

*“All-in-One Is All You Need.”*

ALL-IN-ONE

CompTIA<sup>®</sup>  
**Server+**<sup>™</sup>  
Certification

EXAM GUIDE  
EXAM SK0-005

SECOND EDITION

Save 10% of any  
CompTIA exam  
voucher! Coupon  
code inside.

Online content  
includes:

- 200 practice exam questions
- Interactive performance-based questions
- Test engine that provides full-length practice exams and customizable quizzes by chapter or exam objective

*Complete coverage  
of all objectives for  
exam SK0-005*

*Ideal as both a study  
tool and an  
on-the-job reference*

*Filled with practice exam  
questions and in-depth  
explanations*

**Mc  
Graw  
Hill**

**DANIEL LACHANCE**

CompTIA Server+, CompTIA A+<sup>™</sup>, CompTIA Security+<sup>™</sup>,  
CompTIA Network+<sup>™</sup>, CompTIA Cloud+<sup>™</sup>

ALL ■ IN ■ ONE

CompTIA  
**Server+**<sup>TM</sup>  
Certification

EXAM GUIDE

Second Edition (Exam SK0-005)

---

## ABOUT THE AUTHOR

**Daniel Lachance**, CompTIA Security+, CompTIA A+, CompTIA Network+, CompTIA Server+, CompTIA Cloud Essentials, CompTIA Cloud+, as well as various Microsoft Azure and Amazon Web Services certifications, is the owner of Lachance IT Consulting, Inc., based in Halifax, Nova Scotia, Canada. He is the author of *CompTIA Cloud Essentials+ Certification Study Guide, Second Edition*, and co-author of *CompTIA Security+ Certification Practice Exams*.

Since the early 1990s, he has worked in various capacities as a computer programmer, network and server technician, and security analyst. He is also an experienced trainer, having delivered IT training online, in Canada, and in the Caribbean since the 1990s on topics ranging from Microsoft enterprise products (Active Directory, Hyper-V, System Center Configuration Manager, and Azure) to Amazon Web Services, UNIX, Linux, security, and networking.

He has recorded tech support videos for products such as Microsoft Azure, Amazon Web Services, and Microsoft System Center Configuration Manager and has recorded videos covering various cybersecurity and mobility topics.

He enjoys spending time with his spouse, Tammy; their children, Roman, Trinity, Abby, and Jacob; families and friends; and the family dogs, Dori and Louis. He also enjoys jogging, reading nonfiction, and listening to and playing various styles of music.

### About the Technical Editor

**S. Russell Christy** is a technical trainer in Memphis, Tennessee covering a wide variety of products and specializing in computer maintenance, network and security, and Microsoft Office applications. For over 20 years, he has deployed new desktops and operating systems, servers, and network hardware and software, while simultaneously troubleshooting various hardware and software issues.

He holds a bachelor's degree in business administration from the University of Memphis. He has also gained industry certifications in CompTIA A+, CompTIA Network+, CompTIA Server+, CompTIA Security+, CompTIA CySA+, Cisco CCNA CyberOps, MTA Windows Server Administration Fundamentals, Network Fundamentals, Security Fundamentals, Windows OS Fundamentals, Microsoft 365 Identity and Services, and Adobe Education Trainer.

ALL ■ IN ■ ONE

CompTIA  
**Server+™**  
Certification

EXAM GUIDE

Second Edition (Exam SK0-005)

Daniel Lachance



New York Chicago San Francisco  
Athens London Madrid Mexico City  
Milan New Delhi Singapore Sydney Toronto

McGraw Hill is an independent entity from CompTIA® and is not affiliated with CompTIA in any manner. This publication and accompanying media may be used in assisting students to prepare for the CompTIA Server+™ exam. Neither CompTIA nor McGraw Hill warrants that use of this publication and accompanying media will ensure passing any exam. CompTIA and CompTIA Server+ are trademarks or registered trademarks of CompTIA in the United States and/or other countries. All other trademarks are trademarks of their respective owners. The CompTIA Marks are the proprietary trademarks and/or service marks of CompTIA and its affiliates used under license from CompTIA.

Copyright © 2021 by McGraw Hill. All rights reserved. Except as permitted under the United States Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of the publisher, with the exception that the program listings may be entered, stored, and executed in a computer system, but they may not be reproduced for publication.

ISBN: 978-1-26-046992-9

MHID: 1-26-046992-1

The material in this eBook also appears in the print version of this title: ISBN: 978-1-26-046991-2,

MHID: 1-26-046991-3.

eBook conversion by codeMantra

Version 1.0

All trademarks are trademarks of their respective owners. Rather than put a trademark symbol after every occurrence of a trademarked name, we use names in an editorial fashion only, and to the benefit of the trademark owner, with no intention of infringement of the trademark. Where such designations appear in this book, they have been printed with initial caps.

McGraw-Hill Education eBooks are available at special quantity discounts to use as premiums and sales promotions or for use in corporate training programs. To contact a representative, please visit the Contact Us page at [www.mhprofessional.com](http://www.mhprofessional.com).

Information has been obtained by McGraw Hill from sources believed to be reliable. However, because of the possibility of human or mechanical error by our sources, McGraw Hill, or others, McGraw Hill does not guarantee the accuracy, adequacy, or completeness of any information and is not responsible for any errors or omissions or the results obtained from the use of such information.

## TERMS OF USE

This is a copyrighted work and McGraw-Hill Education and its licensors reserve all rights in and to the work. Use of this work is subject to these terms. Except as permitted under the Copyright Act of 1976 and the right to store and retrieve one copy of the work, you may not decompile, disassemble, reverse engineer, reproduce, modify, create derivative works based upon, transmit, distribute, disseminate, sell, publish or sublicense the work or any part of it without McGraw-Hill Education's prior consent. You may use the work for your own noncommercial and personal use; any other use of the work is strictly prohibited. Your right to use the work may be terminated if you fail to comply with these terms.

THE WORK IS PROVIDED "AS IS." MCGRAW-HILL EDUCATION AND ITS LICENSORS MAKE NO GUARANTEES OR WARRANTIES AS TO THE ACCURACY, ADEQUACY OR COMPLETENESS OF OR RESULTS TO BE OBTAINED FROM USING THE WORK, INCLUDING ANY INFORMATION THAT CAN BE ACCESSED THROUGH THE WORK VIA HYPERLINK OR OTHERWISE, AND EXPRESSLY DISCLAIM ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. McGraw-Hill Education and its licensors do not warrant or guarantee that the functions contained in the work will meet your requirements or that its operation will be uninterrupted or error free. Neither McGraw-Hill Education nor its licensors shall be liable to you or anyone else for any inaccuracy, error or omission, regardless of cause, in the work or for any damages resulting therefrom. McGraw-Hill Education has no responsibility for the content of any information accessed through the work. Under no circumstances shall McGraw-Hill Education and/or its licensors be liable for any indirect, incidental, special, punitive, consequential or similar damages that result from the use of or inability to use the work, even if any of them has been advised of the possibility of such damages. This limitation of liability shall apply to any claim or cause whatsoever whether such claim or cause arises in contract, tort or otherwise.

---

# CONTENTS AT A GLANCE

<b>Chapter 1</b>	Introduction to CompTIA Server+ Essentials.....	1
<b>Chapter 2</b>	Server Hardware.....	27
<b>Chapter 3</b>	Server Operating Systems and Server Roles .....	71
<b>Chapter 4</b>	Server Storage .....	133
<b>Chapter 5</b>	Server Network Communications.....	173
<b>Chapter 6</b>	Server and Network Security .....	221
<b>Chapter 7</b>	Troubleshooting and Performance Optimization .....	277
<b>Chapter 8</b>	Disaster Recovery Planning.....	325
<b>Appendix A</b>	About the Hands-on Exercises and Lab Setup.....	361
<b>Appendix B</b>	Objective Map .....	369
<b>Appendix C</b>	About the Online Content.....	373
	Glossary .....	377
	Index.....	405

*This page intentionally left blank*

---

# CONTENTS

	Acknowledgments .....	xv
	Introduction .....	xvii
<b>Chapter 1</b>	<b>Introduction to CompTIA Server+ Essentials .....</b>	<b>1</b>
	Why This Book Is Relevant .....	1
	Server Hardware Basics .....	2
	Server Form Factors .....	3
	Server CPUs .....	3
	Server Memory .....	3
	Buses and Slots .....	4
	Virtualization .....	4
	Environmental Factors .....	5
	Storage .....	6
	Network Concepts .....	7
	Cables and Connectors .....	7
	IP .....	7
	Ports and Protocols .....	7
	Server Operating Systems and Server Roles .....	8
	Server Roles .....	9
	Maintenance .....	9
	Monitoring .....	10
	Security Considerations .....	10
	Hardening .....	11
	Network Security .....	11
	Troubleshooting and Optimizing Performance .....	13
	Optimizing Performance .....	13
	Troubleshooting .....	13
	Preparing for the Worst .....	15
	Cloud Computing .....	15
	Cloud Computing Characteristics .....	15
	Cloud Computing Deployment Models .....	16
	Cloud Computing Service Models .....	17
	Chapter Review .....	18
	Questions .....	18
	Questions and Answers .....	21



<b>Chapter 2</b>	<b>Server Hardware</b>	<b>27</b>
	Server Form Factors	27
	Tower Servers	28
	Rack-mounted Equipment	29
	Blade Servers	32
	Server Components	33
	BIOS	34
	UEFI	35
	CPUs	36
	GPUs	38
	Memory	39
	Buses	40
	NICs	43
	Storage	43
	USB	44
	Power	46
	Voltage	46
	Wattage	48
	Uninterruptible Power Supply	48
	Environmental Controls	50
	Temperature	50
	Airflow	51
	Humidity	51
	Fire Suppression	51
	Hands-on Exercises	53
	Chapter Review	56
	Server Form Factors	56
	Server Components	56
	Power	57
	Environmental Controls	57
	Questions	57
	Questions and Answers	62
<b>Chapter 3</b>	<b>Server Operating Systems and Server Roles</b>	<b>71</b>
	Server Roles	71
	Infrastructure Roles	72
	Other Server Roles	80
	Virtualization Servers	85
	Hypervisor Types	85
	Hypervisor Host Configuration	86
	Virtual Machine Guest Configuration	87
	Server Installation	89
	Server Licensing	90
	Installing a Type 1 Hypervisor	90
	Installing a Server Operating System	91

Server Administration Methods	94
KVM	95
Out-of-Band Remote Administration	95
In-Band Remote Administration	96
Server Documentation	100
Asset Life Cycle	101
Asset Inventory	101
Service Level Agreements	104
Other Documentation	104
Maintaining Servers	105
Patch Management	105
Proactive Maintenance	106
Reactive Maintenance	106
Hands-on Exercises	106
Chapter Review	116
Server Roles	116
Virtualization Servers	117
Server Installation	117
Server Administration Methods	118
Server Documentation	118
Server Maintenance	119
Questions	119
Questions and Answers	124
<b>Chapter 4</b> Server Storage	133
Storage Technologies	133
Magnetic Hard Disks	134
Solid-State Drives	134
Hybrid Drives	136
Storage Tiers	136
Disk Interfaces	137
Optical Drives	138
Cloud Storage	139
Direct-Attached Storage	139
Network-Attached Storage	140
Storage Area Networks	142
Storage Capacity and Future Growth	143
Base 2 vs. Base 10	144
Where Did All the Disk Space Go?	144
Using Less Disk Space	145
RAID Configuration	148
RAID Levels	149
Storage Device Installation	152
MBR and GPT	152
File Systems	153

	Sample Scenario 1 .....	155
	Sample Scenario 2 .....	155
	Hands-on Exercises .....	156
	Chapter Review .....	160
	Storage Device Characteristics .....	160
	Disk Interfaces .....	161
	Local and Network Storage .....	161
	Storage Capacity Planning .....	161
	RAID .....	161
	Disk Initialization and File Systems .....	162
	Questions .....	162
	Questions and Answers .....	166
<b>Chapter 5</b>	<b>Server Network Communications .....</b>	<b>173</b>
	The OSI Model .....	173
	Cable Installation and Management .....	175
	Cable Placement .....	175
	Cable Labeling .....	176
	Cable Types .....	177
	Network Hardware .....	182
	Network Interface Cards .....	182
	Network Switches .....	184
	Routers .....	186
	Configuring IPv4 .....	187
	PAT .....	187
	Static NAT .....	188
	IP .....	188
	IPv4 Addressing .....	188
	Subnet Mask .....	189
	Reserved Internal IP Address Ranges .....	189
	When to Use Subnetting .....	190
	Configuring IPv6 .....	193
	IPv6 Addressing .....	193
	IPv6 Settings .....	194
	IPv6 Transition Technologies .....	195
	Network Infrastructure Services .....	195
	Default Gateway .....	195
	DNS Servers .....	196
	WINS Servers .....	197
	DHCP .....	197
	TCP and UDP .....	198
	TCP .....	198
	UDP .....	199
	Hands-on Exercises .....	202

Chapter Review	207
The OSI Model	207
Cables and Connectors	207
Network Interface Cards	207
IPv4 and IPv6	208
IP, TCP, and UDP	208
Questions	208
Questions and Answers	212
<b>Chapter 6</b> Server and Network Security	221
Physical Security Measures	221
Premises Access	221
The Human Element	223
Authentication	224
Under Lock and Key	228
Logical Access Control	228
Groups	229
Dynamic Access Control	229
Roles	231
Rights and Permissions	231
File System Permissions	231
Peripheral Devices	235
Network Security	237
NAC	237
VLANs	238
Firewalls	239
Security Zones	243
PKI	244
IPSec	246
VPNs	248
Intrusion Detection and Prevention Systems	248
Hardening	249
Operating System Hardening	250
Hardware Hardening	251
Application Hardening	251
Data Security	252
Data and Mobile Devices	253
Encrypting Data at Rest	253
Tape Encryption	256
Secure Media Disposal	257
Hands-on Exercises	259
Chapter Review	264
Physical Security	264
Authentication	264
Logical Access Control	264

	Network Security .....	264
	Firewalls .....	265
	PKI .....	265
	IPSec .....	265
	VPNs .....	265
	Intrusion Detection and Prevention .....	266
	Hardening .....	266
	Data Security .....	266
	Secure Media Disposal .....	266
	Questions .....	266
	Questions and Answers .....	270
<b>Chapter 7</b>	<b>Troubleshooting and Performance Optimization .....</b>	<b>277</b>
	Troubleshooting Methodology .....	277
	Identify the Problem .....	278
	Establish a Theory of Probable Cause .....	282
	Test the Theory .....	282
	Establish a Plan of Action .....	282
	Implement a Solution or Escalate .....	283
	Verify Functionality .....	283
	Perform Root Cause Analysis .....	283
	Document the Solution .....	284
	Hardware Problems and Solutions .....	285
	Software Problems and Solutions .....	285
	Storage Problems and Solutions .....	292
	Windows Tools .....	292
	Linux Tools .....	294
	Network Problems and Solutions .....	296
	Name Resolution Issues .....	298
	Security Problems and Solutions .....	299
	Malware Troubleshooting .....	299
	Too Few Permissions .....	301
	Too Many Permissions .....	303
	Too Much Running .....	303
	Confidentiality and Integrity .....	304
	Performance Optimization .....	305
	Hardware Optimization .....	306
	Software Optimization .....	306
	Network Optimization .....	306
	Hands-on Exercises .....	307
	Chapter Review .....	310
	Troubleshooting Methodology .....	310
	Hardware Problems and Solutions .....	311
	Software Problems and Solutions .....	311
	Storage Problems and Solutions .....	312

	Network Problems and Solutions .....	312
	Security Problems and Solutions .....	312
	Performance Optimization .....	313
	Questions .....	313
	Questions and Answers .....	317
<b>Chapter 8</b>	<b>Disaster Recovery Planning .....</b>	<b>325</b>
	Disaster Recovery .....	325
	Alternate Sites .....	325
	Data Replication .....	329
	Business Impact Analysis .....	332
	Business Continuity .....	334
	Disaster Recovery Plan .....	334
	Business Continuity Plan .....	335
	Data Backup .....	335
	Backup Types .....	336
	Backup Media .....	340
	On-premises Backup .....	341
	Cloud Backup .....	342
	Backup and Restore Best Practices .....	343
	Hands-on Exercises .....	344
	Chapter Review .....	348
	Disaster Recovery Sites .....	348
	Data Replication .....	348
	Business Impact .....	349
	Disaster Recovery Plan .....	349
	Business Continuity Plan .....	349
	Data Backups .....	350
	Questions .....	351
	Questions and Answers .....	354
<b>Appendix A</b>	<b>About the Hands-on Exercises and Lab Setup .....</b>	<b>361</b>
	Lab Exercise Overview .....	361
	Microsoft Windows .....	361
	Linux .....	362
	Requirements for Hands-on Exercises .....	362
	Hardware Requirements .....	363
	VMware Workstation 16 Pro .....	363
	Acquiring Windows and Linux Installation Media .....	365
	Lab Exercise Miscellany .....	366
<b>Appendix B</b>	<b>Objective Map .....</b>	<b>369</b>
	Exam SK0-005 .....	369
	CompTIA Server+ Certification Exam SK0-005 .....	369

<b>Appendix C About the Online Content</b> .....	373
System Requirements .....	373
Your Total Seminars Training Hub Account .....	373
Privacy Notice .....	373
Single User License Terms and Conditions .....	373
TotalTester Online .....	375
Performance-Based Questions .....	375
Technical Support .....	375
 Glossary .....	 377
 Index .....	 405

---

## ACKNOWLEDGMENTS

With each book project, I am reminded of the immense collective skill from the pool of talented folks required to make it all happen. I would like to thank McGraw Hill for providing the opportunity to create this book. Emily Walters, Lisa McClain, and Rachel Fogelberg were the guiding lights as to what I should be doing and when I should be doing it. I would also like to thank the rest of the wonderful team at McGraw Hill for making this project a great experience, resulting in a solid and fun book.

This book is chock-full of technical goodies, but credit must be given to S. Russell Christy for keeping the technical content consistent, accurate, and relevant—thanks, Russ!

I would like to say a special thank you to my family for enduring the many hours I spent preparing for and creating this book, as well as the endless techno-rambling I tend to fall prey to—love you guys!



*This page intentionally left blank*

---

# INTRODUCTION

The CompTIA Server+ certification exam is becoming increasingly popular and showing up in IT job listings as a desired certification, not only for the management of physical servers but also the management of virtual servers running on premises and in the cloud. Topics formerly included in the discontinued CompTIA Storage+ certification are now included in CompTIA Server+. This change makes sense because of the close relationship servers have with storage, especially network storage.

CompTIA Server+ certification candidates must demonstrate a solid understanding of topics ranging from server hardware to server operating systems, networking and security, storage, virtualization, and cloud-based solutions. Today's server administrators are responsible for remote administration of potentially hundreds of physical and virtual servers, whether in a small server room or in a large data center.

The primary skill required by CompTIA Server+ technicians is effective troubleshooting. Yes, it's important to understand concepts and the sequence of steps required to yield a desired result, but the timely and efficient resolution of server-related issues is what distinguishes a casual server technician from an outstanding server technician. This book includes entire chapters dedicated to troubleshooting and disaster recovery.

## Exam Details

You can book your English-language CompTIA Server+ SK0-005 exam online at [www.pearsonvue.com/comptia/](http://www.pearsonvue.com/comptia/), or by phone with your local Pearson Vue testing center as of May 18, 2021. Additional languages will be added at a later date. Consult the web site to find the testing center nearest you, including directions to the center, or you can opt to have the test proctored for you in the comfort of your office or home.

You can use an exam voucher if you have one; either way, the cost is USD338. Payment methods include various debit and credit cards.

## Preparing for the Exam

People learn in different ways. Reading comprehension is important, especially when it comes to end-of-chapter and master exam questions included with this book. Some folks learn best by doing; that's why you'll find hands-on exercises using both Windows and Linux at the end of each chapter.

It's been proven that doing something related to a discussed topic helps our brains to understand how the topic is relevant and thus aids in retention.

The following list provides some general test-taking tips:

- Have a positive attitude.
- Learn the material in earnest; it is interesting and it can help your career.

- Get a full night's sleep for a few days prior to taking the exam.
- Eat properly before taking the exam.
- Don't cram right before the exam.
- Arrive at the testing center 15 minutes early.
- Read all the questions and *all the answer choices* thoroughly.
- With complex questions, ask yourself, "What is the question really asking?"

## Taking the Exam at Your Location

When you schedule the exam online, you can choose to take the exam from anywhere you have an Internet connection and a computer with a microphone and web camera. You will receive instructions regarding how to prepare your exam room, and you will be prompted to upload a photo ID as well as pictures of the exam room. You will also receive instructions for ensuring that your computer meets the requirements for remote exam-taking. When the exam begins, the exam proctor will most likely ask you to show him or her your exam room through your computer's web camera. You will not be allowed to have a watch on, phones within reach, or refreshments on your desk, nor will you be granted restroom breaks. Importantly, being interrupted by others could forfeit your exam, so make sure that you are completely isolated during the exam.

## Taking the Exam at the Testing Center

After booking your exam, if you will be taking the exam at a testing center, take the time to ensure that you know how to get to the center on time, including the schedule of any public transportation you might take, and find out where you can park if you plan on driving.

You will need two pieces of ID at the testing center, including a photo ID. The testing center will also take a picture of you for your exam score report. You are not allowed to have a smartphone, books, or notes in the testing room.

## Making the Grade

The required passing score for the CompTIA SK0-005 exam is 750 (the range is 100 to 900), and the exam consists of 90 multiple-choice and performance-based questions that you must answer within 90 minutes. Performance-based questions, or PBQs, test more than just your knowledge of CompTIA Server+ topics; they test that you can apply your Server+ knowledge to scenarios to solve problems. As has always been the case, there might be questions included on the exam that do not count toward your score. CompTIA does this occasionally to test out new content or question formats, but you won't know which question(s) this applies to.

Table 1 lists the exam domains and their weight against the overall exam. Notice the emphasis on server administration and troubleshooting. For a detailed breakdown of CompTIA Server+ SK0-005 exam objectives, visit [www.certification.comptia.org/certifications/server](http://www.certification.comptia.org/certifications/server).

Exam Domain	Percentage of the Exam
1.0 Server Hardware Installation and Management	18%
2.0 Server Administration	30%
3.0 Security and Disaster Recovery	24%
4.0 Troubleshooting	28%
<b>Total</b>	<b>100%</b>

**Table 1** CompTIA Server+ SK0-005 Exam Domains

If you are taking the exam at a testing center, you can use a provided marker and laminated paper to write notes and calculate values, such as valid subnet ranges. Each question has a Prev (previous) and Next button you can click to allow navigation throughout the entire exam. You can also flag certain questions for review; an item review screen will appear after the last question, where you can review some or all of your answers before ending the exam.

The exam contains multiple-choice questions that ask you to choose one or more correct answers, and you can expect scenario performance-based questions to solve some kind of server-based problem.

When studying for this exam, always keep the following in mind:

- What is the most efficient way of completing an IT configuration task?
- What sequence of steps must be followed to achieve a goal?
- What is the first step in troubleshooting a specific issue?
- What can be done to improve the performance or security of an IT system?

You will know right away whether you've passed the exam. If you fail on your first attempt to pass the exam, CompTIA does not require a waiting period before making your second attempt at the exam, but you will have to wait 14 calendar days if you want to make a third attempt. Using this book properly should not require a second exam attempt.

## Using the Book

Welcome to *CompTIA Server+ Certification All-in-One Exam Guide, Second Edition (Exam SK0-005)*! By reading and understanding each chapter (including the questions), completing the included lab exercises, and taking the companion online practice exams, you will greatly increase your likelihood of passing the CompTIA Server+ exam. Refer to Appendix A for information about setting up the hands-on lab exercises.

This book maps to the official CompTIA SK0-005 exam objectives to help you prepare for the exam, but it also goes beyond the objectives to provide insight that will prove valuable when you are working in server-related environments.

## Using the Objectives Map

The objectives map included in Appendix B has been constructed to help you cross-reference the official exam objectives from CompTIA with the relevant coverage in the book. The objectives map lists each exam objective exactly as CompTIA presents it and provides the corresponding chapter and section in which that objective is covered.

## Online Practice Exams

This book includes access to practice exams that feature the Total Tester Online exam software test engine, which allows you to generate a complete practice exam or to generate quizzes by chapter module or by exam domain. This book also includes access to ten performance-based questions. For more information about the accompanying software and instructions on how to access the exam tool, see Appendix C.

# Introduction to CompTIA Server+ Essentials

In this chapter, you will

- Learn about server hardware components
- Learn how storage systems can be configured and provisioned
- Review basic network concepts
- Review the basics of monitoring and maintaining server operating systems
- Review basic security concepts
- Learn about troubleshooting and optimizing performance
- Learn about disaster recovery
- Explore cloud computing concepts

This chapter provides an overview of topics (all of the details follow in Chapters 2–8) that you're sure to see on the CompTIA Server+ Certification Exam (SK0-005). We'll also cover these topics in such a way that you'll be well armed working in the IT field, whether you're discussing network storage for cloud-based virtualized servers, determining specific server hardware that must be ordered, or working in a data center.

Data centers are neat—these big facilities host a wide array of IT services that are consumed by clients located hundreds, even thousands, of miles away. This is important with cloud computing. The equipment that cloud services run on is housed in data centers. You can bet that, among other certifications, a CompTIA Server+ certification is often preferred for data center jobs.

## Why This Book Is Relevant

Because you're reading this book, you probably already see the value in learning about the CompTIA Server+ Certification Exam. CompTIA is respected globally in the IT industry for its certifications for A+, Network+, Server+, Linux+, and many others. If you check out IT job listings on your favorite job-hunting site or ask people working at IT academic institutions, you'll learn that many jobs require CompTIA certifications in one form or another.

Being a server expert is much different today from what it was in the 1990s, and that's going back only 25 years or so! If you've been working in IT for a while, you might agree that there was a time when

- We could know everything about a server operating system.
- Arguably, overall server support was *simpler* because the software wasn't doing as many things.
- Arguably, overall server support was *harder* because we didn't have great Internet search engines and video tutorials.
- Applying patches didn't occupy nearly as much of our time.
- Malware infections weren't nearly as ubiquitous.
- Storage was contained physically inside the server case.
- IPv4 was groundbreaking.

The CompTIA certification exam helps ensure that candidates really know what they're talking about, including much more than just knowledge of servers themselves. You need to be familiar with server types and components, virtualization, IPv4 and IPv6 networking, cloud computing, operating systems, network storage, security, and troubleshooting. If all this interests you (and you need this certification to get a job!), then you are in the right place.

What kinds of jobs relate to the CompTIA Server+ certification? Countless, but let's list a few general categories:

- Data center IT technician
- Server administrator
- Network technician
- Help desk technician

Remember that an understanding of the body of knowledge presented in this book is crucial for IT server and network technicians working for companies in any industry.

## Server Hardware Basics

Everybody uses servers in one way or another, even if they don't realize it—including teenagers posting content to social media from their smartphones, shoppers making online purchases, companies backing data up to the cloud, and Wall Street executives making decisions based on server-supplied data. Although the role of a server (the centralized serving of content and services to concurrent users over a network) hasn't changed much over the decades, the scale of clients demanding quick and reliable access, the amount of data processed, as well as how servers are implemented, *have changed*. It's not enough to know about servers themselves; we need to be aware of the entire ecosystem.

Servers offer some kind of a service to clients over the network, such as a web server offering a web site to a client's web browser. The server's role and number of clients it serves dictate how much horsepower is needed. The interesting part is that this is true whether you're working with physical or virtual servers, on-premises or in the cloud.

The CompTIA Server+ Certification Exam will test your knowledge on the best hardware configurations given specific server scenarios. Even virtual machine servers, otherwise called *instances*, deployed in the cloud can be configured using different instance types, or sizes, which consist of factors like the number of virtual CPU cores (vCPUs), the amount of RAM, the number of attached disks, and so on. The great thing about virtual machines is that you can scale down, or reduce the amount of horsepower when appropriate, and in the cloud this saves you money.

## Server Form Factors

Physical servers come in different sizes and shapes, as you'll see in Chapter 2. For instance, tower servers take up more space than their slimmer cousins, blade and rack-mount servers. Blade servers take up the least amount of space.

Arranging physical servers requires knowledge of server racks with sliding rails for equipment, including, but not limited to, servers. This is a much better use of space than tower servers arranged on the floor. Blade enclosures make even better use of space. A blade enclosure is mounted into a rack, and blade servers slide into the enclosure for the utmost in space savings (what techies like to call "increasing data center server density").

## Server CPUs

Although a single CPU chip may suffice for a desktop, servers often have multiple physical CPU chips plugged into CPU sockets on the server motherboard. Each CPU chip consists of one or more cores. A CPU core has the same functionality as a physical CPU, and multiple cores can be embedded on a single CPU chip.

CPUs use high-speed L1, L2, and L3 cache memory for data and instructions to speed up processing. A 4 gigahertz (GHz) CPU slows down when data needs to get on a bus to get to a different component in the server—for example, down to 600 megahertz (MHz).

Where desktops and servers commonly use the more powerful Complex Instruction Set Computing (CISC) processors, mobile and consumer electronic devices need less power, which means less heat and cooling are required, so they tend to use Reduced Instruction Set Computing (RISC) processors. RISC chips are designed to use simple instruction sets, where a single instruction can be executed within a single clock cycle. Bear in mind that CPUs can have speeds of millions (megahertz) or billions (gigahertz) of clock cycles per second. A clock cycle is an opportunity for CPU instructions to be executed.

## Server Memory

Scaling a database server to support larger datasets means adding memory (RAM). Virtual servers can be configured to use more RAM as needed, and in the cloud this is called *resizing*, which may require the virtual machine to be restarted. The extra RAM becomes



available from other virtual machines running on the same underlying physical host that have RAM to spare at the time.

Memory is either static (SRAM) or dynamic (DRAM). SRAM is used in smaller amounts and is faster than DRAM, which requires a constant refresh of electricity. SRAM is used for L1, L2, and L3 CPU cache memory and *not* the main system memory in your server—that would be DRAM.

You need to know about the various types of RAM chips and which ones will work in your specific hardware. Some servers will accept only double data rate 3 (DDR3) memory, others use DDR4, and so on. Memory chips may also have to be added in pairs for efficient use by server motherboards with multiple CPU sockets. Modern motherboards color-code memory chip sockets to facilitate pairing.

Many servers use memory chips that can detect and fix memory corruption errors; this is known as error correcting code (ECC) memory. That's why memory chips from a desktop, even though they may physically fit on a server motherboard, may not work correctly; they're usually non-ECC chips.

## Buses and Slots

Then there are expansion slots. Let's say you need to add a 10Gb Ethernet card in your physical server. All slots are not created equally; an older PCI Express (PCIe) ×2 card will fit and work in a PCIe ×16 slot, for example, but it won't work the other way around!

Cards have their own form factor, and you may even have run into this at home; some expansion cards seem to fit into a slot on the motherboard, but the card is too tall and the case won't fit back on. Expansion cards need a bus to move data into and out of the card; servers use a variety of buses to move data around the system.

Thinner servers such as rack-mounts and blades won't accommodate standard height expansion cards, but you can plug a daughter card into some rack-mount models that include slots that are oriented in such a way that standard-size cards can be used.

Blades get their additional capabilities such as storage and networking from other components plugged into the blade enclosure; you won't find PCI network cards plugged into a blade server.

## Virtualization

Virtual servers run as guests on a host (physical) computer, but they are still configured with virtual hardware. Some hypervisors (virtualization software) enable virtual guests to access physical hardware components directly, whereas others emulate hardware for the guest. Consider, for instance, how many vCPUs a virtual machine is configured with. Each vCPU maps to a CPU core and not a physical CPU chip.

A physical multiprocessor system enables each virtual machine to be configured with multiple vCPUs, but more vCPUs doesn't necessarily mean better performance. This is because the hypervisor can get bogged down finding physical CPU time slots for all of the vCPUs; it really depends on what is running within the virtual machine—sometimes less is more.

Bumping up the number of vCPUs can also cause another problem: licensing. Many software vendors license their server-based products based on the number of physical or logical processors; talk to your software vendor to ensure your compliance with their license rules. Ideally, you'll have an automated way to periodically inventory both hardware and software in use on servers; this is especially useful in a data center.

## Environmental Factors

Adding virtual servers to a hypervisor host doesn't make as much difference to power consumption and heat as does adding physical server hardware to a server room or data center. Careful arrangement of equipment can allow for optimal airflow, and ensuring adequate temperature and humidity control goes a long way toward ensuring that hardware runs at peak efficiency and has the longest life possible.

### Heat and Air Flow

Imagine cramming dozens of physical servers into a tiny closet with no airflow—excessive heat and servers do not get along well. In Chapter 2 we will explore how to determine whether you have sufficient heating/ventilation/air conditioning (HVAC) for your hardware.

The more equipment you have, the greater the power draw, which means more heat is generated, which means more cooling is required. Using blade enclosures and server virtualization are two ways to reduce power consumption and heat.

When you're working with a lot of equipment, it's important to ensure that cool air is fed into devices and that the warmer resulting air is exhausted elsewhere and properly taken out of the room or cooled down again. Data centers normally send cool air from the floor up to rack components. Arranging racks of equipment in rows facilitates the creation of hot and cold aisles.

CPUs generate a lot of heat the faster they run. Heat sinks are placed over the CPU chip to dissipate heat away from the chip, and this can be accelerated with a fan to suck the warm air away from the chip. Liquid cooling is another server cooling option that uses cool water brought into the server. The heated water then leaves the server and is cooled externally.

### Static Charges

An often overlooked aspect of environmental control includes electrostatic discharge (ESD) and fire suppression. Of course, large server rooms and data centers are designed with these things in mind, but smaller environments might simply convert the janitor's closet into a server/wiring closet.

ESD is the frenzied rush of electrons between differently charged objects; this is bad news because it can damage sensitive electronic components. There are ways to reduce the possibility of ESD—basically, take steps to ensure that objects coming into contact with each other are equally charged.

## Fire Suppression

In the event of an electrical fire, we do *not* want to use water as a fire suppressant. Server room and data center designs call for special construction and fire-suppression mechanisms not only to extinguish fires but also to minimize damage to electrical equipment. Chapter 2 will dive into proper ESD and fire-suppression practices.

## Storage

Most individuals and businesses like to keep data around for a while. Servers can play the role of a centralized file repository. You need to know the variations on how disk subsystems can be configured and provisioned. This topic in all its glory was formerly covered in a separate CompTIA certification (Storage+) that no longer exists; now, all of the storage details are a part of the CompTIA Server+ certification.

Traditionally, servers housed local disks within their physical enclosures, otherwise known as direct-attached storage (DAS). This has evolved to the point where server storage needs are addressed by accessing storage over a storage area network (SAN) and in some cases in the cloud. A server might boot the operating system from DAS and store data on the SAN. Various SAN standards will be covered in Chapter 4.

There are various types of storage media, including traditional magnetic hard disks and the much sought-after solid-state drives (SSDs). SSDs have no moving parts, so they are quieter, consume less power, and are pretty quick; the downside is their price. But as with many other things, all storage media are not alike. Disk interfaces such as universal serial bus (USB), serial-attached SCSI (SAS), and Fibre Channel are important considerations when working with servers, as are disk characteristics such as input/output operations per second (IOPS) and storage capacity.

Redundant array of independent disks, or, depending on who you ask, redundant array of *inexpensive* disks (RAID) groups multiple physical disks together for two potential purposes: better performance and fault tolerance. If your hardware supports RAID, you can configure it at the firmware level; alternatively, software RAID organizes disks to work as a logical unit within the supported operating system. You need to understand RAID levels so that you can configure disks to best suit a particular server need. For instance, RAID level 1 (disk mirroring) increases the resilience to failure, but RAID level 0 (disk striping) improves performance. We'll talk about how this works in Chapter 4.



**EXAM TIP** Rest assured that successful CompTIA Server+ Certification Exam candidates are expected to know when to apply specific RAID configurations given a particular scenario. It's not enough to know just the theory; make sure you have experience configuring and using various RAID configurations, even at the software level. A hands-on exercise in Chapter 4 can help guide you.

Once a server connects to storage of any kind, it's business as usual—partitioning, formatting, and setting user permissions. In some cases, disk or file encryption may be needed to secure data further.

# Network Concepts

Networking makes computing devices especially useful—the immediate accessing of services and the sharing of data. You may remember a time (the 1980s) when sharing data involved copying files to tape cassettes or floppy disks. In the early 1990s, we had modems with dial-up bulletin boards at what we now consider to be painfully slow transmission rates. How things have changed over the decades!

## Cables and Connectors

A solid understanding of computer networks includes network hardware as well as network software. Ask yourself how many miles of cables or wireless transmissions are involved when you check the local news web site each morning—it really is amazing how it works, since there are so many different technologies actually talking to each other.

Cables and connectors matter as much to the average home Internet user as they do to IT pros installing cabling in office buildings, data centers, and provider networks. Chapter 5 will provide the nitty-gritty for fiber-optic and copper-based transmission media.

## IP

On top of correctly configured network hardware, we have what most of us will configure and troubleshoot during our IT careers—Transmission Control Protocol/Internet Protocol (TCP/IP), usually referred to simply as IP. IPv4 and the newer IPv6 are collections of protocols that make things work on a network, but only if we've configured things correctly, such as the IP address, subnet mask, default gateway, and DNS server, among other settings.



**NOTE** There is no IPv5; IPv6 is newer than the 1970s-era IPv4. IPv6 adoption on the Internet and in the enterprise continues to grow.

IPv6 uses a 128-bit address space compared to IPv4's 32-bit address space. This means IPv6 has many more unique IP addresses compared to IPv4. IPv6 also includes built-in mechanisms that improve network quality of service (QoS) as well as security. IPv4 is used everywhere on the Internet, but IPv6 is not ubiquitous yet, so currently we can use tunneling technologies to get IPv6 traffic routed through the Internet.

It is only with a clear understanding of IP addressing and protocols in a multisubnet environment that you can properly support TCP/IP. For example, for a client workstation unable to connect to a server on a remote subnet, you might verify that name resolution and the default gateway (router) are configured correctly and reachable by the client.

## Ports and Protocols

The application layer of the TCP/IP model contains protocols that enable direct interaction with users over the network, and they depend on lower layer transport protocols such

as User Datagram Protocol (UDP) and Transmission Control Protocol (TCP). Think of a web server listening for requests from client web browsers. Hypertext Transfer Protocol (HTTP) is an application protocol that listens for client requests on TCP port 80. You can think of a port as a listening channel that uniquely identifies a network service.

Port numbers are 16-bit values tied to TCP and/or UDP. Mathematically, there are 65,536 possible port numbers ( $2^{16}$ ); however, port 0 is reserved for internal use when assigning dynamic port numbers, and since we are starting at zero, the last valid port number is 65,535. The first 1024 (ports 0–1023) are called *well-known ports* and are reserved for common network services such as HTTP.

Other application protocols such as Simple Mail Transfer Protocol (SMTP) and Domain Name Service (DNS) use different port numbers—TCP 25 and UDP 53, respectively. Some network software will allow you to specify the port number that it listens on, while others will not.



**EXAM TIP** Know the application layer protocols and their port numbers. Scenario-based questions may indirectly test your knowledge in this area. A much more exhaustive list will be presented in Chapter 5.

## Server Operating Systems and Server Roles

When it comes to servers, a myriad of questions need answers. What is the server's purpose? How many users will be connecting at once? How critical is it that the server stays up and running 24/7? Planning servers before jumping into the fun stuff is paramount.

Chapter 3 will address these questions and more:

- Will the server operating system be installed onto physical hardware or in a virtual machine?
- Does the server support basic input-output system (BIOS) or unified extensible firmware interface (UEFI)?
- Are we installing a Windows, UNIX, or Linux operating system variant?
- What type of disk partitioning and file systems should we use?
- How will we remotely connect to the server?
- How will patches be applied?
- How will we secure the server?

Servers provide services and content to users. Companies can have their own servers running on their own equipment (on premises), but installing, configuring, and running various server roles in the cloud is becoming more and more common. The great thing about provisioning virtual server roles in the public cloud (on provider equipment) is how quickly you can get servers up and running without an investment in hardware. And, generally, you pay only for the time the virtual server is running—no wonder its popularity has soared!



**NOTE** Throughout this book you will learn about not only Windows Server operating systems but also Ubuntu Linux, in both physical and virtual environments. This includes hands-on exercises and troubleshooting tips as well.

## Server Roles

Infrastructure services such as Network Time Protocol (NTP), DNS, and Dynamic Host Configuration Protocol (DHCP) must be in place before anything else will work on the network. NTP assures that network devices agree on the time. DHCP delivers a valid IP configuration to hosts, while DNS is used to resolve friendly names (such as `www.mheducation.com`) to an IP address.

Directory servers provide a central network database of configuration objects, such as user accounts used for authentication. In a Microsoft Active Directory Domain Services (AD DS) environment, DNS must be functional for AD to be functional. In Chapter 3, we will install an AD domain and then explore the various ways that computers can be joined to the domain for centralized administration.

Application servers and web servers—how are they different? Traditionally, application servers run a server-side application that clients can access over the network in various ways. These days, the connection often uses HTTP or its secure counterpart, HTTPS. This means a web server needs to be on the server to handle these requests. The application server comes into play when developers build a business-specific solution on the web server—in other words, business logic, such as a payment processing service for e-commerce transactions. Planning for and configuring a web or application server have different requirements than an internal file server would.

File servers are file repositories in the enterprise. As IT administrators, we grant user permissions and set quotas on disk space usage. Having enough fast disks is a primary concern here, as is the network speed for file access. These days, cloud storage is all the rage; individuals and enterprises can provision storage as needed on provider equipment and pay only for the amount of disk space used. There are also hybrid solutions in which enterprise file servers can synchronize and even do backups to the cloud.

## Maintenance

Once servers are humming along, we want to keep them going. Maintenance tasks, such as applying patches and firmware updates and ensuring that the server complies with organizational configuration settings and security policy, require continual attention, so just because the server is installed and configured doesn't mean our work is complete. Your approach to these tasks will differ in a small company with a handful of servers compared to a large data center with thousands of servers. You'll also need a way to connect remotely to servers through either a software or a hardware solution.

One way to ensure server uptime is to configure failover clustering. This involves at least two server nodes working together to provide the same service. If one cluster node fails, users are automatically switched over to the service running on a remaining node. Network load balancing (NLB) can be used to achieve optimal network performance,

such as for a busy web site. It does this by distributing incoming traffic to the least busy back-end server hosting the web site. In the case of public cloud computing, failover and load balancing are often enabled with minimal effort, and in some cases they are automatically enabled!

## Server Management

Long gone are the days of logging in directly at the physical server console. Instead, administrators connect remotely over the network to manage servers. This can be done using software or hardware solutions.

Software remote management solutions such as Secure Shell (SSH) depend on the server operating system functioning properly. But what if the server hangs and doesn't respond? Then what?

Hardware remote management solutions such as Intelligent Platform Management Interface (IPMI) and Integrated Dell Remote Access (iDRAC) run independently of the operating system—so a server that hangs isn't a problem. Hardware management solutions require a valid IP configuration so that the server can be accessed remotely.

## Monitoring

Monitoring the performance of servers is an ongoing task. You'll normally monitor some aspect of CPU, memory, disk, and network use. You might even enable alerts based on configured thresholds ("We're almost out of RAM—do something!"). In a virtualization (and certainly a data center) environment, monitoring can be complex and must be configured carefully. We should monitor not only the aspects of physical virtualization hosts (hypervisors) but also the virtual machine guests that run on them.

## Service Level Agreements

A service level agreement (SLA) is a contract between a provider and consumer stating expected levels of service, including details such as guaranteed uptime and response time. Data centers hosting IT services for numerous customers must diligently monitor server performance and network availability to ensure adherence to their SLA guarantees to their customers. A great example of this is an SLA for a public cloud service, such as Amazon Web Services (AWS) Simple Storage Service (S3) bucket cloud storage. In this case, the SLA is a contract between Amazon and its customer that guarantees a monthly uptime percentage such as 99.9 percent; if this uptime is not maintained, the customer receives a service credit, which is applied to their next bill.

## Security Considerations

Servers can run on networks that are completely isolated from the Internet (air-gapped), or they can in some way be connected. Either way, servers need to be secured and the effectiveness of security controls continuously evaluated. In fact, a large percentage of security breaches actually occur from *within* the network.



## Hardening

*Hardening* is an all-encompassing term used to describe how we lock down or tighten security. For servers this includes, but is not limited to, the following:

- Placing servers in locked rooms/racks
- Disabling unnecessary services
- Applying patches
- Running antimalware software
- Adhering to the principle of least privilege
- Enabling multifactor authentication
- Auditing the use of sensitive data
- Encrypting data in motion (network)
- Encrypting data at rest (storage)

In a larger network, instead of hardening each server manually, we can apply security settings from a central configuration. Many tools can be used to accomplish this for Windows, UNIX, and Linux servers, such as Microsoft PowerShell Desired State Configuration (DSC) and Microsoft System Center Configuration Manager (SCCM).

Larger networks need an easy way to determine which servers are not compliant with organizational security policies. They also need an easy way to bring these devices into compliance. Enterprise tools such as SCCM have these problems covered. There are also cloud-specific tools that do the same thing, such as Microsoft Azure Policy compliance reports.



**EXAM TIP** Knowing how to harden a handful of servers is very important. The CompTIA Server+ certification also applies to server management in large environments (such as data centers), so always keep in mind how you can apply configurations on a large scale.

## Network Security

Everything would be so much more secure if computers didn't connect to a network, but this is not reality. Fortunately, there are plenty of hardware and software security solutions to address network security.

The first step in protecting a network is to be very selective about which users and devices connect to the network. The Institute of Electrical and Electronics Engineers (IEEE) 802.1X security standard is focused on controlling access to a network. The idea is that network edge devices (virtual private network [VPN] concentrators, wireless access points, and Ethernet switches) do not perform authentication of connecting devices, but instead forward those requests to a central authentication server on a protected network.

In the case of a LAN device not supporting IEEE 802.1X, it can be disabled at the switch port level.



## Authentication

Authentication is the proving of one's identity. This applies to users as well as computing devices. For instance, a smartphone may be required to authenticate itself to a VPN before allowing user credentials.

Multifactor authentication combines at least two different categories of authentication, such as *something you know* (such as a PIN) and *something you have* (such as a smartcard). Combined with firewalls that control the flow of traffic into and out of a network, this is a good first step to securing your network.

## PKI

Public Key Infrastructure (PKI) is a hierarchy of trusted digital certificates used for security. Certificates are issued from a trusted authority to users, devices, or services, and they can exist as files on a disk or settings located in a secured storage location (such as the Windows certificate store), or they may be written to a smartcard.

Certificates can be used in many ways:

- Authenticate a smartphone to a VPN appliance
- Digitally sign an app before it is published to an app store
- Encrypt sensitive e-mail messages before transmission

You can get certificates from a trusted third party on the Internet, or you can create self-signed certificates for internal use. Chapter 6 will demonstrate how to install and configure a PKI using Windows Certificate Services.

## Access Control

After successful authentication, authorization to use a network resource is granted based on access control lists (ACLs). There are many types of ACLs, such as those allowing access to a network and those that allow or deny access to a specific resource such as a file or a web site.

In the real world and on the CompTIA Server+ Certification Exam, you should be able to determine which permissions must be set to accomplish a specific goal, and you must also understand permission *inheritance* and *precedence*. The principle of least privilege states that only the permissions required to complete necessary tasks should be granted, and this must always be followed. We'll go through this in detail, including a hands-on exercise, in Chapter 6.

## Data at Rest

Securing transmissions to and from servers is always a good idea, but what about data once it reaches a destination and is stored on media? We keep hearing media reports about how millions of customers' personal data has been compromised, and in some cases this involves a malicious user gaining physical access to a storage device.

Encrypting entire disk volumes or individual files and folders adds another layer of security. In the case of a stolen physical storage device, if encryption has been implemented properly, you needn't worry—the device could serve as a lovely paperweight instead of as a source of valuable data. Encryption of data at rest is very useful with public

cloud computing, where data center administrators are prevented from accessing cloud tenant data.

You can use tools to encrypt specific files and folders, or you may choose to encrypt entire disk volumes. Data backups should also be encrypted for additional protection.

When data reaches the end of its useful life, it needs to be disposed of in a secure manner, which may be required for regulatory compliance. This includes the remote wiping of mobile devices, soft and hard wiping of data, and the physical destruction of storage media.

## Troubleshooting and Optimizing Performance

Solid troubleshooting stems from truly understanding the underlying technologies and the proper application of a troubleshooting methodology. Your ability to reproduce a problem and determine its scope is important to resolving issues quickly.

### Optimizing Performance

Optimizing the performance of servers and their surrounding ecosystem can prevent negative incidents from occurring in the first place. In the enterprise, troubleshooting can involve poor performance. A server running out of disk space will slow the system to a crawl and could cause services to freeze, yet this could be caused by a lack of memory (RAM). Chapter 7 deals with both performance and troubleshooting, since often they are related.

### Troubleshooting

We can apply the CompTIA troubleshooting methodology to real-world situations as well as to scenario-based exam questions. For instance, change only one thing at a time and observe the results of that change. Organizations use a variety of tools to document problems and their eventual resolutions, from recording this information in a spreadsheet to a full-fledged enterprise help desk ticketing system. Over time, this type of knowledge base can prove very valuable.

### Hardware and Software Troubleshooting

One area you must be able to troubleshoot is hardware. All hardware eventually fails, in some cases prematurely due to excessive heat or perhaps ESD. Think of the dreaded “Operating system not found” message on a server. Is that a hardware problem, or is it a software problem such as a corrupt file? You have to ask yourself, “Self, what has changed since this last worked correctly?”

You’re more than likely going to encounter software rather than hardware issues in your IT career. Consider software misconfigurations. A specific server configuration may work well in one environment for a particular use, but it may fail spectacularly under different circumstances—it’s all about meeting business needs.

Other software problems can be much more difficult to troubleshoot—take random freezes, for example. You can still successfully troubleshoot unresponsive systems by viewing log files over time to determine whether or not a pattern exists, or to see what else on the system might have caused the problem.

## Network Troubleshooting

Problems with server network communications can be a real pain in the neck; in some cases, such as with physical servers, we may need to be physically present to solve a problem. Hardware and software remote control solutions do nothing if we can't remotely communicate with the server in the first place.

An enterprise IPv4 or IPv6 environment can be tricky to troubleshoot because the problem is specific to that implementation. Technical knowledge and network documentation (or details about network configurations including infrastructure) are central to snuffing out issues quickly when they arise. As always, we need to determine who (or what) is experiencing network communication issues. If it's a single user or station, we know our infrastructure is good, so that enables us to focus our attention on the most likely problem sources.

Just because we can't connect a network resource by name doesn't mean the Internet is broken. Chapter 7 is chock-full of tips and tools at our disposal.



**EXAM TIP** One way to ensure your success in passing the CompTIA SK0-005 exam is to know *when* to use a particular network troubleshooting tool or command. Of course, you are also expected to know *how* to use a tool to solve a specific problem.

## Storage Troubleshooting

Because the former CompTIA Storage+ certification body of knowledge is now included in the CompTIA Server+ Certification Exam, not only do you have to understand the subject of storage, but you also have an obligation to know how to fix it when it's broken or performing badly.

Equaling the other troubleshooting topics in Chapter 7, here you are required to know how to use the correct tools to solve storage problems that are interwoven with hardware (inappropriate RAID configurations), software and network issues (inability to connect to network storage), as well as matters of security (most often permissions).

## Security Troubleshooting

This topic is interweaved with some form of hardware, software, or network troubleshooting. Think of an expired PKI certificate on a laptop that prevents that laptop from connecting to a VPN—sounds like a network issue, right? Well, it is, sort of, but it's not the *source* of the problem.

The exam could include questions related to users being unable to access some type of network resource because of misconfigured permissions. Again, it *sounds* like it could be a network issue, but if other users can access that same resource, the problem may lie elsewhere.

Security-related hardware and software have log files that we can peruse to determine problem causes. But really, we can't solve problems with silo thinking—in other words, we need to know about surrounding configurations (hardware, software, and network) and business processes to solve security problems effectively.

Chapter 7 will dive into malware and firewall issues, group policy configurations, and many other problems (and resolutions) associated with security.

## Preparing for the Worst

Bad things happen. We can't always predict how or when, but we *can* plan in advance for potential future negative incidents. This is the theme in Chapter 8.

Disaster recovery (DR) and business continuity relate to matters such as

- Determining the impact of interruptions
- Duplication (removing single points of failure)
- Data replication
- Data backup
- Alternate locations to resume business operations

We can't prevent all adverse situations, but we can minimize the unfavorable impacts against servers and business operations. This can succeed only with up-front planning and ensuring that technicians know their roles when it comes to incident response handling.

## Cloud Computing

With cloud computing, everything old is once again new, at least when it comes to IT. All of the items discussed in this chapter (servers, storage, virtualization, networking, security, troubleshooting and performance optimization, and preparing for the worst) are IT topics that you may have years of experience with in an on-premises environment. If you use cloud computing services, some or all of this runs elsewhere on somebody else's equipment and is accessible over a network. The same care in planning, managing, and securing IT services that applies to an on-premises environment applies in the cloud.

### Cloud Computing Characteristics

Running a few virtual machines on a server on your network does not, in and of itself, constitute a cloud network. According to the National Institute of Standards and Technology (NIST) Special Publication (SP) 800-145, five characteristics define cloud computing:

- **Measured service** Also called *metered usage*; cloud service usage is tracked for billing purposes much like electricity is.
- **Rapid elasticity** Cloud services can automatically expand to accommodate spikes and contract when things quiet down. This is referred to as *autoscaling*.
- **Resource pooling** Cloud provider network, compute, and storage infrastructure is grouped together and made available to cloud customers.
- **Broad network access** Cloud services are accessible over a network using any type of device. Most cloud services do not require special software to be installed on client devices.
- **On-demand self-service** Cloud users can provision or deprovision cloud resources without involving the cloud service provider.

## Cloud Computing Deployment Models

There are a few different cloud deployment models, although they all exhibit the same five cloud computing characteristics just mentioned. Your agency or organization may choose to use one or more of these cloud types.

### Public Clouds

A public cloud service provider (CSP) owns and manages the underlying equipment that cloud services such as storage, virtual machines, and databases run on. Examples of public CSPs include Microsoft Azure, Amazon Web Services (AWS), Google Cloud Platform (GCP), and Alibaba Cloud. This underlying equipment is housed in data centers around the globe. For example, AWS consists of data centers located in more than 190 countries.

Public clouds are accessible over the Internet to anybody who wants to sign up with a new cloud provider account, some of which are free and some of which are not. At the enterprise level, public cloud providers offer different types of cloud subscriptions with different service use and support options, and of course customers pay for the cloud resources they use, such as amount of cloud storage or how long cloud-based virtual machines are left running. At the individual level, some cloud services such as Dropbox, Gmail, and Microsoft OneDrive have free variations.

There are benefits for organizations that choose to use public cloud services:

- Outsourcing IT infrastructure acquisition and management to the cloud provider enables the organization to focus more on using technology to solve business problems.
- CSPs offer less up-front capital expenditures (CAPEX) in IT equipment acquisition, configuration, and management.
- Consumers pay ongoing monthly operating expenditures (OPEX) for cloud service use as opposed to CAPEX when hosting all IT services on-premises.
- Public cloud providers are constantly undergoing third-party audits to attain security accreditations.

### Private Clouds

In a private cloud, the underlying equipment that provides cloud services is owned and used by a single organization, which constitutes CAPEX initially. Management of the private cloud infrastructure can be done by the owning organization, or it can be outsourced to IT consulting firms. The primary benefit of a private cloud is the owning organization having full configuration control of the entire infrastructure.

Remember that all cloud deployment models must adhere to the NIST's five cloud computing characteristics. The *measured service* characteristic tracks cloud service usage. In a private cloud, this can be done simply to track usage by various users, departments, or projects within the organization, or it can be used for departmental chargeback, where the department in the organization pays the IT department for the use of private cloud services. Unlike a public cloud, private clouds are not available to anybody over the Internet willing to sign up with a new account.

## Community Clouds

A community cloud is designed for limited access by groups with similar IT service and security needs, such as government agencies and specific industries such as finance or medicine that must comply with very specific regulations. Examples of community clouds include Microsoft Cloud for Healthcare and AWS GovCloud for American government agencies.

## Hybrid Clouds

Hybrid clouds include any combination of public, private, and community clouds. An example is an organization hosting a private cloud on premises that uses cloud bursting to consume public cloud resources when private cloud resources are depleted.

## Cloud Computing Service Models

Cloud services are grouped into categories, or models, based on the type of cloud service and where responsibility of managing that service lies. Anything as a Service (XaaS) is a general category of cloud-computing services, where the *X* represents a variable that could be a number of things. Table 1-1 describes the three most common cloud service models.

Cloud Service Model	Description	Responsibility
Infrastructure as a Service (IaaS)	Storage, network components, firewalls, and virtual machines deployed in the cloud	CSP is responsible for the underlying infrastructure; cloud customer is responsible for the configuration, and in the case of virtual machines, patching the OS and any installed apps
Platform as a Service (PaaS)	Software developer tools, network directory services, and databases where the underlying virtual machines and software are automatically provisioned	CSP is responsible for the underlying infrastructure and deployment and management of virtual machines and software to support the service; cloud customer is responsible for configuring developer tools, directory service objects, and databases
Software as a Service (SaaS)	End-user productivity apps such as word processors and spreadsheets, where the underlying virtual machines and software are automatically provisioned	CSP is responsible for the underlying infrastructure and deployment and management of virtual machines and software to support the service; cloud customer has minimal configuration options and is responsible for the management of data resulting from the use of the cloud service

**Table 1-1** Common Cloud Service Models

## Chapter Review

This chapter provided insight into what the rest of this book offers. The body of knowledge covered for the CompTIA Server+ Certification Exam (SK0-005) is valuable for anyone working in an IT-related job.

### Questions

1. Which server form factor occupies the most space?
  - A. Blade
  - B. Virtual
  - C. Rack-mount
  - D. Tower
2. What term is used to describe a virtualized operating system?
  - A. Host
  - B. Guest
  - C. Hypervisor
  - D. Load balancer
3. Which two of the following items are related to HVAC?
  - A. Perimeter fencing
  - B. Ventilation
  - C. Dual power supplies
  - D. Temperature control
4. How can ESD be reduced?
  - A. Storing hardware components in a freezer
  - B. Ensuring that objects coming into contact with one another have different charges
  - C. Ensuring that objects coming into contact with one another have equal charges
  - D. Storing hardware components in plastic bags
5. What is the difference between DAS and a SAN?
  - A. DAS enables the server to use network storage as if it were local. SAN storage is local to the server.
  - B. DAS is local storage, and SAN is network storage.
  - C. DAS is local storage specifically for virtual servers. SAN refers to cloud storage.
  - D. DAS and SAN are separate terms that refer to exactly the same thing.

6. What benefits do SSDs provide over magnetic hard disks? Choose two.
  - A. Stronger encryption
  - B. Increased file integrity
  - C. Less power consumption
  - D. Quicker file access
7. You would like to improve disk I/O performance for data residing on a physical server. In accordance with your organizational policy, a server must remain running even if an operating system disk fails. The server is currently configured with RAID 1 for the operating system and data. What should you configure?
  - A. Nothing; RAID 1 already offers the best performance.
  - B. Remove the existing RAID configuration and enable RAID 0.
  - C. Add disks to the server. Configure RAID 0 and create a file system. Move the existing data to the newly created file system.
  - D. Add disks to the server. Configure RAID 0 and create a file system. Move the existing operating system file to the newly created file system.
8. What role does a default gateway perform in an IPv4/IPv6 environment?
  - A. It prevents local area network broadcasts from reaching other subnets.
  - B. It is a router that allows traffic into and out of a network.
  - C. It prevents multicasts from reaching other subnets.
  - D. It is a router that performs domain name-to-IP address resolution.
9. Roman is the server administrator for an international insurance company. The London office currently has Active Directory domain controller servers configured with IPv6. The Toronto office domain controllers are also configured with IPv6, yet replication over the Internet is failing between Toronto and London. What should Roman do to enable domain controller replication between the two sites?
  - A. Enable IPv6 replication through group policy.
  - B. Configure an IPv6 tunneling solution.
  - C. Make sure the server PKI certificates allow domain controller replication.
  - D. Ensure that servers at both locations are configured with the same default gateway.
10. While capturing network traffic, you notice an excessive amount of traffic destined to a particular host on TCP port 25. What is one possible explanation for this?
  - A. A mail server is being flooded with spam.
  - B. A network broadcast attack is taking place.



- C. Users' devices on the network happen to be issuing DNS queries at the same time.
  - D. A web server is experiencing a large volume of concurrent visitors.
11. Which of the following demonstrates an advantage of cloud computing? Choose two.
- A. Web servers can be scaled to handle heavy volumes of traffic.
  - B. SLAs provide fault tolerance.
  - C. Virtual machines and storage can be rapidly provisioned.
  - D. SLAs provide guaranteed uptime.
12. Your company's e-commerce site is experiencing excessive traffic during the holiday shopping season. Being the server specialist, you have been directed to configure the site to improve user response time during peak loads. What should you configure?
- A. Network load balancing
  - B. Failover clustering
  - C. VPN
  - D. Operating system virtualization
13. Industry regulations require multifactor authentication to be used for sensitive server systems. Currently, users authenticate to these systems with a username and password. What should be done to ensure regulatory compliance?
- A. Nothing; multifactor authentication is already in use.
  - B. Configure a PIN requirement in addition to current authentication settings.
  - C. Use PKI certificates to secure authentication further.
  - D. Enforce periodic password changes.
14. What category of processor is used by mobile devices?
- A. CISC
  - B. BISC
  - C. RISC
  - D. DISC
15. You plan to use extra DDR3 memory chips for your server, which currently supports DDR4. What should you consider?
- A. DDR3 chips can be plugged into DDR4 sockets.
  - B. DDR3 chips cannot be plugged into DDR4 sockets.
  - C. The server BIOS will have to be updated.
  - D. The server must use UEFI.

16. You are manually deploying virtual machines into a private cloud. Which type of cloud service model does this apply to? Choose the best answer.
- A. SaaS
  - B. XaaS
  - C. PaaS
  - D. IaaS

## Questions and Answers

1. Which server form factor occupies the most space?
- A. Blade
  - B. Virtual
  - C. Rack-mount
  - D. Tower

**D.** Tower servers take up the most space. A, B, and C are incorrect. Blade servers are essentially circuit boards that plug into a single server chassis. Virtual is not a server form factor. Rack-mount servers are installed in a tall metal rack along with other rack-mount servers.

2. What term is used to describe a virtualized operating system?
- A. Host
  - B. Guest
  - C. Hypervisor
  - D. Load balancer

**B.** Virtualized operating systems are known as guests. A, C, and D are incorrect. Hosts are the physical machines on which guests run, and this is the same with hypervisor. Load balancers distribute incoming network traffic destined for a network service (such as a web site) to improve network performance.

3. Which two of the following items are related to HVAC?
- A. Perimeter fencing
  - B. Ventilation
  - C. Dual power supplies
  - D. Temperature control

**B, D.** HVAC stands for heating, ventilation, and air conditioning, which encompasses ventilation and temperature control. A and C are incorrect. Perimeter fencing is a physical security measure, and dual power supplies provide hardware redundancy.

4. How can ESD be reduced?

- A. Storing hardware components in a freezer
- B. Ensuring that objects coming into contact with one another have different charges
- C. Ensuring that objects coming into contact with one another have equal charges
- D. Storing hardware components in plastic bags

**C.** Equalizing the charge between objects that will come into contact with each other reduces the flow of electrons, which can damage sensitive electronic components. A, B, and D are incorrect. ESD is not reduced by storing components in plastic bags or freezers.

5. What is the difference between DAS and a SAN?

- A. DAS enables the server to use network storage as if it were local. SAN storage is local to the server.
- B. DAS is local storage, and SAN is network storage.
- C. DAS is local storage specifically for virtual servers. SAN refers to cloud storage.
- D. DAS and SAN are separate terms that refer to exactly the same thing.

**B.** Direct-attached storage (DAS) means storage local to the server. A storage area network (SAN) is network storage, not the other way around; they are not the same thing. A, C, and D are incorrect. DAS can be used by physical or virtual servers. SANs are not specifically related to cloud computing. DAS and SAN are not one and the same.

6. What benefits do SSDs provide over magnetic hard disks? Choose two.

- A. Stronger encryption
- B. Increased file integrity
- C. Less power consumption
- D. Quicker file access

**C, D.** Solid-state drives (SSDs) have no moving parts and thus consume less power and generally provide quicker file access than magnetic hard disks. A and B are incorrect. Encryption and file integrity are normally configured at the software level and are not specific to SSDs.

7. You would like to improve disk I/O performance for data residing on a physical server. In accordance with your organizational policy, a server must remain running even if an operating system disk fails. The server is currently configured with RAID 1 for the operating system and data. What should you configure?
- A. Nothing; RAID 1 already offers the best performance.
  - B. Remove the existing RAID configuration and enable RAID 0.
  - C. Add disks to the server. Configure RAID 0 and create a file system. Move the existing data to the newly created file system.
  - D. Add disks to the server. Configure RAID 0 and create a file system. Move the existing operating system file to the newly created file system.

**C.** RAID 0 is disk striping. The best course of action is to add disks to the server, configure RAID 0 and create a file system, and move the existing data to the newly created file system. A, B, and D are incorrect. The existing RAID 1 mirror does not provide a performance benefit, and it should be left alone for operating system file fault tolerance, but data should be moved to the new file system on the RAID 0 array.

8. What role does a default gateway perform in an IPv4/IPv6 environment?
- A. It prevents local area network broadcasts from reaching other subnets.
  - B. It is a router that allows traffic into and out of a network.
  - C. It prevents multicasts from reaching other subnets.
  - D. It is a router that performs domain name-to-IP address resolution.

**B.** The default gateway is a router that allows traffic into and out of a network, in some cases based on conditions. A, C, and D are incorrect. Although routers do not forward broadcasts, that is not the role of the IP default gateway from an IP perspective. Multicast (group) traffic can normally traverse routers. Domain name-to-IP address resolution is a function of a DNS server.

9. Roman is the server administrator for an international insurance company. The London office currently has Active Directory domain controller servers configured with IPv6. The Toronto office domain controllers are also configured

with IPv6, yet replication over the Internet is failing between Toronto and London. What should Roman do to enable domain controller replication between the two sites?

- A. Enable IPv6 replication through group policy.
- B. Configure an IPv6 tunneling solution.
- C. Make sure the server PKI certificates allow domain controller replication.
- D. Ensure that servers at both locations are configured with the same default gateway.

**B.** Roman should configure an IPv6 tunneling solution that will enable the IPv6 packets to traverse the Internet (an IPv4 network). A, C, and D are incorrect. There is no such thing as IPv6 replication. PKI will not enable replication; it is used for security purposes. Servers at each site should point to their local default gateway, not the same default gateway overall.

10. While capturing network traffic, you notice an excessive amount of traffic destined to a particular host on TCP port 25. What is one possible explanation for this?
- A. A mail server is being flooded with spam.
  - B. A network broadcast attack is taking place.
  - C. Users' devices on the network happen to be issuing DNS queries at the same time.
  - D. A web server is experiencing a large volume of concurrent visitors.

**A.** TCP port 25 is normally reserved for SMTP mail servers. Excessive traffic destined for the mail server could indicate spamming is taking place. B, C, and D are incorrect. Packets destined for port 25 are not broadcast packets; neither are they DNS queries, which use UDP port 53, or web server requests, which would use TCP port 80 or 443.

11. Which of the following demonstrates an advantage of cloud computing? Choose two.
- A. Web servers can be scaled to handle heavy volumes of traffic.
  - B. SLAs provide fault tolerance.
  - C. Virtual machines and storage can be rapidly provisioned.
  - D. SLAs provide guaranteed uptime.

**C, D.** Although virtual machines can be provisioned without a cloud environment, one characteristic of cloud computing is the rapid deployment of services by users. Service level agreements (SLAs) can guarantee uptime, among other details; SLAs are contracts between cloud providers and consumers. A and B are incorrect. Web servers can be configured to handle heavy traffic volumes without cloud computing. SLAs themselves do not provide fault tolerance.

- 12.** Your company's e-commerce site is experiencing excessive traffic during the holiday shopping season. Being the server specialist, you have been directed to configure the site to improve user response time during peak loads. What should you configure?
- A.** Network load balancing
  - B.** Failover clustering
  - C.** VPN
  - D.** Operating system virtualization

**A.** Network load balancing (NLB) distributes incoming traffic to the least busy of multiple back-end servers offering the same service. B, C, and D are incorrect. Failover clustering enables network services to continue even if servers fail, as long as at least one cluster node remains running. VPNs enable encrypted connectivity to a private network over an untrusted network. Operating system virtualization enables multiple OSs to run in their own virtual environments on the same physical host.

- 13.** Industry regulations require multifactor authentication to be used for sensitive server systems. Currently, users authenticate to these systems with a username and password. What should be done to ensure regulatory compliance?
- A.** Nothing; multifactor authentication is already in use.
  - B.** Configure a PIN requirement in addition to current authentication settings.
  - C.** Use PKI certificates to secure authentication further.
  - D.** Enforce periodic password changes.

**C.** Multifactor authentication requires multiple authentication "categories"; in this case, something you know (username and password) and something you have (PKI certificate). A, B, and D are incorrect. Multifactor authentication is not in place when only the *something you know* category is in use, which is also true even if you enable PINs or enable periodic password changes.

14. What category of processor is used by mobile devices?

- A. CISC
- B. BISC
- C. RISC
- D. DISC

C. Reduced Instruction Set Computing (RISC) processors are often used by mobile devices because of their low power requirements. A, B, and D are incorrect. Complex Instruction Set Computing (CISC) processors tend to be used in desktops and servers where more complex processing is required. BISC and DISC are fictitious terms.

15. You plan to use extra DDR3 memory chips for your server, which currently supports DDR4. What should you consider?

- A. DDR3 chips can be plugged into DDR4 sockets.
- B. DDR3 chips cannot be plugged into DDR4 sockets.
- C. The server BIOS will have to be updated.
- D. The server must use UEFI.

B. DDR memory chips are not interchangeable with other DDR version standards. A, C, and D are incorrect. DDR3 chips cannot be physically plugged into DDR4 sockets. The BIOS and UEFI are irrelevant to whether DDR3 chips can be used in DDR4 slots.

16. You are manually deploying virtual machines into a private cloud. Which type of cloud service model does this apply to? Choose the best answer.

- A. SaaS
- B. XaaS
- C. PaaS
- D. IaaS

D. Infrastructure as a Service (IaaS) includes cloud-based storage, network components, and virtual machines that are manually deployed and managed by the cloud customer. A, B, and C are incorrect. Software as a Service (SaaS) refers to end-user productivity software accessed over a network in a cloud. Anything as a Service (XaaS) is a catch-all phrase describing a cloud service accessed over a network. Platform as a Service (PaaS) refers to cloud services related to software development, network directory services, and databases where the underlying virtual machines and software are automatically provisioned and not directly under the control of the cloud customer.

# Server Hardware

In this chapter, you will

- Learn about server form factors: tower servers, rack-mounted servers, and blade servers
- Learn how BIOS works
- Learn how UEFI works
- Review the basics of CPUs and how they handle data
- Review the different types of memory
- Learn about bus types
- Learn about NICs
- Review storage types
- Learn about power and environmental controls

Server hardware isn't just about the components inside the server—it's also about the size of the server case and its components. This chapter will cover what you need to know when ordering and replacing server components.

## Server Form Factors

Physical servers and their parts come in a variety of shapes and sizes, or *form factors*. Not all components will fit into every server. For example, adapter cards that fit well within a server tower will not fit into a blade server, and a tower server can't be rack-mounted in a server room rack.

Form factors also apply to computer cases, power supplies, motherboards, expansion cards, and so on. The dimensions of these pieces determine which ones fit together properly. Data centers need to fit as many servers as possible within a finite amount of space, so increasing server density is possible using smaller server form factors.

Other components such as PCI Express (PCIe) expansion cards adhere to industry-standard form factors. Perhaps you want to add a 10 Gbps (gigabits per second) Ethernet network card to your rack-mounted server. This could be a problem, because a standard PCIe card won't fit within most rack-mounted servers, which are much thinner and smaller than tower servers.

Server technicians need to know all the details related to the types of servers they are responsible for. Let's say it's the first day on the job for a CompTIA Server+ certified tech.



How can this person know what servers are in place? Clearly, one way is to inventory servers physically in server rooms or data centers, but larger enterprises will have automated solutions that inventory physical and virtual servers and store the results in a database that can be queried and reported on. Virtual servers, of course, don't have form factors.

## Tower Servers

Tower servers have been around for a long time, and the tower is what most people think of when they hear the word “server.” Powerful desktop computers (such as those used by gamers) are often towers. This server form factor isn't screwed into any type of mounting device; it is a standalone computer that can be easily moved without removing screws or sliding it out of a rack.

In this server form factor, all server components are housed within a single case that can sit directly on the floor or on a desk. Server components are easy to find, because this is a tried-and-true hardware technology.

Adding dedicated components such as disk storage to the server is easy, because there is plenty of physical space inside a tower server. This isn't the case with other server form factors such as blade servers, however, although blades can use rack-mounted storage devices; the storage isn't physically installed inside the blade. Imagine trying to accommodate hundreds or thousands of tower servers in a server room or data center: the cost of real estate alone would be tremendous! Suffice it to say that tower servers don't scale well. They take up a lot of space and they can't be rack-mounted, as you can see in Figure 2-1.

**Figure 2-1**  
Dell PowerEdge  
T620 tower  
server, courtesy  
of Dell, Inc.



Tower servers are often used in smaller offices. If IT budgets are stretched, towers might be an attractive option, because even a standard desktop PC could be configured with a server operating system—but, of course, this isn't designed for large-scale use.

Another possibility is a central IT office that preconfigures servers to be used in branch offices. Perhaps only a single server is needed at a branch office to localize user access to server services, so it might make sense to ship a tower server to that location instead of purchasing expensive server racks and a single-rack mounted device.

When it comes to component redundancy, most tower servers fall short. They can accommodate standard hardware components, but they don't often come with redundant power supplies. There's also the issue of power and data cables. If your server room contains only towers, you may find it tricky to organize all of the cables. Server racks have conduits into which cables are easily and neatly arranged, which makes labeling and troubleshooting easier. Main distribution frames (MDFs) are cable racks that organize data cabling from intermediary distribution frames (IDFs), which organize cables plugged into equipment. (Just make sure you label both ends of each cable!)

## Rack-mounted Equipment

The 1990s produced the server rack form factor courtesy of Compaq, which was acquired by Hewlett-Packard in 2002. Computing and the Internet were taking off, and many companies realized they needed servers on premises. Nowadays that has shifted to running servers in the public cloud on somebody else's equipment.

Rack-mounting increases the potential server density in a server room or a data center, and using this equipment can increase security, because most racks have front and back doors that can be locked. In a large data center that accommodates multiple customers or tenants, controlling physical access to rack-mounted servers and equipment is important. The size of your enterprise and its data requirements, and the size of the data center, determine how many racks you can use. Server technicians will sometimes need to identify a physical server in a server rack; unit identification (UID) LED lights on the front and back of a rack-mounted server will blink when the server is being managed remotely or when the UID button is physically pressed on the server. This is useful when technicians must access the back of a rack-mounted server and want to ensure they are working on the correct server.

Special cases and rails are used so that rack-mounted equipment can be easily inserted and removed from racks. Rack-mounted servers ship with their cases, and sometimes they also include rails and the screws needed to secure the rails on the rack.

## Rack-mounted Servers

Rack-mounted servers will appeal to those who prefer tidiness and organization; they are essentially thin computers (from 2 to 12 inches wide) that are designed to be stacked vertically in a metal framework, or rack. This keeps things tidy and uses a minimum of space, so this form factor is definitely scalable. Rack-mounted equipment can be blocky, with sharp edges, however.

To assist in removing single points of failure, rack-mounted servers normally have dual power supplies. There are normally at least two network cards and in some cases



**Figure 2-2** Dell PowerEdge R515 rack server, courtesy of Dell, Inc.

management ports, and all of these connections are on the back of the device. Figure 2-2 shows ports on the front of the server; these allow a keyboard, mouse, video display, or a KVM (keyboard, video, mouse) switch to be connected.

## Racks


Full metal racks can be heavy, and server room and data center construction must account for that by including floors that can safely handle the weight. And because you don't want racks tipping over, balancing and bolting them to the floor is recommended.

Most racks are 19 inches wide (they do come wider), and they often use metal sliding rails that the servers and other types of appliances (such as network firewalls, storage enclosures, and so on) neatly fit into. Racks may not come with rails, so you may have to acquire rail kits. Rack-mounted equipment such as servers usually comes with a rail kit. (Be careful not to get pinched when inserting into or removing equipment from these sliding rails.)

Not all racks are equal, and they are available in different widths and heights. You may use one rack to accommodate servers and another to accommodate storage arrays. Rack-mountable devices use their own units of measurement for height, called rack units (U)—1U is 1 3/4 inches, 2U is 3 1/2 inches, 4U is 7 inches, and so on. This measurement is a standard that refers to the vertical distance between the holes in the rack to which rails and rack-mounted equipment are secured. A single rack-mounted device may be from 1U up to 7U in height. Most racks have a maximum of 42U (see Figure 2-3). When you're ordering rack-mountable servers, the specs will detail how many Us the enclosure requires so you can plan placement within the rack.

Now, because you can place many servers together in a rack, it just makes sense to place server storage appliances, power sources (including UPSs), and network cables in the rack, too. Too many times I've seen messy racks with cables hanging everywhere, both in front of and behind the rack; not only is this a safety hazard, but tracing cables when troubleshooting will be next to impossible!

**Figure 2-3**  
A 42U server  
rack can contain  
various pieces of  
equipment.



Power distribution units (PDUs) provide power outlets to racks in server rooms and data centers. To eliminate a single point of failure, redundant PDUs should be plugged into separate circuits. To extend this point, redundant server power supplies should each plug into separate PDUs. Data centers normally have alternate sources or providers of power, such as diesel generators, in the case of a power outage.

Because many different types of items can draw power from PDUs, you should check your PDU's rating to ensure that your equipment doesn't draw more power than the PDU's load capacity can accommodate.

### Cable Management Arm

A cable management arm is a metal or plastic folding component that is attached to the back of a rack-mounted device. All cables from the device (power, network, and so on) are fitted into the arm, which serves as a conduit or trench in which the cables are placed. You would also normally use cable zip ties to bundle together cables from a device. (A standard rack-mounted server has dual power supplies and at least two network cards, so that's four cables right away for a single device.) When you pull out a rack-mounted device (on the sliding rails), the cable arm expands so that you don't pull the cables out.

After your rack-mounted devices and cabling are in place, use rack fillers (blanking panels) to cover empty spaces in the rack. These, in addition to cable management arms, ensure that fan intake vents are not blocked. This can improve airflow, which can also save money by saving energy in the long run. Rack fillers are also measured using the U system and are available with venting holes.

## Blade Servers

Blade servers make me think of *Star Trek*: a technician inserts a highly sophisticated card into a slot to prevent the destruction of the *USS Enterprise*. And that's essentially what the blade server form factor is—a circuit board containing its own processors, memory, and network capabilities, and, in some cases, a small amount of storage, but no power supply or cooling mechanisms. Blade servers cannot run on their own. Most blades have a USB connector on the front in case you want to connect external components such as a DVD drive.

This small server form factor (Figure 2-4) will most likely replace rack-mounted servers at some point. Large data centers can increase their server density using blades, so scalability is not a problem. Like everything in IT, it's a tradeoff—sure, you can fit more blades than towers in a fixed amount of space, but towers are cheaper and easier to expand if you need expansion cards or additional storage.

### Blade Enclosure

A blade enclosure is a proprietary chassis that can house several blade servers from the same vendor, and it can measure from 6U to 12U. Blade servers slide into the blade enclosure. The enclosure provides

- Temperature control mechanisms, including fans
- Power

**Figure 2-4**  
Dell PowerEdge  
M520 blade  
server, courtesy  
of Dell, Inc.



- Network connectivity
- Storage connectivity
- Server remote management connections

Within the enclosure, the *backplane* connects server and I/O blades. I/O blades (or cards) can provide faster network connectivity, storage for blade servers, management capabilities, and other things. The *midplane* is a printed circuit board (PCB) with server blades that connect on one side (the front) and other components accessible on the other side (the back). If, for example, 10Gb network switching is required, you would have to use a specific midplane with this support.

Note that some manufacturers provide backplane and midplane redundancy to reduce single points of failure. Blanking panels are used where there are empty slots in the enclosure for better cooling and airflow.

## Server Components

Whether you are working with tower, rack-mounted, or blade servers, they all have components that give them specific functionality:

- Multiple processors (each with multiple cores)
- Memory (RAM)
- Storage (local and/or network accessible)
- Network connectivity (servers often have multiple cards)
- Management capabilities (for blade systems or hardware-level remote control)

*Firmware* is essentially software stored in a chip, and it's used all over the place—a server's motherboard BIOS, smartphones, and expansion cards, to name a few. Like operating system or application software, firmware comes in different versions with different capabilities and needs to be updated periodically. Updating server motherboard firmware can sometimes update the unique 128-bit universal unique identifier (UUID) that is used in firmware as a global server identifier for asset-tracking purposes.

Hardware problems can sometimes masquerade as software issues. For instance, flawed firmware code could cause server operating system instability. The solution is sometimes as simple as downloading and applying a firmware update from the manufacturer's web site.

Even though it is considered obsolete, some manufacturers *still* supply an MD5 hash value on the download web page that you can recompute after downloading to verify that the file, such as a firmware update, hasn't been changed. MD5 has been superseded by Secure Hashing Algorithm (SHA), so you're more likely to see SHA hashes than MD5. Plus, when updating firmware, you need to be sure you're applying the correct version of the update. The big guys (Dell, HP, IBM) often offer rollback options if you don't like the applied firmware update.

## BIOS

The basic input-output system (BIOS) is firmware built into a circuit board such as a motherboard or a RAID (Redundant Array of Inexpensive Disks) disk controller. BIOS has been around for decades. This is the magic that kicks in the moment you turn on the power for your server or when a card initializes.

When starting up, the server BIOS checks critical hardware components such as power, CPU, RAM, and video to make sure they are in place and functional. If the components are not functional, you'll get various beep codes or error numbers and messages, provided video is working. This is called the power on self test (POST). Assuming things are working, the BIOS then checks the master boot record (MBR) to hand control over to an installed operating system.

Disks are initialized on a computer as either MBR or GUID Partition Table (GPT), depending on operating system support. There is only one copy of the MBR on an MBR disk, and it sits on the first sector of the disk before the first disk partition. GPT disks store multiple copies of this data throughout the disk for additional resiliency.

MBR disks are limited to four partitions. GPT disks can have up to 128 partitions on a Windows system. So on a unified extensible firmware interface (UEFI) system, booting from a GPT disk is possible as part of the POST.

The complementary metal oxide semiconductor (CMOS) is essentially the specific configuration of hardware settings supported by the BIOS. For example, you might change the boot order on your server to boot from USB (requiring a password of course!) first, and then from a local hard disk. Or you might enable or disable the CPU execute disable (XD) bit. With this option enabled, the processor flags specific areas of memory where data can reside and code cannot execute. The BIOS has the capabilities, and the CMOS retains your configuration of those capabilities. Figure 2-5 shows a basic BIOS configuration screen.

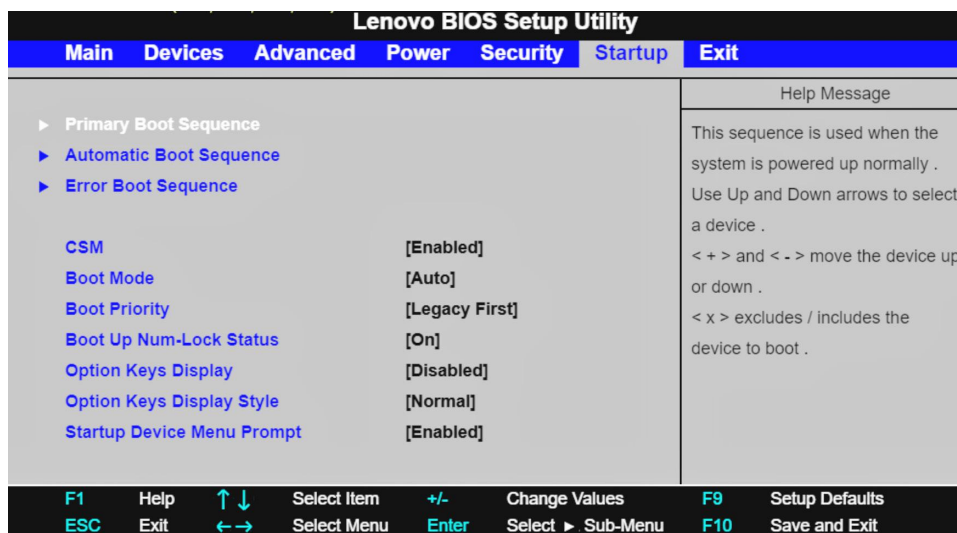


Figure 2-5 BIOS screen