads/ds101160>

### 4.5.11 Lenovo XClarity Integrators

Lenovo XClarity supports integration with VMware vSphere and Microsoft System Center by using XClarity Integrators. The upward integration provides automation and resiliency in virtual environments. Each integrator integrates hardware PFA and microcode management and diagnostics into standard hypervisors, which provides the following capabilities:

- ▶ Manage resources from virtualization console
- ▶ Non-disruptive server updates
- ▶ Non-disruptive server restarts
- ▶ Evacuate workloads on predicted hardware failure

The Lenovo XClarity Integrator plug-ins are available to the following virtualization platforms:

- ▶ Lenovo XClarity Integrator for VMware vCenter; Lenovo customized ESXi 5.0 (U1, U2, and U3), 5.1 (U1, U2, and U3), 5.5 (U1, U2, and U3) and 6.0 images
- ▶ Lenovo XClarity Integrator for Microsoft System Center 2012, 2012 R2, or later

### 4.5.12 Upward Integration Module for Zenoss

Lenovo System x Upward Integration Modules (UIM) for Zenoss extends the Zenoss Resource Manager heterogeneous monitoring tool with System x and Flex System infrastructure inventory information and events.

The solution helps IT Operations transform from an endpoint management focus to a service-oriented focus by automatically discovering and modeling heterogeneous hardware and software resources in the environment, and immediately creating an asset-service relationship context. This end-to-end visualization helps IT Operations isolate and address problems faster, and reduce mean time to resolution (MTTR).

UIM includes the following key features:

- ▶ Agent-less software plug-ins that are installed on the Zenoss Resource Manager server
- ▶ Automatically discovers System x servers and Flex System in the environment, and makes them available for visualization and monitoring from the Zenoss Resource Monitor graphical user interface (GUI)
- ▶ Visibility into System x and Flex System hardware and firmware inventory data from a unified monitoring console
- ▶ Simple Network Management Protocol v3 (SNMPv3) based inventory collection, availability and health monitoring, and trap handling
- ▶ Faster incident resolution through holistic and precise root cause analysis, with automated error and exceptions handling, such as restarting system components

For more information about the UIM for Zenoss, see this website:

<http://www.zenoss.com/solution/lenovo>

## 4.6 Lenovo Storage offerings

The Lenovo comprehensive approach to SAP system storage needs addresses the requirements for high performance, scalability, maximum security, reliability, and low TCO. Lenovo also recognizes the importance of sharing data across platforms to reduce data duplication, increase data availability, and improve network response time.

SAP solutions include the following Lenovo Storage service offerings:

- ▶ Virtualization solutions from Lenovo to help improve flexibility in the storage infrastructure.
- ▶ FlashCopy solutions that handle large volumes of business-critical data and help SAP customers maintain high application-server availability to help minimize the affect on revenue-generating applications. Working with SAP, Lenovo tested several backup and recovery and high availability scenarios that were designed to maintain data consistency and minimize user downtime.

### 4.6.1 SAP customer requirements

Whether they operate in retail, finance, pharmaceuticals, or many other vertical markets, today's competitive businesses depend on SAP to run many of their core operational processes, such as procurement, sales, finance, customer relationship management, and supply chain management.

Applications, such as SAP ERP, give managers enterprise-wide visibility and control. However, with the increasing complexity of enterprise software systems comes an ever-growing pool of data that must be kept secure and available 24 hours, 7 days a week. SAP addresses this challenge by running on a 2-tier or 3-tier architecture with servers that are running the applications and a dedicated server for the database.

Above all, SAP customers need storage solutions that maximize performance while keeping implementation and maintenance costs down. Lenovo offers the following storage solutions:

- ▶ Lenovo Entry SAN storage offerings:
  - Lenovo Storage S2200
  - Lenovo Storage S3200
- ▶ IBM Storwize offerings for Lenovo:
  - Storwize V3700
  - Storwize V5000
  - Storwize V7000

The potential results for customers include low integration costs and fast solution deployment. Numerous SAP solutions were tested with Lenovo storage solutions, including the full SAP Business Suite and applications, such as SAP ERP and SAP Advanced Planning and Optimization and SAP.

Lenovo and SAP have a global partnership. For more information about Lenovo and their global partnership with SAP, see this website:

<http://www.lenovo.com/solutions/sap>

## 4.6.2 Lenovo Entry SAN Storage

The S2200 and S3200 are Lenovo's latest feature-rich, high-performance storage area network (SAN) offerings that are equipped to manage demanding workloads in any environment. Configured with hybrid connectivity options, dual controllers, redundant power supplies, and hot swap drives and fans, the S2200 and S3200 SANs are built to provide the always on, dependable availability and scalability that is required to grow with a business. Lenovo Storage brings enterprise-class features, such as Intelligent Real-time Tiering, SSD Caching, Thin Provisioning, Snapshots, and Rapid RAID Rebuild. The S2200 and S3200 are described in the following sections.

### Lenovo Storage S2200

The Lenovo Storage S2200 SAN array offers simplicity, speed, and scalability. With its easy-to-use management interface, the S2200 makes complex administrative storage tasks simple. There is no need for another dedicated resource to manage this array. Setup also is simple by using the intuitive navigation; it is unpacked and deployed in under an hour. The Lenovo S2200 offers a great balance between simplicity and performance.

Figure 4-30 shows the Lenovo Storage S2200:



Figure 4-30 Lenovo Storage S2200

### **S2200 key features**

The S2200 includes the following key features:

- ▶ Ease of use: Ready for use setup and easy deployment.
- ▶ Support for up to 96 drives.
- ▶ 254 Snapshots per volume.
- ▶ Includes the Lenovo SAN Manager, which is an intuitive easy to use GUI.
- ▶ Data Tiering: Automatically moves data between storage layers, which increases overall storage performance.
- ▶ Thin provisioning: Supports business applications that must grow dynamically while using only the space that is actually used.
- ▶ SSD Read Caching: Enables priority access to “hot” data accelerating read performance.
- ▶ Rapid RAID Rebuild: Minimizes recovery time and risk factors with rapid data restoration.
- ▶ Snapshot: Provides point-in-time copies of live data, which eliminates the need for restore and performance hits.
- ▶ Storage Pooling: Virtualizes storage across multiple drive types, which improves I/O performance without affecting applications.

### **S2200 specifications**

The S2200 features the following specifications:

- ▶ Dual controller or single controllers with 6 GB of cache
- ▶ 6/12 Gbps SAS or 1/10 Gbps iSCSI or 8 Gbps Fibre Channel
- ▶ Support for 12x 3.5-inch large form factor or 24x 2.5-inch small form factor drives, hot-swappable 6 Gb SAS, NL SAS, or SSD drives
- ▶ Support for RAID 0, 1, 5, 6, and 10
- ▶ Scales up to three expansion units with 12 x 3.5-inch drives or 24 x 2.5-inch drives
- ▶ LUNs per system: 1024 maximum and 128 TB maximum LUN size
- ▶ Includes: 64 snapshots that can be upgraded to 256 or 512 snapshots

For more information about the S2200, see *Lenovo Storage S2200*, TIPS1298, which is available at this website:

<http://lenovopress.com/tips1298>

### **Lenovo Storage S3200**

The Lenovo Storage S3200 SAN is a feature-rich, high-performing mid-range SAN ready for demanding workloads for any environment. Offering hybrid configurations and real-time tiering, the S3200 can provide near All-Flash-Array (AFA) performance and up to 100,000 IOPS at a fraction of the cost. Along with extreme performance, the S3200 provides an easy-to-use GUI, thin provisioning, Virtual Snapshots, Flash Cache, Virtualized Storage, Rapid RAID Rebuild and more. By combining high performance and incredible features, the S3200 is the perfect mid-range SAN to support the needs of a business.

Figure 4-31 shows the S3200.



Figure 4-31 Lenovo Storage S3200

### **S3200 key features**

The S3200 includes the following key features:

- ▶ Ease of use: Read for use setup and easy deployment.
- ▶ Comes with the Lenovo SAN Manager, which is an intuitive easy to use GUI.
- ▶ Data Tiering: Automatically moves data between storage layers, which increases overall storage performance.
- ▶ Thin provisioning: Support business applications that must grow dynamically while using only the space that is actually used.
- ▶ SSD Read Caching: Enables priority access to “hot” data, which accelerates read performance.
- ▶ Rapid RAID Rebuild: Minimizes recovery time and risk factors with rapid data restoration.

- ▶ Snapshot: Provides point-in-time copies of live data, which eliminates the need for restore and performance hits.
- ▶ Storage Pooling: Virtualizes storage across multiple drive types, which improves I/O performance without affecting applications.

### **S3200 specifications**

The S3200 features the following specifications:

- ▶ Dual controller or single controllers with 6 GB of cache
- ▶ 6/12 Gbps SAS or 1/10 Gbps iSCSI or 8/16 Gbps Fibre Channel connectivity
- ▶ Support for 12x 3.5-inch large form factor or 24x 2.5-inch small form factor drives, hot-swappable 6 Gb SAS, NL SAS, or SSD drives
- ▶ Support for RAID 0, 1, 5, 6, and 10
- ▶ Scales up to seven expansion units with 12 x 3.5-inch drives or 24 x 2.5-inch drives
- ▶ LUNs per system: 1024 maximum and 128 TB maximum LUN size
- ▶ Includes 128 snapshots, which can be upgraded to 512 or 1024 snapshots

For more information about the S3200, see *Lenovo Storage S3200*, TIPS1299, which is available at this website:

<http://lenovopress.com/tips1299>

## **4.6.3 IBM Storwize for Lenovo**

The IBM Storwize family of products are designed for software-defined environments. The Storwize family includes built-in functions, such as Real-time Compression and Easy Tier technology, which combines flash and HDDs to deliver extraordinary levels of efficiency and high performance.

Storwize is available in a wide range of storage systems, including the V3700, V5000, and V7000. The Storwize family delivers sophisticated capabilities that are easy to deploy and help control costs for growing businesses.

### **Storwize V3700**

Storwize V3700 is an entry level disk storage system that is designed with sophisticated capabilities that are unusual for a system of this class. It offers efficiency and flexibility through built-in thin provisioning and nondisruptive migration of data from existing storage. Built upon the innovative technology in the Storwize family, Storwize V3700 addresses block storage requirements of small and midsize organizations at an affordable price.

Figure 4-32 shows the Storwize V3700.



Figure 4-32 Storwize V3700



### ***Storwize V3700 key features***

The Storwize V3700 includes the following key features:

- ▶ Ease of use: The new generation GUI and point-and-click system management capabilities offer easy-to-use data management.
- ▶ FlashCopy: Creates instant copies of data for backup or application testing.
- ▶ Thin provisioning: Supports business applications that must grow dynamically while using only the space that is actually used.
- ▶ Easy Tier: Delivers better performance while controlling costs.
- ▶ Internal virtualization: Helps you quickly respond to changing requirements.
- ▶ Non-disruptive migration: Speeds implementation and minimizes disruption.
- ▶ Remote Mirror: Helps create copies of data at remote locations for disaster recovery.
- ▶ IP replication: Reduces remote mirroring costs with innovative network optimization.
- ▶ Turbo performance: Increases system maximum IOPS and maximum throughput.

### ***Storwize V3700 specifications***

The Storwize V3700 features the following specifications:

- ▶ Dual-active intelligent array node canisters with up to 8 GB cache per canister
- ▶ 6 Gb SAS and 1 Gb iSCSI connectivity standard with optional 8 Gb Fibre Channel (FC) or 10 Gb iSCSI or Fibre Channel over Ethernet (FCoE) connectivity
- ▶ Support for 12 3.5-inch large form factor or 24 2.5-inch small form factor drives, hot-swappable 6 Gb SAS disk drives
- ▶ Support for RAID 0, 1, 5, 6, and 10
- ▶ Scalable up to 240 drives per system with the attachment of Storwize V3700 expansion units
- ▶ Internal virtualization, thin provisioning, one-way data migration, and FlashCopy (up to 64 targets, which can be upgraded to 2,040 targets)

For more information about the Storwize V3700, see this website:

<http://lenovopress.com/tips1300>

### **Storwize V5000**

Storwize V5000 is a mid-range disk storage system that is a highly flexible, easy to use virtualized storage system that enables mid-sized organizations to meet the challenges of rapid data growth and limited IT budgets. As an intermediate Storwize family offering, Storwize V5000 enables organizations to consolidate and provide new capabilities to their existing infrastructures.

Figure 4-33 shows the Storwize V5000.



Figure 4-33 Storwize V5000

### ***Storwize V5000 key features***

The Storwize V5000 includes the following key features:

- ▶ Higher flexibility: Easily customizable with flexible software options.
- ▶ Ease of use: Simplifies management with industry-leading graphical user interface.
- ▶ Easy Tier: Advanced technology for automatically migrating data between storage tiers based on real-time usage analysis patterns.
- ▶ External virtualization: Consolidates and provides Lenovo Storwize V5000 capabilities to storage infrastructures.
- ▶ Replication over IP: Improves network usage for remote mirroring with innovative technology.
- ▶ Clustered systems: Provides the ability for block systems to scale up and out for performance and capacity.

### ***Storwize V5000 specifications***

The Storwize V5000 features the following specifications:

- ▶ Dual-active intelligent array node canisters with up to 8 GB cache per canister (32 GB with two-way clustered systems)
- ▶ 1 Gb iSCSI, 6 Gb SAS and 8 Gb FC, or 10 Gb iSCSI/FCoE
- ▶ Support for 12x 3.5-inch large form factor or 24x 2.5-inch small form factor drives, hot-swappable 6 Gb SAS disk drives, or 24x 2.5-inch SSDs
- ▶ Support for RAID 0, 1, 5, 6, and 10
- ▶ Up to 19 Storwize V5000 expansion enclosures (maximum of 480 drives per system and 960 drives in two-way clustered systems)
- ▶ Included: Internal virtualization, thin provisioning, one-way data migration, and dual-system clustering.

For more information about the Storwize V5000, see this website:

<http://lenovopress.com/tips1301>

### **Storwize V7000**

Designed for software-defined storage environments, the Storwize V7000 Storage System is a virtualized, enterprise-class storage system that provides the foundation for implementing an effective storage infrastructure and transforming the economics of data storage. Designed to complement virtual server environments, this modular storage system delivers the flexibility and responsiveness that is required for changing business needs.

Figure 4-34 shows the Storwize V7000.



Figure 4-34 Storwize V7000

### ***Storwize V7000 key features***

The Storwize V7000 includes the following key features:

- ▶ Delivers sophisticated, enterprise-class storage functionality for businesses.
- ▶ Supports your growing business requirements while controlling costs.
- ▶ Provides up to three times performance improvement by moving as little as 5% of data to flash storage.
- ▶ Enables storing up to 5x more active primary data in the same physical disk space by using Real-time Compression.
- ▶ Improves network usage for remote mirroring with innovative replication technology.
- ▶ Consolidates block and file storage for simplicity, greater efficiency, and ease of management.
- ▶ Enables near-continuous availability of applications.
- ▶ Incorporates easy-to-use data management that is designed with a GUI and point-and-click system management capabilities.
- ▶ Includes Metro Mirror and Global Mirror for replicating data synchronously or asynchronously between systems for backup efficiency.
- ▶ Has flash storage for applications that demand high speed and quick access to data.

### ***Storwize V7000 specifications***

The Storwize V7000 features the following specifications:

- ▶ Dual-active intelligent array node canisters with 32 or 64 Gb of cache per controller, 64 or 128 Gb per enclosure, up to 512 Gb per cluster
- ▶ SAN-attached 8 Gbps FC, 1 Gbps iSCSI, and optional 10 Gbps iSCSI/FCoE network-attached storage (NAS) attached 1 Gbps and 10 Gbps Ethernet
- ▶ Support for 12x 3.5-inch large form factor or 24x 2.5-inch small form factor drives, hot-swappable 6 Gb SAS, 7.2k NL SAS disk drives, or 2.5-inch SSDs
- ▶ Up to 504 drives per control enclosure; 1,056 per clustered system
- ▶ Support for RAID 0, 1, 5, 6, and 10
- ▶ Easy Tier, FlashCopy, and thin provisioning included
- ▶ Other features include:
  - Remote mirroring
  - Real-time Compression (RTC)
  - Encryption
  - Virtualization of external storage

For more information about the Storwize V7000, see this website:

<http://lenovopress.com/tips1302>

## 4.7 Lenovo Networking offerings

The next generation of data centers seek to improve performance and usage of their resources, simplify management, and reduce complexity of their networks, while simultaneously trying to reduce capital and operational expenses. To meet these requirements, Lenovo Networking provides solutions that increase ease of use and management, help scale-out architectures to meet the growing needs of businesses, reduce power consumption and cost, and provide high performance and virtualization aware networking.

The following models of Lenovo Top-of-Rack Ethernet switches can be used in Lenovo solutions for SAP Business Suite:

- ▶ **Lenovo RackSwitch™ G7028**

The G7028 is a Layer 2 1/10 Gb Ethernet switch with 24x 10/100/1000BASE-T RJ45 ports and four 10 Gigabit Ethernet SFP+ ports.

- ▶ **Lenovo RackSwitch G7052**

The G7052 is a Layer 2 1/10 Gb Ethernet switch with 48x 10/100/1000BASE-T RJ-45 ports and four 10 Gigabit Ethernet SFP+ ports.

- ▶ **Lenovo RackSwitch G8052**

The G8052 is a Layer 2/3 1/10 Gb Ethernet switch with 48x 10/100/1000 BASE-T RJ45 ports and four 10 GbE SFP+ ports.

- ▶ **Lenovo RackSwitch G8124E**

The G8124 is a Layer 2/3 10 Gb Ethernet switch with 24x 10 GbE SFP+ ports.

- ▶ **Lenovo RackSwitch G8264**

The G8264 is a Layer 2/3 10/40 Gb Ethernet switch with 48x 10 GbE SFP+ ports and four 40 GbE QSFP+ ports.

- ▶ **Lenovo RackSwitch G8272**

The G8272 is a Layer 2/3 10/40 Gb Ethernet switch with 48x 10 GbE SFP+ ports and six 40 GbE QSFP+ ports.

- ▶ **Lenovo RackSwitch G8264CS**

Lenovo G8264CS has 36 SFP/SFP+ ports for 1 Gb or 10 Gb Ethernet, 12 Omni Ports™, each of which can operate as a 10 Gb Ethernet or auto-negotiating 4/8 Gb FC, and for QSFP+ ports for 40 Gb Ethernet.

- ▶ **Lenovo RackSwitch G8296**

Lenovo G8296 is a Layer 2/3 10/40 Gb Ethernet switch. It has 86 SFP/SFP+ ports for 1 Gb or 10 Gb Ethernet, 8 QSFP+ ports for 40 Gb Ethernet and 2 QSFP+ ports for either 40 Gb or 4x 10Gb Ethernet connections.

- ▶ **Lenovo RackSwitch G8332**

The G8332 is a high-density 40 Gigabit Ethernet (GbE) switch with 32 ports for 40 GbE.

### 4.7.1 Lenovo RackSwitch G7028

The RackSwitch G7028 is an inexpensive entry level Layer 2 1/10 Gb Ethernet switch that is designed for the data center. Data center customers continue to deploy applications by using 1 GbE networking switches while 10 GbE continues to grow for customer solutions that require greater I/O bandwidth. The RackSwitch G7028 is an ideal solution for data center customers who need an economical 1 Gb and 10 Gb connectivity solution.

The RackSwitch G7028 is a Top-of-Rack data center switch that delivers excellent non-blocking, wire speed performance at an attractive price point. It supports 24 10/100/1000BASE-T RJ45 ports and four 10 Gigabit Ethernet SFP+ ports standard and often uses only 45 watts of power. The RackSwitch G7028 is also designed with rear-to-front airflow, which allows for flexible mounting of the switch in a rack cabinet and provides convenient cable management and significant savings in cooling costs.

Figure 4-35 shows the front view of the RackSwitch G7028.



Figure 4-35 RackSwitch G7028

### **G7028 key benefits**

The G7028 includes the following key benefits:

- ▶ Increased network performance  
The RackSwitch G7028 provides up to 128 Gbps of switching throughput, supports four SFP+ 10 Gb uplink ports for zero oversubscription, and has a port-to-port latency of 3.3 microseconds.
- ▶ Lower power and exceptional cooling  
The RackSwitch G7028 typically uses only 45 W of power. Unlike side-cooled switches that can cause heat recirculation and reliability concerns, the rear-to-front cooling design of the RackSwitch G7028 switch reduces data center air conditioning costs by matching airflow to the server's configuration in the rack.
- ▶ Fault tolerance  
These switches learn alternative paths automatically and perform faster convergence if there is a link, switch, or power failure. The switch uses technologies, such as Virtual Link Aggregation, L2 trunk failover, and HotLinks.
- ▶ Seamless interoperability  
RackSwitch switches interoperate seamlessly with the upstream switches of other vendors.

### **G7028 ports**

The G7028 features the following ports:

- ▶ A total of 24 10/100/1000 Mb Ethernet ports with RJ-45 connectors.
- ▶ Four SFP/SFP+ ports for 1 Gb or 10 Gb Ethernet transceivers (support for 1000BASE-SX, 1000BASE-LX, 1000BASE-T, 10GBASE-SR, or 10GBASE-LR) or SFP+ direct-attach copper (DAC) cables.
- ▶ One Ethernet port (RJ-45 connector) for out of band (OOB) management.
- ▶ One RS-232 serial port (RJ-45 connector) that provides another means to configure the switch module.
- ▶ One USB port for mass storage devices.

Figure 4-36 shows the ports on the G7028.




Figure 4-36 G7028 ports

For more information about the G7028, see *Lenovo RackSwitch G7028*, TIPS1268, which is available at this website:

<http://lenovopress.com/tips1268>

## 4.7.2 Lenovo RackSwitch G7052

RackSwitch G7052 is a low-cost, entry-level Layer 2 1/10 Gb Ethernet switch that is designed for the data center. Data center customers continue to deploy applications by using 1 GbE networking switches while 10 GbE continues to grow upstream in the network as clients require greater I/O bandwidth. The RackSwitch G7052 is an ideal solution for data center customers who need an economical 1 Gb and 10 Gb connectivity solution.

The RackSwitch G7052 is a top-of-rack data center switch that delivers excellent non-blocking, wire speed performance at an attractive price point. It supports 48x 10/100/1000BASE-T RJ-45 ports and four 10 Gigabit Ethernet SFP+ ports standard, and typically uses only 76 watts of power. The RackSwitch G7052 is also designed with rear-to-front airflow, which allows for flexible mounting of the switch in a rack cabinet and provides convenient cable management and significant savings in cooling costs.

Figure 4-37 shows the G7052.



Figure 4-37 RackSwitch G7052 front view

### **G7052 key benefits**

The G7052 includes the following key benefits:

- ▶ Increased network performance

The RackSwitch G7052 provides up to 176 Gbps of switching throughput, supports four SFP+ 10 Gb uplink ports, and has a port to port latency of 3.3 microseconds.

- ▶ Lower power and exceptional cooling
 

The RackSwitch G7052 often uses only 76 W of power. Unlike side-cooled switches that can cause heat recirculation and reliability concerns, the rear-to-front cooling design of the RackSwitch G7052 switch reduces data center air conditioning costs by matching airflow to the server's configuration in the rack.
- ▶ Fault tolerance
 

These switches learn alternative paths automatically and perform faster convergence if there is a link, switch, or power failure. The switch uses technologies, such as Virtual Link Aggregation, L2 trunk failover, and HotLinks.
- ▶ Seamless interoperability
 

RackSwitch switches interoperate seamlessly with the upstream switches of other vendors.

### G7052 ports

The G7052 features the following ports:

- ▶ A total of 48 10/100/1000 Mb Ethernet ports with RJ-45 connectors.
- ▶ Four SFP/SFP+ ports for 1 Gb or 10 Gb Ethernet transceivers (support for 1000BASE-SX, 1000BASE-LX, 1000BASE-T, 10GBASE-SR, or 10GBASE-LR) or SFP+ direct-attach copper (DAC) cables. SFP+ modules or DAC cables are not included and must be purchased separately.
- ▶ One Ethernet port (RJ-45 connector) for OOB management.
- ▶ One RS-232 serial port (RJ-45 connector) that provides another means to configure the switch module.
- ▶ One USB port for mass storage devices.

Figure 4-38 shows the G7052 ports.




Figure 4-38 G7052 ports

For more information about the G7052, see *Lenovo RackSwitch G7052*, TIPS1269, which is available at this website:

<http://lenovopress.com/tips1269>

### 4.7.3 Lenovo RackSwitch G8052

This switch is a Top-of-Rack switch that is designed for a data center. It combines great performance, server-like airflow for cooling, and low-power consumption in a virtualization-ready package.



Figure 4-39 shows the Lenovo RackSwitch G8052 Top-of-Rack switch.




Figure 4-39 *Lenovo RackSwitch G8052*

The RackSwitch G8052 is a Top-of-Rack data center switch that delivers unmatched line-rate Layer 2/3 performance at an attractive price. It has 48 10/100/1000 BASE-T RJ45 ports and four 10 Gigabit Ethernet SFP+ ports, and includes hot-swap redundant power supplies and fans standard, which minimizes your configuration requirements. Unlike most rack equipment that cools from side-to-side, the RackSwitch G8052 has rear-to-front or front-to-rear airflow that matches server airflow.

For 10 Gb uplinks, there is a choice of SFP+ transceivers (SR or LR) for longer distances or more cost-effective and lower-power usage options, such as SFP+ direct-attached cables (DAC or Twinax cables), which can be 1 - 7 meters and are ideal for connecting to another Top-of-Rack switch or an adjacent rack.

### ***G8052 key benefits***

The G8052 includes the following key benefits:

► High performance

The RackSwitch G8052 provides up to 176 Gbps throughput and supports four SFP+ 10 G uplink ports for a low oversubscription ratio, in addition to a low latency of 1.8 microseconds.

► Lower power and better cooling

The RackSwitch G8052 typically uses only 130 W of power, which is a fraction of the power consumption of most competitive offerings. Unlike side-cooled switches, which can cause heat recirculation and reliability concerns, the G8052's rear-to-front or front-to-rear cooling design reduces data center air conditioning costs by matching airflow to the server's configuration in the rack. Variable speed fans help to automatically reduce power consumption.

► VM-aware network virtualization

VMready® software on the switch simplifies configuration and improves security in virtualized environments. VMready automatically detects VM movement between physical servers and instantly reconfigures each VM's network policies across VLANs to keep the network up and running without interrupting traffic or affecting performance. VMready works with all leading hypervisors, such as VMware, Citrix Xen, Red Hat KVM, Microsoft Hyper-V, and IBM PowerVM hypervisor.

► Layer 3 functionality

The RackSwitch G8052 includes Layer 3 functionality, which provides security and performance benefits plus the full range of Layer 3 static and dynamic routing protocols, including Open Shortest Path First (OSPF) and Border Gateway Protocol (BGP) for enterprise customers at no extra cost.



► Fault tolerance

These switches learn alternative routes automatically and perform faster convergence if there is a link, switch, or power failure. The switch uses proven technologies, such as L2 trunk failover, VLAN-based failover, VRRP, HotLink, IGMP V3 snooping, and OSPF.

► OpenFlow enabled

The RackSwitch G8052 offers benefits of OpenFlow. OpenFlow is the new open application programming interface (API) that enables the network administrator to easily configure and manage virtual networks that control traffic on a “per-flow” basis. It creates multiple independent virtual networks and related policies without dealing with the complexities of the underlying physical network and protocols. The G8052 can also be used with any industry-compliant OpenFlow controller.

### **G8052 ports**

The G8052 features the following ports:

- A total of 48 auto-sensing 10/100/1000 Mb Ethernet ports with RJ-45 connectors.
- Four SFP/SFP+ ports for 1 Gb or 10 Gb Ethernet transceivers (support for 1000BASE-SX, 1000BASE-LX, 1000BASE-T, 10GBASE-SR, 10GBASE-LR, or 10GBASE-ER) or SFP+ direct-attach copper (DAC) cables.
- One RS-232 serial port (Mini-USB connector) that provides an extra means to configure the switch module.
- One USB port for mass storage devices.

Figure 4-40 shows the G8052 ports.




Figure 4-40 G8052 ports

For more information about the G8052, see *Lenovo RackSwitch G8052*, TIPS1270, which is available at this website:

<http://lenovopress.com/tips1270>

## **4.7.4 Lenovo RackSwitch G8124E**

The Lenovo RackSwitch G8124E is a 10 Gigabit Ethernet switch that is specifically designed for the data center and provides a virtualized, cooler, and easier network solution. The G8124E offers 24 10 Gigabit Ethernet ports in a high-density, 1U footprint. Designed with top performance in mind, the RackSwitch G8124E provides line-rate, high-bandwidth switching, filtering, and traffic queuing without delaying data and large data center grade buffers to keep traffic moving.

Figure 4-41 shows the RackSwitch G8124E Top-of-Rack switch.



Figure 4-41 Lenovo RackSwitch G8124E

### **G8124E key benefits**

The G8124E includes the following key benefits:

- ▶ High performance

The 10 GbE low latency (as low as 570 nanoseconds) switch provides the best combination of extremely low latency, non-blocking line-rate switching, and ease of management.

- ▶ Lower power and better cooling

The G8124E uses as little power as two 60 W light bulbs, which is a fraction of the power consumption of most competitive offerings. Unlike side-cooled switches, which can cause heat recirculation and reliability concerns, the G8124E rear-to-front cooling design reduces data center air conditioning costs by having airflow match the servers in the rack. In addition, variable speed fans help to automatically reduce power consumption.

- ▶ Virtual Fabric

Virtual Fabric can help customers address I/O requirements for multiple NICs while also helping reduce cost and complexity. Virtual Fabric allows for the carving up of a physical NIC into multiple virtual NICs (2 - 8 vNICs) and creates a virtual pipe between the adapter and the switch for improved performance, availability, and security while reducing cost and complexity.

- ▶ VM-aware networking

VMready software on the switch helps reduce configuration complexity while significantly improving security levels in virtualized environments. VMready automatically detects VM movement from one physical server to another, and instantly reconfigures each VM's network policies across VLANs to keep the network up and running without interrupting traffic or affecting performance. VMready works with all leading VM providers, such as VMware, Citrix, Xen, IBM PowerVM, and Microsoft Hyper-V.

- ▶ Layer 3 functionality

The switch includes Layer 3 functionality, which provides security and performance benefits as inter-VLAN traffic stays within the chassis. This switch also provides the full range of Layer 3 protocols from static routes for technologies, such as OSPF and BGP for enterprise customers.

- ▶ Seamless interoperability

RackSwitch switches interoperate seamlessly with the upstream switches of other vendors. For more information, see *Tolly Reports: Tolly Functionality and Certification: RackSwitch G8000 and G8124 and Cisco Catalyst Interoperability Evaluation*, which is available at this website:

<http://ibm.co/ZRwiR0>

► Fault tolerance

These switches learn alternative routes automatically and perform faster convergence if there is a link, switch, or power failure. The switch uses proven technologies, such as L2 trunk failover, advanced VLAN-based failover, VRRP, Hot Links, IGMP V3 snooping, and OSPF.

► Converged fabric

The switch supports Converged Enhanced Ethernet/Data center bridging (CEE/DCB) and connectivity to FCoE gateways. CEE helps enable clients to combine storage, messaging traffic, VoIP, video, and other data on a common data center Ethernet infrastructure. FCoE helps enable highly efficient block storage over Ethernet for consolidating server network connectivity. As a result, clients can deploy a single-server interface for multiple data types, which can simplify deployment and management of server network connectivity while maintaining the high availability and robustness that is required for storage transactions.

► Transparent networking capability

With a simple configuration change to Easy Connect mode, the RackSwitch G8124E becomes a transparent network device, which is invisible to the core and eliminates network administration concerns of Spanning Tree Protocol configuration and interoperability, VLAN assignments, and avoids any possible loops. By emulating a host NIC to the data center core, it accelerates the provisioning of VMs by eliminating the need to configure the typical access switch parameters.

### **G8124E ports**

The G8124 features the following ports:

- 24 SFP/SFP+ ports for 1 Gb or 10 Gb Ethernet transceivers (support for 1000BASE-SX, 1000BASE-LX, 1000BASE-T, 10GBASE-SR, 10GBASE-LR, and 10GBASE-ER) or SFP+ direct-attach copper cables.
- Two 10/100/1000 Ethernet port (RJ-45 connector) for OOB management.
- One RS-232 serial port (Mini-USB connector) that provides another means to configure the switch.

Figure 4-42 shows the G8124E ports.




Figure 4-42 G8124E ports

For more information about the G8124E, see *Lenovo RackSwitch G8124E*, TIPS1271, which is available at this website:

<http://lenovopress.com/tips1271>

## 4.7.5 Lenovo RackSwitch G8264

The RackSwitch G8264 is ideal for latency-sensitive applications, such as high performance computing clusters and financial applications, and is a key piece when you are connecting multiple chassis together with Pure System Solutions.

The G8264 supports stacking for up to eight switches for simplified switch management and Lenovo Virtual Fabric to enable clients to reduce the number of I/O adapters by creating virtual NICs out of a 10 Gb adapter, which helps reduce cost and complexity. The G8264 supports the newest protocols, including CEE/DCB) for support of FCoE, in addition to iSCSI and NAS.

Figure 4-43 shows the RackSwitch G8264.



Figure 4-43 Lenovo RackSwitch G8264

### **G8264 key benefits**

The G8264 includes the following key benefits:

- ▶ High performance

The 10 Gb/40 Gb switch provides the best combination of low latency, non-blocking line-rate switching, and ease of management. It has a throughput of 1.28 Tbps.

- ▶ Lower power and better cooling

The RackSwitch G8264 uses as little as 330 W of power, which is a fraction of the power consumption of most competitive offerings. Unlike side-cooled switches, which can cause heat recirculation and reliability concerns, the front-to-rear or rear-to-front cooling design of the G8264 reduces data center air conditioning costs by having airflow match the servers in the rack. In addition, variable speed fans help to automatically reduce power consumption.

- ▶ Stacking

With the G8264, a single switch image and configuration file can be used for up to eight switches, which shares only one IP address and one management interface.

- ▶ Virtual Fabric

The G8264 can help customers address I/O requirements for multiple NICs while reducing cost and complexity. Virtual Fabric allows you to carve a physical NIC into multiple virtual NICs (2 - 8 vNIC™) and to create a virtual pipe between the adapter and the switch for improved performance, availability, and security.

- ▶ VM-aware networking

VMready software on the switch simplifies configuration and improves security in virtualized environments. VMready automatically detects VM movement between physical servers and instantly reconfigures each VM's network policies across VLANs to keep the network up and running without interrupting traffic or affecting performance. VMready works with all leading VM providers, such as VMware, Citrix, Xen, and Microsoft Hyper-V.

- ▶ **Layer 3 functionality**  
The G8264 includes Layer 3 functionality, which provides security and performance benefits, as inter-VLAN traffic stays within the switch. This switch also provides the full range of Layer 3 protocols from static routes for technologies, such as OSPF and BGP for enterprise customers.
- ▶ **Seamless interoperability**  
RackSwitch switches perform seamlessly with the upstream switches of other vendors.
- ▶ **Fault tolerance**  
RackSwitch switches learn alternative routes automatically and perform faster convergence if there is a link, switch, or power failure. The switch uses proven technologies, such as L2 trunk failover, advanced VLAN-based failover, VRRP, and Hot Links.
- ▶ **Multicast**  
The G8264 supports IGMP Snooping v1, v2, and v3 with 2K IGMP groups. It also supports Protocol Independent Multicast (PIM), such as PIM Sparse Mode or PIM Dense Mode.
- ▶ **Converged fabric**  
The G8264 switch supports CEE and connectivity to FCoE gateways. CEE helps enable clients to combine storage, messaging traffic, VoIP, video, and other data on a common data center Ethernet infrastructure. FCoE helps enable highly efficient block storage over Ethernet for consolidating server network connectivity. As a result, clients can deploy a single-server interface for multiple data types, which can simplify deployment and management of server network connectivity while maintaining the high availability and robustness required for storage transactions.
- ▶ **OpenFlow enabled**  
The RackSwitch G8264 is the first 10 GbE switch to offer benefits of OpenFlow. OpenFlow is the new open API that enables the network administrator to easily configure and manage virtual networks that control traffic on a “per-flow” basis. It creates multiple independent virtual networks and related policies without dealing with the complexities of the underlying physical network and protocols. The G8264 is also the ideal switch to use with industry-compliant OpenFlow controllers.
- ▶ **Transparent networking capability**  
With a simple configuration change to Easy Connect Mode, the RackSwitch G8264 becomes a transparent network device that is invisible to the core and eliminates network administration concerns of Spanning Tree Protocol configuration and interoperability, VLAN assignments, and avoids any possible loops.

### **G8264 ports**

The G8264 features the following ports:

- ▶ A total of 48 ports for 1 Gb or 10 Gb Ethernet SFP/SFP+ transceivers (support for 1000BASE-SX, 1000BASE-LX, 1000BASE-T, 10GBASE-SR, 10GBASE-LR, or 10GBASE-ER) or SFP+ direct-attach copper cables.
- ▶ Four ports for 40 Gb Ethernet QSFP+ transceivers, QSFP+ to QSFP+ DAC cables, or QSFP+ to 4x 10 Gb SFP+ break-out cables.
- ▶ One 10/100/1000 Ethernet port (RJ-45 connector) for OOB management.
- ▶ One RS-232 serial port (mini-USB connector) that provides another means to configure the switch.
- ▶ One USB port for mass storage devices.

Figure 4-44 shows the G8264 ports.




Figure 4-44 G8264 ports

For more information about the G8264, see *Lenovo RackSwitch G8264*, TIPS1272, which is available at this website:

<http://lenovopress.com/tips1272>

## 4.7.6 Lenovo RackSwitch G8264CS

Many clients successfully use Ethernet and FC connectivity from their servers to their LAN and SAN. These clients are seeking ways to reduce the cost and complexity of these environments by using the capabilities of Ethernet and FC convergence.

The RackSwitch G8264CS Top-of-Rack switch offers the benefits of a converged infrastructure. As part of its forward-thinking design, this switch has flexibility for future growth and expansion. This switch is ideal for clients who want to connect to SANs and clients who want native FC connectivity, in addition to support for such protocols as Ethernet, FCoE, and iSCSI.

Figure 4-45 shows the front of the G8264CS top-of-rack switch.



Figure 4-45 Lenovo RackSwitch G8264CS

### **G8264CS key benefits**

The G8264CS includes the following key features:

- Lowers the TCO by using consolidation

By consolidating LAN and SAN networks and converging to a single fabric, clients can reduce the equipment that is needed in their data centers. This benefit significantly reduces energy and cooling, management and maintenance, and capital costs.

- ▶ Improves performance and increases availability

The G8264CS is an enterprise-class and full-featured data center switch that offers high-bandwidth performance with 36 1/10 Gb SFP+ connections, 12 Omni Ports that can be used for 10 Gb SFP+ connections, 4/8 Gb Fibre Channel connections (or both), and four 40 Gb QSFP+ connections. The G8264CS switch delivers full line rate performance on Ethernet ports, which makes it an ideal choice for managing dynamic workloads across the network.

This switch also provides a rich Layer 2 and Layer 3 feature set that is ideal for many of today's data centers. Combined with redundant hot-swappable power and fans, and numerous high availability features, this switch comes fully equipped to handle the demands of business-sensitive traffic.

- ▶ High performance

The 10 Gb/40 Gb switch provides the best combination of low latency, non-blocking line-rate switching, and ease of management. It has a throughput of up to 1.28 Tbps.

- ▶ Lower power and better cooling

The G8264CS uses as little as 330 W of power, which is a fraction of the power consumption of most competitive offerings. Unlike side-cooled switches, which can cause heat recirculation and reliability concerns, the front-to-rear or rear-to-front cooling design of the G8264CS switch reduces the costs of data center air conditioning by having airflow match the servers in the rack. In addition, variable speed fans help to automatically reduce power consumption.

- ▶ Support for Virtual Fabric

The G8264CS can help customers address I/O requirements for multiple NICs while reducing cost and complexity. By using Virtual Fabric, you can carve a physical dual-port NIC into multiple vNICs (2 - 8 vNICs) and create a virtual pipe between the adapter and the switch for improved performance, availability, and security. FCoE also is supported; two vNIC can be configured as CNAs to allow for more cost savings through convergence.

- ▶ VM-aware networking

VMready software on the switch simplifies configuration and improves security in virtualized environments. VMready automatically detects VM movement between physical servers and instantly reconfigures the network policies of each VM across VLANs to keep the network up and running without interrupting traffic or affecting performance. VMready works with all leading VM providers, such as VMware vSphere, Citrix Xen, IBM PowerVM, and Microsoft Hyper-V.

- ▶ Layer 3 functionality

The G8264CS includes Layer 3 functionality, which provides security and performance benefits because inter-VLAN traffic stays within the switch. This switch also provides the full range of Layer 3 protocols from static routes for technologies, such as OSPF and BGP for enterprise customers.

- ▶ Seamless interoperability

The G8264CS switches perform seamlessly with the upstream switches of other vendors.

- ▶ Fault tolerance

The G8264CS switches learn alternative routes automatically and perform faster convergence if there is a link, switch, or power failure. The switch uses proven technologies, such as L2 trunk failover, VLAN-based failover, VRRP, and Hot Links.



► Multicast support

The switch supports IGMP Snooping v1, v2, and v3 with 2K IGMP groups. It also supports PIM, such as PIM Sparse Mode or PIM Dense Mode.

► Transparent networking capability

With support for Easy Connect mode, the RackSwitch G8264CS becomes a transparent network device that is invisible to the core and eliminates network administration concerns of Spanning Tree Protocol configuration and interoperability and VLAN assignments, and avoids any possible loops. By emulating a host NIC to the data center core, it accelerates the provisioning of VMs by eliminating the need to configure the access switch settings.

**G8264CS ports**

The G8264CS features the following ports (SFP+ modules and DAC cables are not included and must be purchased together with the switch):

- A total of 36 ports for 1 Gb or 10 Gb Ethernet SFP/SFP+ transceivers (support for 1000BASE-SX, 1000BASE-LX, 1000BASE-T, 10GBASE-SR, 10GBASE-LR, or 10GBASE-ER) or SFP+ direct-attach copper cables.
- A total of 12 Omni Ports, each of which can operate as a 10 Gb Ethernet (support for 10GBASE-SR, 10GBASE-LR, 10GBASE-ER, or 10 GbE SFP+ DAC cables), or auto-negotiating 4/8 Gb Fibre Channel, depending on the SFP+ transceiver that is installed in the port.

**Note:** Omni Ports do not support 1 Gb Ethernet SFP transceivers.

- Four ports for 40 Gb Ethernet QSFP+ transceivers, QSFP+ to QSFP+ DAC cables, or QSFP+ to 4x 10 Gb SFP+ break-out cables. QSFP+ modules and DAC cables are not included and must be purchased separately.
- One 10/100/1000 Ethernet port (RJ-45 connector) for OOB management.
- One RS-232 serial port (mini-USB connector) for configuring the switch.
- One USB port for mass storage devices.

Figure 4-46 shows the G8264CS ports.




Figure 4-46 G8264CS ports

For more information about the G8264CS, see *Lenovo RackSwitch G8264CS*, TIPS1273, which is available at this website:

<http://lenovopress.com/tips1273>



## 4.7.7 Lenovo RackSwitch G8272

The Lenovo RackSwitch G8272 that uses 10Gb SFP+ and 40Gb QSFP+ Ethernet technology is specifically designed for the data center. It is ideal for today's big data, cloud, and optimized workload solutions. It is an enterprise class Layer 2 and Layer 3 full-featured switch that delivers line-rate, high-bandwidth switching, filtering, and traffic queuing without delaying data. Large data center-grade buffers help keep traffic moving, while the hot-swap redundant power supplies and fans (along with numerous high-availability features) help provide high availability for business sensitive traffic.

The RackSwitch G8272 is ideal for latency sensitive applications, such as high-performance computing clusters and financial applications. In addition to the 10 GbE and 40 GbE connections, the G8272 can use 1 GbE connections. The G8272 supports the newest protocols, including CEE/DCB for FCoE, iSCSI, and NAS.

The Lenovo RackSwitch G8272 is shown in Figure 4-47.



Figure 4-47 Lenovo RackSwitch G8272

### **G8272 key benefits**

The G8272 includes the following key benefits:

- ▶ High performance

The 10 Gb/40 Gb switch provides the best combination of low latency, non-blocking line-rate switching, and ease of management. It has a throughput of 1.44 Tbps.

- ▶ Lower power and better cooling

The RackSwitch G8272 uses as little as 123 W of power, which is a fraction of the power consumption of most competitive offerings. Unlike side-cooled switches, which can cause heat recirculation and reliability concerns, the front-to-rear or rear-to-front cooling design of the G8272 reduces data center air conditioning costs by having airflow match the servers in the rack. In addition, variable speed fans help to automatically reduce power consumption.

- ▶ Lenovo Virtual Fabric

The G8272 can help customers address I/O requirements for multiple NICs while reducing cost and complexity. By using Virtual Fabric, you can carve a physical NIC into multiple virtual ports (2 - 8 vPorts) and create a virtual pipe between the adapter and the switch for improved performance, availability, and security.

- ▶ VM-aware networking

VMready software on the switch simplifies configuration and improves security in virtualized environments. VMready automatically detects VM movement between physical servers and instantly reconfigures each VM's network policies across VLANs to keep the network up and running without interrupting traffic or affecting performance. VMready works with all leading VM providers, such as VMware, Citrix Xen, and Microsoft Hyper-V.

- ▶ Layer 3 functionality

The G8272 includes Layer 3 functionality provides security and performance benefits as inter-VLAN traffic stays within the switch. This switch also provides the full range of Layer 3 protocols from static routes for technologies, such as OSPF and BGP for enterprise customers.

- ▶ Seamless interoperability

RackSwitch switches perform seamlessly with the upstream switches of other vendors.

- ▶ Fault tolerance

RackSwitch switches learn alternative routes automatically and perform faster convergence if there is a link, switch, or power failure. The switch uses proven technologies, such as L2 trunk failover, advanced VLAN-based failover, VRRP, and Hot Links.

- ▶ Multicast

The G8272 supports IGMP Snooping v1, v2, and v3 with 2 K IGMP groups. It also supports PIM, such as PIM Sparse Mode or PIM Dense Mode.

- ▶ Management

Network Address Translation (NAT) allows a single device, such as a router, to act as an agent between the Internet (or “public network”) and a local (or “private”) network. By using this configuration, only a single, unique IP address is required to represent an entire group of local IP addresses to the external network. NAT also provides security by acting as a firewall between internal and external networks.

- ▶ Converged fabric

The G8272 switch supports CEE and connectivity to FCoE gateways. CEE helps enable clients to combine storage, messaging traffic, VoIP, video, and other data on a common data center Ethernet infrastructure. FCoE helps enable highly efficient block storage over Ethernet for consolidating server network connectivity. As a result, clients can deploy a single-server interface for multiple data types, which can simplify the deployment and management of server network connectivity while maintaining the high availability and robustness that is required for storage transactions.

- ▶ OpenFlow enabled

The RackSwitch G8272 offers benefits of OpenFlow. OpenFlow is the new open API that enables the network administrator to easily configure and manage virtual networks that control traffic on a “per-flow” basis. It creates multiple independent virtual networks and related policies without dealing with the complexities of the underlying physical network and protocols. The G8272 is also the ideal switch to use with industry-compliant OpenFlow controllers.

### **G8272 ports**

The G8272 features the following ports:

- ▶ A total of 48 ports for 1 Gb or 10 Gb Ethernet SFP/SFP+ transceivers (support for 1000BASE-SX, 1000BASE-LX, 1000BASE-T, 10GBASE-SR, 10GBASE-LR, or 10GBASE-ER) or SFP+ direct-attach copper cables.
- ▶ Six ports for 40 Gb Ethernet QSFP+ transceivers, QSFP+ to QSFP+ DAC cables, or QSFP+ to 4x 10 Gb SFP+ breakout cables.
- ▶ One 10/100/1000 Ethernet port (RJ-45 connector) for OOB management.
- ▶ One RS-232 serial port (mini-USB connector) that provides another means to configure the switch.
- ▶ One USB port for mass storage devices.

Figure 4-48 shows G8272 ports.




Figure 4-48 G8272 ports

For more information about the G8272, see *Lenovo RackSwitch G8272*, TIPS1267, which is available at this website:

<http://lenovopress.com/tips1267>

#### 4.7.8 Lenovo RackSwitch G8296

The Lenovo RackSwitch G8296 incorporates 10 Gb SFP+ and 40 Gb QSFP+ Ethernet technology and is specifically designed for the data center. It is ideal for today's big data, cloud, and optimized workload solutions. It is an enterprise class Layer 2 and Layer 3 full-featured switch that delivers line-rate, high-bandwidth switching, filtering, and traffic queuing without delaying data. Large data-center grade buffers help keep traffic moving, while the redundant hot-swap power supplies and fans, along with numerous availability features, provide high availability for business sensitive traffic.

The RackSwitch G8296 is ideal for latency-sensitive applications, such as high performance computing clusters and financial applications. In addition to the 10 Gb Ethernet (GbE) and 40 GbE connections, the G8296 also has the capability for traditional 1 GbE connections. The G8296 supports the newest protocols including Data Center Bridging/Converged Enhanced Ethernet (DCB/CEE) for Fibre Channel over Ethernet (FCoE), in addition to iSCSI and network-attached storage (NAS).

Figure 4-49 shows the front view of the G8296 RackSwitch.



Figure 4-49 Lenovo RackSwitch G8296

## **G8296 key benefits**

The G8296 includes the following key benefits:

- ▶ High performance  
The 10 Gb/40 Gb switch provides the best combination of low latency, non-blocking line-rate switching, and ease of management. It has a throughput of 2.56 Tbps.
- ▶ Lower power and better cooling  
The RackSwitch G8296 uses as little as 210 W of power, and the front-to-rear or rear-to-front cooling design of the G8296 reduces data center air conditioning costs by having airflow match the servers in the rack. In addition, variable speed fans help to automatically reduce power consumption.
- ▶ Virtual Fabric  
The RackSwitch G8296 can help customers address I/O requirements for multiple NICs while reducing cost and complexity. Virtual Fabric allows you to carve a physical dual-port NIC into multiple virtual ports (between 2 - 8 vPorts) and to create a virtual pipe between the adapter and the switch for improved performance, availability, and security.
- ▶ VM-aware networking  
VMready software on the switch simplifies configuration and improves security in virtualized environments. VMready automatically detects virtual machine movement between physical servers and instantly reconfigures each VM's network policies across VLANs to keep the network up and running without interrupting traffic or impacting performance. VMready works with all leading VM providers, including VMware, Citrix Xen, and Microsoft Hyper-V.
- ▶ Layer 3 functionality  
The RackSwitch G8296 includes Layer 3 functionality, which provides security and performance benefits, as inter-VLAN traffic stays within the switch. This switch also provides the full range of Layer 3 protocols from static routes for technologies such as Open Shortest Path First (OSPF) and Border Gateway Protocol (BGP) for enterprise customers.
- ▶ Seamless interoperability  
RackSwitch switches perform seamlessly with other vendors' upstream switches.
- ▶ Fault tolerance  
RackSwitch switches learn alternate routes automatically and perform faster convergence in the unlikely case of a link, switch, or power failure. The switch uses proven technologies like L2 trunk failover, advanced VLAN-based failover, VRRP, and Hot Links.
- ▶ Multicast  
The RackSwitch G8296 supports IGMP Snooping v1, v2, and v3 with 2 K IGMP groups. It also supports Protocol Independent Multicast (PIM), such as PIM Sparse Mode or PIM Dense Mode.
- ▶ Management  
Network Address Translation (NAT) allows a single device, such as a router, to act as an agent between the Internet (or "public network") and a local (or "private") network. This means that only a single, unique IP address is required to represent an entire group of local IP addresses, to the external network. In addition to reducing the number of public IP addresses that are required, NAT also provides security by acting as a firewall between internal and external networks.

► Converged fabric

The RackSwitch G8296 switch supports CEE and connectivity to FCoE gateways. CEE helps enable clients to combine storage, messaging traffic, VoIP, video, and other data on a common data center Ethernet infrastructure. FCoE helps enable highly efficient block storage over Ethernet for consolidating server network connectivity. As a result, clients can deploy a single server interface for multiple data types, which can simplify both deployment and management of server network connectivity, while maintaining the high availability and robustness required for storage transactions.

► OpenFlow enabled

The RackSwitch G8296 offers benefits of OpenFlow. OpenFlow is an open API that enables the network administrator to easily configure and manage virtual networks that control traffic on a “per-flow” basis. It creates multiple independent virtual networks and related policies without dealing with the complexities of the underlying physical network and protocols. The RackSwitch G8296 is also the ideal switch to use with industry-compliant OpenFlow controllers.

**G8296 ports**

The G8296 features the following ports:

- 86x SFP/SFP+ port connectors to attach SFP/SFP+ transceivers for 1 Gb or 10 Gb Ethernet connections or DAC cables for 10 Gb Ethernet connections.
- 8x QSFP+ ports to attach QSFP+ transceivers or DAC cables for 40 Gb Ethernet connections (no 4x 10 GbE support).
- 2x QSFP+ ports to attach QSFP+ transceivers or DAC cables for 40 Gb or 4x 10 Gb Ethernet connections.
- One RJ-45 10/100/1000 Mb Ethernet port for out-of-band management.
- One mini-USB RS-232 console port for configuring the switch.
- One USB port for mass storage devices.

Figure 4-50 shows the ports on the RackSwitch G8296.




Figure 4-50 G8296 ports

For more information about the G8296, see *Lenovo RackSwitch G8296*, TIPS1266, which is available at this website:

<https://lenovopress.com/tips1266>

## 4.7.9 Lenovo RackSwitch G8332

The RackSwitch G8332 provides low latency, lossless performance, and a feature-rich design with key virtualization features, such as CEE/DCB, high availability, and enterprise class Layer 2 and Layer 3 functions. The RackSwitch G8332 has 32 QSFP+ ports and is suitable for clients that use 10 Gigabit Ethernet or 40 Gigabit Ethernet connectivity (or both).

With latency below 600 nanoseconds, the RackSwitch G8332 is among industry leaders. This low latency rate and the 2.56 Tbps throughput makes the RackSwitch G8332 an ideal offering for latency-sensitive applications, such as high-performance computing, financial applications, hosting environments, and cloud designs in Enterprise, Web 2.0, and VM Centric data centers. In addition, the G8332 supports VMready with Virtual Vision, which enables the network to be VM aware and provides the capability to have a virtualization environment that is simpler and less expensive with exceptional performance.

Figure 4-51 shows the front view of the G8332 RackSwitch.



Figure 4-51 *Lenovo RackSwitch G8332*

### **G8332 key benefits**

The G8332 includes the following key benefits:

► High performance

This 40 GbE low latency switch with 2.56 of Tbps throughput provides the best combination of low latency, non-blocking switching, and ease of management. The RackSwitch G8332 is also a single ASIC design, which ensures consistent lower port-to-port latency.

► Lower power and better cooling

The RackSwitch G8332 uses as little as 270 W of power, which is a fraction of the power consumption of previous 40 GbE offerings. The front-to-rear or rear-to-front cooling design reduces data center air conditioning costs by having airflow match the servers in the rack. In addition, variable speed fans help reduce power consumption.

► High availability

The G8332 comes standard with hot-swap redundant power supplies and fans, which makes the switch highly reliable and easy to service if there is a failure.

- ▶ VM-aware networking

VMready software on the switch simplifies configuration and improves security in virtualized environments. VMready automatically detects VM movement between physical servers and instantly reconfigures each VM's network policies across VLANs to keep the network running without interrupting traffic or affecting performance. VMready works with all leading VM providers, including VMware, Microsoft Hyper-V, and Red Hat KVM.

- ▶ Layer 3 functionality

The switch includes Layer 3 functionality, which provides security and performance benefits, as inter-VLAN traffic stays within the switch. This switch also provides the full range of Layer 3 protocols from static routes for technologies, such as OSPF and BGP for enterprise customers.

- ▶ Seamless interoperability

The G8332 switch interoperates seamlessly with the upstream switches of other vendors.

- ▶ Fault tolerance

The G8332 switch learns alternative routes automatically and performs faster convergence if there is a link, switch, or power failure. The RackSwitch G8332 uses technologies, such as L2 trunk failover, VLAN-based failover, VRRP, and Hot Links.

- ▶ OpenFlow enabled

The RackSwitch G8332 offers the benefits of OpenFlow. OpenFlow is the new open API that enables the network administrator to easily configure and manage virtual networks that control traffic on a "per-flow" basis. It creates multiple independent virtual networks and related policies without dealing with the complexities of the underlying physical network and protocols.

- ▶ Multicast

Multicast supports IGMP Snooping v1, v2, and v3 with 3K IGMP groups, and PIM, such as PIM Sparse Mode or PIM Dense Mode.

- ▶ Management

NAT allows a single device, such as a router, to act as an agent between the Internet (or "public network") and a local (or "private") network. By using this configuration, only a single, unique IP address is required to represent an entire group of local IP addresses to the external network. NAT also provides security by acting as a firewall between internal and external networks.

### ***G8332 ports***

The G8332 features the following ports:

- ▶ A total of 32 ports for 40 Gb Ethernet QSFP+ transceivers (40GBASE-SR4 or 40GBASE-LR4), QSFP+ to QSFP+ DAC cables (40GBASE-CR4), or QSFP+ to 4x 10 Gb SFP+ break-out cables (ports 2 - 25 only).
- ▶ One 10/100/1000 Ethernet port (RJ-45 connector) for OOB management.
- ▶ One RS-232 serial port (mini-USB connector) that provides another means to configure the switch.
- ▶ One USB port for mass storage devices.

Figure 4-52 shows the ports on the RackSwitch G8332.




Figure 4-52 G8332 ports

For more information about the G8332, see *Lenovo RackSwitch G8332, TIPS1274*, which is available at this website:

<http://lenovopress.com/tips1274>



# Abbreviations and acronyms

<b>ABAP</b>	Advanced Business Application Programming	<b>IMM2</b>	Integrated Management Module II
<b>AES</b>	Advanced Encryption Standard	<b>IOPS</b>	I/O operations per second
<b>AFA</b>	All-Flash-Array	<b>KVM</b>	Kernel-based Virtual Machine
<b>API</b>	application programming interface	<b>LETS</b>	Lenovo Expert Technical Sales
<b>APIC</b>	advanced programmable interrupt controller	<b>MCA</b>	Machine Check Architecture
<b>APO</b>	Advanced Planner and Optimizer	<b>MSCS</b>	Microsoft Cluster Server
<b>ASCS</b>	ABAP central services instance	<b>MTTR</b>	mean time to resolution
<b>AVX</b>	Advanced Vector Extensions	<b>NAS</b>	network-attached storage
<b>AVX2</b>	Advanced Vector Extensions 2	<b>NAT</b>	Network Address Translation
<b>BGP</b>	Border Gateway Protocol	<b>NUMA</b>	non-uniform memory access
<b>CEE</b>	Converged Enhanced Ethernet	<b>NVMe</b>	Non-Volatile Memory Express
<b>CMMs</b>	Chassis Management Modules	<b>OOB</b>	out of band
<b>CRC</b>	cyclic redundancy checking	<b>OS</b>	operating system
<b>CRM</b>	Customer Relationship Management	<b>OSPF</b>	Open Shortest Path First
<b>CTIO</b>	Cut Through I/O	<b>PCH</b>	PCIe Controller Hub
<b>DAC</b>	direct-attach copper	<b>PCIe</b>	PCI Express
<b>DB</b>	database	<b>PFA</b>	Predictive Failure Alert
<b>DCS</b>	Data Center Services	<b>PIM</b>	Protocol Independent Multicast
<b>DDDC</b>	Double Device Data Correction	<b>PLM</b>	Product Lifecycle Management
<b>DEV</b>	Development system	<b>PRD</b>	production system
<b>DWPD</b>	drive writes per day	<b>QA</b>	Quality Assurance
<b>ECC</b>	Error Correction Code	<b>QPI</b>	quick path interconnect
<b>EPM</b>	Enterprise Performance Management	<b>RAS</b>	reliability, availability, and serviceability
<b>ERP</b>	Enterprise Resources Planning	<b>RDIMMs</b>	registered memory DIMMs
<b>ERS</b>	Enqueue Replication Server instance	<b>RFID</b>	radio frequency identification
<b>FAMM</b>	full array memory mirroring	<b>RHEL</b>	Red Hat Enterprise Linux
<b>FC</b>	Fibre Channel	<b>RHEV</b>	Red Hat Enterprise Virtualization
<b>FCoE</b>	Fibre Channel over Ethernet	<b>ROC</b>	RAID-on-Chip
<b>GPU</b>	Graphics Processing Unit	<b>ROI</b>	return on investments
<b>GRC</b>	governance, risk, and compliance	<b>SAN</b>	Storage Area Networks
<b>GUI</b>	graphical user interface	<b>SAPS</b>	SAP Application Performance Standard
<b>GbE</b>	Gigabit Ethernet	<b>SCM</b>	Supply Chain Management
<b>HCL</b>	Hardware Compatibility List	<b>SCS</b>	services instance
<b>HDD</b>	hard disk drive	<b>SDDC</b>	Single Device Data Correction
<b>ICM</b>	Internet Communication Manager	<b>SDM</b>	Software Deployment Manager
<b>IGS</b>	Internet Graphics Service	<b>SLES</b>	SUSE Linux Enterprise Server
<b>IIO</b>	Integrated I/O	<b>SMI</b>	Scalable memory interconnect
		<b>SMI2</b>	Scalable Memory Interconnect generation 2

<b>SNMPv3</b>	Simple Network Management Protocol v3
<b>SOA</b>	service-oriented architecture
<b>SPOF</b>	single point of failure
<b>SRM</b>	Supplier Relationship Management
<b>SSD</b>	solid-state drive
<b>TCA</b>	total cost of acquisition
<b>TCO</b>	total cost of ownership
<b>TPC</b>	Transaction Processing Performance Council
<b>TRN</b>	Training system
<b>TSX</b>	Transactional Synchronization eXtensions
<b>UEFI</b>	Unified Extensible Firmware Interface
<b>UIM</b>	Upward Integration Module
<b>VM</b>	virtual machine
<b>VMs</b>	virtual machines
<b>WfMC</b>	Workflow Management Coalition
<b>eMCA</b>	Enhanced MCA
<b>eMCA2</b>	Enhanced Machine Check Architecture Gen2
<b>eMLC</b>	enterprise-grade MLC

# Related publications

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

## Lenovo Press publications

For more information, see the following publication, which might be available in softcopy only:

- ▶ *Lenovo System x3850 X6 and x3950 X6 Planning and Implementation Guide*, SG24-8208
- ▶ *Flex System Products and Technology*, SG24-8255
- ▶ *Lenovo Flex System X6 Compute Node Planning and Implementation Guide*, SG24-8227
- ▶ *In-memory Computing with SAP HANA on Lenovo X6 Systems*, SG24-8086

You can search for, view, download or order these documents and other books, papers, product guides, draft, and other materials at this website:

<http://www.lenovopress.com>

## Online resources

The following websites are also relevant as further information sources:

- ▶ Lenovo System x solutions for SAP environments:  
<http://www.lenovo.com/solutions/sap/>
- ▶ SAP on Linux:  
<http://scn.sap.com/community/linux>
- ▶ SAP on Microsoft Windows:  
<http://scn.sap.com/community/windows>
- ▶ SAP Standard Application Benchmarks:  
<http://www.sap.com/benchmark>
- ▶ SAP Certification: Lenovo systems:  
<http://www.saponwin.com/pub/hardware.asp?l=vendor&s1=43&i=43&la=en>

