



# Advanced Solutions of Microsoft SharePoint Server 2013



Exam Ref

70-332

Michael Doyle

# Exam Ref 70-332



Prepare for Exam 70-332—and help demonstrate your real-world mastery of Microsoft SharePoint Server 2013. Designed for experienced IT professionals ready to advance their status, *Exam Ref* focuses on the critical-thinking and decision-making acumen needed for success at the MCSE level.

## Focus on the expertise measured by these objectives:

- Plan Business Continuity Management
- Plan a SharePoint Environment
- Upgrade and Migrate a SharePoint Environment
- Create and Configure Service Applications
- Manage SharePoint Solutions, BI, and Systems Integration

## This Microsoft *Exam Ref*:

- Organizes its coverage by exam objectives.
- Features strategic, what-if scenarios to challenge you.
- Includes a 15% exam discount from Microsoft. Offer expires 12/31/2017. Details inside.

## Advanced Solutions of Microsoft SharePoint Server 2013

### About the Exam

Exam 70-332 is one of five Microsoft exams that will demonstrate your expertise in helping a company organize, sync, collaborate, and share information across the organization. Earning an MCSE: SharePoint certification will help qualify you for a job as systems or network analyst.

### About Microsoft Certification

The new **Microsoft Certified Solutions Expert** (MCSE) certifications demonstrate your ability to build innovative solutions across multiple technologies, both on-premises and in the cloud.

Exams 70-410, 70-411, 70-412, 70-331, and 70-332 are required for MCSE: SharePoint certification.

See full details at:  
[microsoft.com/learning/certification](http://microsoft.com/learning/certification)

### About the Author

**Michael Doyle**, MCTS, MCSD, MCSE, is a senior architect at a major company. He's been working with SharePoint for more than a decade. Michael is the author of *Microsoft SharePoint 2010: Customizing My Site* and coauthor of *Microsoft SharePoint Foundation 2010 Inside Out*.

[microsoft.com/mspress](http://microsoft.com/mspress)

ISBN: 978-0-7356-7810-1



9 780735 678101

**U.S.A. \$39.99**  
Canada \$41.99  
[Recommended]

Certification/Microsoft SharePoint Server



# **Exam Ref 70-332: Advanced Solutions of Microsoft SharePoint Server 2013**

Michael Doyle

Copyright © 2013 by Michael Doyle

All rights reserved. No part of the contents of this book may be reproduced or transmitted in any form or by any means without the written permission of the publisher.

ISBN: 978-0-7356-7810-1

1 2 3 4 5 6 7 8 9 QG 8 7 6 5 4 3

Printed and bound in the United States of America.

Microsoft Press books are available through booksellers and distributors worldwide. If you need support related to this book, email Microsoft Press Book Support at [mspinput@microsoft.com](mailto:mspinput@microsoft.com). Please tell us what you think of this book at <http://www.microsoft.com/learning/booksurvey>.

Microsoft and the trademarks listed at <http://www.microsoft.com/about/legal/en/us/IntellectualProperty/Trademarks/EN-US.aspx> are trademarks of the Microsoft group of companies. All other marks are property of their respective owners.

The example companies, organizations, products, domain names, email addresses, logos, people, places, and events depicted herein are fictitious. No association with any real company, organization, product, domain name, email address, logo, person, place, or event is intended or should be inferred.

This book expresses the author's views and opinions. The information contained in this book is provided without any express, statutory, or implied warranties. Neither the authors, Microsoft Corporation, nor its resellers, or distributors will be held liable for any damages caused or alleged to be caused either directly or indirectly by this book.

**Acquisitions and Developmental Editor:** Kenyon Brown

**Production Editor:** Kara Ebrahim

**Editorial Production:** Box Twelve Communications

**Technical Reviewer:** Kyle Davis

**Copyeditor:** Box Twelve Communications

**Indexer:** Box Twelve Communications

**Cover Design:** Twist Creative • Seattle

**Cover Composition:** Karen Montgomery

**Illustrator:** Rebecca Demarest

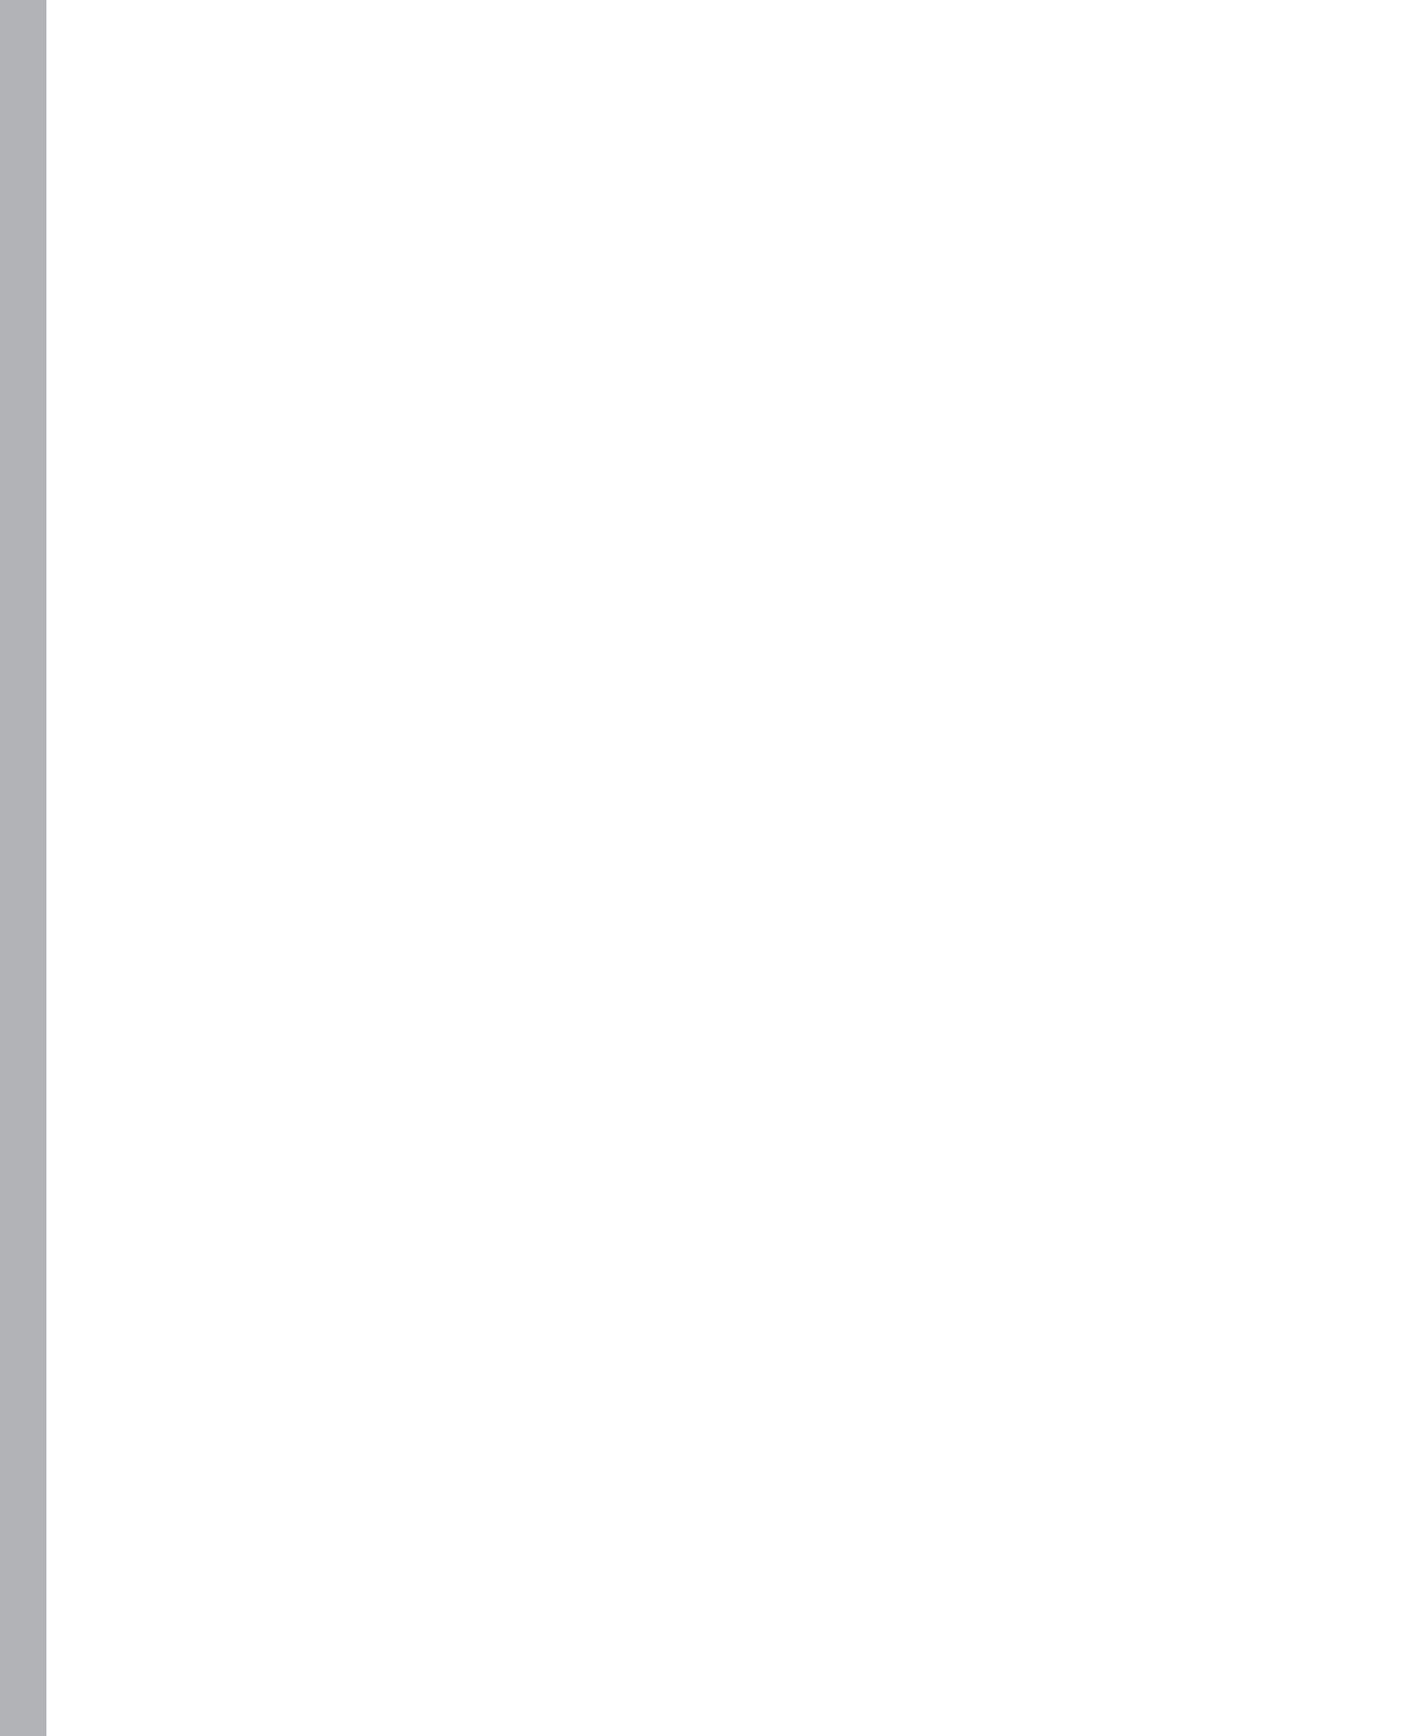
*This book is dedicated to David and Sandy Doyle.*

—MICHAEL DOYLE



# Contents at a glance

	<i>Introduction</i>	<i>xv</i>
	<i>Preparing for the Exam</i>	<i>xviii</i>
CHAPTER 1	<b>Plan business continuity management</b>	<b>1</b>
CHAPTER 2	<b>Plan a SharePoint environment</b>	<b>59</b>
CHAPTER 3	<b>Upgrade and migrate a SharePoint environment</b>	<b>149</b>
CHAPTER 4	<b>Create and configure service applications</b>	<b>209</b>
CHAPTER 5	<b>Manage SharePoint solutions, BI, and systems integration</b>	<b>291</b>
	<i>Index</i>	<i>351</i>
	<i>About the Author</i>	<i>373</i>



# Contents

<b>Introduction</b>	<b>xv</b>
<i>Microsoft certifications</i>	<i>xvi</i>
<i>Acknowledgments</i>	<i>xvi</i>
<i>Errata &amp; book support</i>	<i>xvi</i>
<i>We want to hear from you</i>	<i>xvii</i>
<i>Stay in touch</i>	<i>xvii</i>
<i>Preparing for the Exam</i>	<i>xix</i>
<b>Chapter 1 Plan business continuity management</b>	<b>1</b>
Objective 1.1: Plan for SQL high availability and disaster recovery . . . . .	1
Gathering requirements	2
Planning for SQL Server mirroring	6
Planning for SQL Server clustering	10
Planning for SQL Server AlwaysOn	11
Planning for SQL Server log shipping	15
Planning for storage redundancy	20
Planning for login replication	21
Objective summary	22
Objective review	23
Objective 1.2: Plan for SharePoint high availability . . . . .	24
Planning for service distribution	24
Planning for service instance configuration	26
Planning for physical server distribution	27
Planning for network redundancy	29

---

**What do you think of this book? We want to hear from you!**

Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

[www.microsoft.com/learning/booksurvey/](http://www.microsoft.com/learning/booksurvey/)

Planning for SQL Server load balancing	32
Planning for SQL Server aliases	34
Objective summary	37
Objective review	37
Objective 1.3: Plan backup and restore .....	38
Establishing a SharePoint backup schedule	38
Using Central Administration backup	39
Establishing a SharePoint backup schedule	42
Establishing a SQL Server backup schedule	43
Planning a non-production environment content refresh	43
Planning for farm configuration recovery	44
Planning for service application recovery	46
Planning for content recovery	47
Objective summary	51
Objective review	52
Chapter summary .....	53
Answers .....	54
Objective 1.1: Thought experiment	54
Objective 1.1: Review	54
Objective 1.2: Thought experiment	55
Objective 1.2: Review	56
Objective 1.3: Thought experiment	57
Objective 1.3: Review	57
<b>Chapter 2 Plan a SharePoint environment</b>	<b>59</b>
Objective 2.1: Plan a social workload .....	59
Planning communities	60
Planning My Sites	65
Planning social permissions	71
Planning user profiles	74
Planning activity feeds	77
Planning connections	81
Objective summary	85
Objective review	86

Objective 2.2: Plan and configure a search workload . . . . .	87
Planning and configuring search result relevancy	87
Planning and configuring index freshness	94
Planning and configuring result sources	95
Planning and configuring the end-user experience	99
Planning and configuring a search schema	102
Analyzing search analytics reports	107
Objective summary	109
Objective review	109
Objective 2.3: Plan and configure Web Content Management (WCM) .	110
Planning and configuring channels	111
Planning and configuring product catalogs and topic pages	113
Planning and configuring content deployment	119
Planning and configuring variations	122
Objective summary	127
Objective review	127
Objective 2.4: Plan an Enterprise Content Management (ECM)	
workload . . . . .	128
Planning and configuring eDiscovery	128
Planning and configuring record management	131
Planning and configuring record disposition and retention	135
Planning large document repositories	137
Objective summary	139
Objective review	139
Chapter summary . . . . .	140
Answers . . . . .	141
Objective 2.1: Thought experiment	141
Objective 2.1: Review	141
Objective 2.2: Thought experiment	142
Objective 2.2: Review	143
Objective 2.3: Thought experiment	144
Objective 2.3: Review	145
Objective 2.4: Thought experiment	146
Objective 2.4: Review	147

<b>Chapter 3 Upgrade and migrate a SharePoint environment</b>	<b>149</b>
Objective 3.1: Evaluate content and customizations . . . . .	149
Performing migration pre-check tasks	150
Analyzing content database test results	153
Configuring web application authentication for upgrades	155
Resolving orphan objects	158
Resolving missing files	160
Resolving configuration issues	162
Objective summary	164
Objective review	165
Objective 3.2: Plan an upgrade process . . . . .	166
Planning removal of servers in rotation	166
Configuring parallel upgrades	171
Configuring read-only access for content	173
Configuring upgrade farms	176
Measuring upgrade performance	179
Planning an installation sequence	181
Objective summary	182
Objective review	183
Objective 3.3: Upgrade site collection . . . . .	184
Performing a health check	184
Analyzing and resolving health check results	186
Planning and configuring available site collection modes	188
Planning and configuring site collection upgrade availability	190
Planning and configuring evaluation mode	191
Planning and configuring site collection upgrade queues and throttling	193
Objective summary	198
Objective review	199
Chapter summary . . . . .	200
Answers . . . . .	201
Objective 3.1: Thought experiment	201
Objective 3.1: Review	201
Objective 3.2: Thought experiment	203

Objective 3.2: Review	203
Objective 3.3: Thought experiment	205
Objective 3.3: Review	206

## **Chapter 4 Create and configure service applications 209**

Objective 4.1: Create and configure App Management . . . . .	209
Creating and configuring the App Store	210
Creating and configuring subscriptions	212
Configuring DNS entries	215
Configuring wildcard certificates	216
Objective summary	219
Objective review	219
Objective 4.2: Create and configure productivity services . . . . .	220
Creating and configuring Microsoft Excel Services	221
Creating and configuring Microsoft Access Services	231
Creating and configuring Microsoft Visio Services	239
Creating and configuring Microsoft Word Automation Services	242
Creating and configuring Microsoft PowerPoint Conversion Services	246
Creating and configuring Machine Translation Services	247
Objective summary	252
Objective review	252
Objective 4.3: Configure Service Application Federation . . . . .	253
Planning services to federate	254
Performing a certificate exchange	255
Managing trusts	256
Managing service application permissions	258
Publishing service applications	260
Consuming service applications	261
Objective summary	263
Objective review	264
Objective 4.4: Create and configure a Business Connectivity Service and Secure Store application. . . . .	265
Importing and configuring BCS models	265
Configuring BCS model security	270

Generating a Secure Store master key	272
Managing the Secure Store master key	274
Creating Secure Store target applications	276
Managing Secure Store target application permissions	278
Objective summary	280
Objective review	281
Chapter summary . . . . .	282
Answers . . . . .	283
Objective 4.1: Thought experiment	283
Objective 4.1: Review	283
Objective 4.2: Thought experiment	284
Objective 4.2: Review	284
Objective 4.3: Thought experiment	286
Objective 4.3: Review	287
Objective 4.4: Thought experiment	288
Objective 4.4: Review	288

**Chapter 5 Manage SharePoint solutions, BI, and systems integration 291**

Objective 5.1: Manage SharePoint solutions and applications . . . . .	291
Managing sandboxed solution quotas	292
Configuring Sandboxed Solution Management	295
Deploying farm solutions	296
Upgrading farm solutions	298
Deploying apps	299
Upgrading apps	301
Objective review	303
Objective 5.2: Planning and configuring a BI infrastructure . . . . .	304
Planning and configuring PerformancePoint	304
Planning and configuring Reporting Services	310
Planning and configuring PowerPivot	319
Planning and configuring Excel Services data modeling	326
Planning and configuring Power View	327
Objective summary	330
Objective review	331

Objective 5.3: Creating and configuring work management . . . . .	332
Configuring a connection to Microsoft Project Server	332
Configuring a connection to Exchange for eDiscovery and Resource Sharing	339
Objective summary	341
Objective review	341
Chapter summary . . . . .	343
Answers . . . . .	344
Objective 5.1: Thought experiment	344
Objective 5.1: Review	344
Objective 5.2: Thought experiment	345
Objective 5.2: Review	346
Objective 5.3: Thought experiment	347
Objective 5.3: Review	348
 <i>Index</i>	 351

---

**What do you think of this book? We want to hear from you!**

Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

[www.microsoft.com/learning/booksurvey/](http://www.microsoft.com/learning/booksurvey/)



# Introduction

---

This book is primarily intended to help you prepare for Exam 70-332, “Advanced Solutions of Microsoft SharePoint Server 2013,” but it’s also intended to be a reference that you can consult during your experiences with SharePoint Server 2013. In many cases, the steps to perform a task are shown to help you feel comfortable with related questions on the exam as well provide a reference on how to perform the task in a real-life situation.

The level of detail in this book will often exceed what’s required on the exam because of the very nature of it being an advanced solutions exam. This doesn’t mean you won’t come up with specific questions about the steps required to perform a task or requirements needed to install a service application. It does mean, however, that you don’t need to focus on being able to spell out a command correctly or know exactly what parameter to pass it. You should focus on the concepts, the overall steps involved with a task, and the components needed for a solution. If you focus on these concepts and go through the tasks in this book, you will be well on your way to passing the exam.

This book is generally intended for exam candidates that have four or more years working with SharePoint Server and related technologies, such as SQL Server and Windows Server. Candidates should have hands-on experience with a multi-server SharePoint farm in the capacities of planning, implementing, and maintaining. This includes, but isn’t limited to, the areas of high availability, disaster recovery, capacity planning, and exposure to SharePoint Online. Despite having multiple years of experience with a multi-server SharePoint farm, that you will have experience with all the technologies covered by the exam is doubtful; you should focus on the areas to which you have the least exposure. Also, any feature that has been added to SharePoint Server 2013 will likely receive additional coverage on the exam.

This book will help you prepare for the exam, but nothing can take the place of real-life experience. In an effort to make the exams closer to measuring knowledge of the product, they are going more and more to case studies and getting away from simple multiple-choice questions. You’ll still see a number of traditional multiple-choice questions, but you’ll also see questions where you have to place steps in order and questions where you have to choose the right set of items from a large list of possible items. In these cases, practicing the actual implementation of the functionality covered in this book will help you far more than just trying to memorize what is involved.

This book covers every exam objective, but it does not cover every exam question. Only the Microsoft exam team has access to the exam questions themselves and Microsoft regularly adds new questions to the exam, making it impossible to cover specific questions. You should consider this book a supplement to your relevant real-world experience and other study materials. If you encounter a topic in this book that you do not feel completely comfortable with, use the links you’ll find in text to find more information and take the time to

research and study the topic. Great information is available on MSDN, TechNet, and in blogs and forums.

## Microsoft certifications

---

Microsoft certifications distinguish you by proving your command of a broad set of skills and experience with current Microsoft products and technologies. The exams and corresponding certifications are developed to validate your mastery of critical competencies as you design and develop, or implement and support, solutions with Microsoft products and technologies both on-premise and in the cloud. Certification brings a variety of benefits to the individual and to employers and organizations.

### **MORE INFO** ALL MICROSOFT CERTIFICATIONS

For information about Microsoft certifications, including a full list of available certifications, go to <http://www.microsoft.com/learning/en/us/certification/cert-default.aspx>.

## Acknowledgments

---

I'd like to thank the following people. I would like to especially thank Stepha, Maxson, and Carianna Doyle for being so patient with their father through this process. Also, I would like to thank Dr. L.J. Geiken for putting up with me being book-focused for so long and for being a great partner. I would finally like to thank the following people in the SharePoint community for encouraging my efforts (even if they didn't realize they were doing it): Dan Holme, Randy Williams, Debbie Ireland, Nick Hadlee, Brian Farnhill, and Christian Buckley.

## Errata & book support

---

We've made every effort to ensure the accuracy of this book and its companion content. Any errors that have been reported since this book was published are listed on our Microsoft Press site at:

*<http://aka.ms/ER70-332/errata>*

If you find an error that is not already listed, you can report it to us through the same page.

If you need additional support, email Microsoft Press Book Support at [mssinput@microsoft.com](mailto:mssinput@microsoft.com).

Please note that product support for Microsoft software is not offered through the addresses above.

## **We want to hear from you**

---

At Microsoft Press, your satisfaction is our top priority, and your feedback our most valuable asset. Please tell us what you think of this book at <http://www.microsoft.com/learning/book-survey>.

The survey is short, and we read every one of your comments and ideas. Thanks in advance for your input!

## **Stay in touch**

---

Let's keep the conversation going! We're on Twitter: <http://twitter.com/MicrosoftPress>.



# Preparing for the Exam

---

Microsoft certification exams are a great way to build your resume and let the world know about your level of expertise. Certification exams validate your on-the-job experience and product knowledge. While there is no substitution for on-the-job experience, preparation through study and hands-on practice can help you prepare for the exam. We recommend that you round out your exam preparation plan by using a combination of available study materials and courses. For example, you might use the Exam Ref and another study guide for your "at home" preparation, and take a Microsoft Official Curriculum course for the classroom experience. Choose the combination that you think works best for you.

Note that this Exam Ref is based on publically available information about the exam and the author's experience. To safeguard the integrity of the exam, authors do not have access to the live exam.



# Plan business continuity management

Downtime is something that nobody wants to think about, but it's a reality that any organization needs to plan for. This is the first chapter in this book because business continuity is something that you should plan for before you start building your SharePoint farm. Microsoft SQL Server and Microsoft SharePoint work together to provide various options for business continuity depending on your business needs. With the proper planning, your business can achieve a high degree of disaster recovery no matter what the situation requires. It all depends on the effort put into the planning and the resources that can be allocated to this endeavor.

***important***  
***Have you read page xix?***

It contains valuable information regarding the skills you need to pass the exam.

## Objectives in this chapter:

- Objective 1.1: Plan for SQL high availability and disaster recovery
- Objective 1.2: Plan SharePoint high availability
- Objective 1.3: Plan backup and restore

## Objective 1.1: Plan for SQL high availability and disaster recovery

---

SQL Server is the foundation of any SharePoint 2013 farm. Getting your SQL installation ready is paramount to any disaster recovery or high-availability plan. Here, planning is of the utmost importance and is the area that's most likely at fault when something goes wrong. Poor planning at this stage can bring down the SharePoint farm for hours or even days. You can't plan for every potential disaster, but you can plan for items that you anticipate. Before starting with the objective topics, you need to understand gathering requirements and defining limitations.

**This objective covers the following topics:**

- Plan for SQL Server mirroring
- Plan for SQL Server clustering
- Plan for SQL Server AlwaysOn
- Plan for SQL Server log shipping
- Plan for storage redundancy
- Plan for login replication

## Gathering requirements

Gathering requirements is definitely an art. It's a balance between spending too little time (which ends up in a poorly designed server farm) or spending too much time, causing project delays, and going over budget. The term poorly designed, in this sense, doesn't mean that it won't meet the needs of the organization, but instead that it might not match the business needs but addresses issues that don't really exist. Having a disaster recovery plan that works for everyone would be nice, but that simply isn't the case. To get started with gathering requirements, you must know the relevant terminology. The two most important terms associated with SQL Server high availability are Recovery Time Objective (RTO) and Recovery Point Objective (RPO).



---

**EXAM TIP**

**You must be familiar with these terms and comfortable defining these requirements within the context of the SharePoint implementation.**

---

RTO is usually a number that's heard when talking about uptime, such as 99.9 percent or something higher. Everyone wants a number that approaches 100 percent, but each 9 that's added causes the cost to go up exponentially. The number of 9's required depend heavily on the type of data being stored and who's accessing it from where. SharePoint 2013 isn't used to store transactional data, such as that stored by banks that process thousands of transactions every minute, so it doesn't have the same requirements. However, the content stored in SharePoint is often mission critical and requires a high RTO. When figuring your RTO requirements, you need to keep the business environment in context. Consider the following important questions:

- Does the data need to be available 24 hours a day, seven days a week, or are maintenance windows allowed?
- Is all content to be treated the same, or does different content have different sets of requirements?
- How is data to be archived when it's no longer timely?

RTO can't be determined just from a technological viewpoint. The opinions and requirements of business stakeholders also need to be taken into consideration. Also, the budget allocated to RTO must be considered. All these factors go into determining a realistic RTO.

Truly understanding RTO requires doing some math. Assuming that you're dealing with an organization that runs 24 hours a day, seven days a week, you can do the following math:

- An RTO of 99.9 percent means a downtime total of 8.76 hours in the whole year (365 days × 24 hours per day × .001).
- An RTO of 99.99 percent means a downtime total of less than 53 minutes in a year.
- An RTO of 99.999 percent means that less than 6 minutes of downtime is allowed for a whole year.

You can easily see that even the addition of one 9 can dramatically affect the high availability of your organization. For example, an RTO of 99.99 percent doesn't allow for the cumulative update of a server that might bring down the system. Therefore, you must have a plan that allows servers to be brought down without affecting the functionality of the system.

RPO isn't discussed as often but is just as important during the requirement gathering phase. You must have a true understanding of the data involved to determine the RPO. When determining the RPO, you should take into consideration the following details:

- Amount of data lost (for example, 30 minutes of loss)
- Cost to the company of lost data
- Cost to the company for the amount of time to recover

This means that if an outage causes a loss of 30 minutes of data, requires an hour to come back online, and then requires 30 minutes to replace the lost data, you are looking at an RPO of 2 hours. This reflects true outage time because it calculates the amount of lost productivity.

When you gather requirements, calculating and translating the number of lost hours into a dollar amount often helps. This is simply the RPO times the number of people affected times the average salary per hour of the people affected. You can use this kind of information to help gather support for high-availability initiatives.

## Choosing a version of SQL Server

Because SQL Server comes in many versions, when planning high availability you need to determine which version to use. When installing SQL Server for a SharePoint 2013 farm, you currently have two version choices:

- The 64-bit version of Microsoft SQL Server 2012
- The 64-bit version of SQL Server 2008 R2 Service Pack 1

Both versions have options for high availability and represent viable installations for a comprehensive disaster recovery plan. The standalone version of SharePoint 2013 isn't part of this discussion because it shouldn't be used in an environment that requires high availability.

Future versions and revisions of SQL Server should work with SharePoint 2013, but that's not guaranteed. Any upgrade or update should be tested in a non-production environment before it's installed.

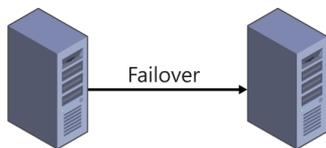
You need to choose a particular SQL Server version for many reasons. In many cases, moving to the most recent version isn't an option for various reasons including cost, ability to provide technical support, or lack of testing by the organization. However, one primary reason to choose SQL Server 2012 in regards to high availability is multi-subnet failover clustering, as defined in the next section.

#### **MORE INFO** HARDWARE AND SOFTWARE REQUIREMENTS FOR SHAREPOINT 2013

See <http://technet.microsoft.com/en-us/library/cc262485.aspx#section4> for more information about hardware and software requirements for SharePoint 2013.

## Understanding multi-subnet failover clustering

Multi-subnet failover clustering allows each cluster node to reside in a completely different subnet. As a result, clusters can be farther away on the network as well as geographically dispersed. Geographically dispersed clusters are often referred to as *stretch clusters*. Because these clusters can't share data, a replication scenario must be enabled for the data in the cluster. Figure 1-1 shows a multi-subnet cluster.



Node 1: 10.10.5.5      Node 2: 20.20.8.8

**FIGURE 1-1** Example of multi-subnet failover cluster

Clustering—just one of several options that enable high availability—is discussed later in this chapter. Having multiple servers in each node is possible and, depending on the configuration, SQL Server can come online as long as one IP address is available.

## Determining limitations

To avoid spending time and money designing a plan that can't be implemented, you need to know all the limitations that can affect your high-availability plan. Limitations can be grouped into two categories: non-technical and technical. Both must be considered in the creation of your high-availability plan.

Most limitations that you encounter have nothing to do with technology and can include but definitely aren't limited to the following:

- **Budget** Your budget can range from pessimistic to optimistic.

- **Power availability** Servers and backup devices require power that might or might not be available.
- **Heat** More servers might overheat the room, causing server failure.
- **Space availability** In the server room, racks require space that might not exist.
- **Training** This is critical to successful high-availability plans.
- **Network bandwidth** Different high-availability scenarios have varying bandwidth requirements.
- **Manpower** Servers require upkeep.
- **Politics** This is always involved.
- **Time** Implementing any high-availability scenario requires time.

Many of these limitations are out of the control of the individuals implementing the high-availability plan. Therefore, knowing these up front and which ones can be altered is important. For example, if the server room is in downtown Tokyo and is limited to a small space that's already full and challenging to cool, you need to take such limitations into consideration if the funds aren't available to rent more office space. The cloud helps alleviate many of these issues and should be considered, especially if one of the non-technical limitations becomes a critical issue.

After you identify the non-technical issues but before you plan for individual high-availability components, you need to look at the technical limitations. Some of these might include the following:

- **The recovery model** The simple recovery model doesn't allow for database mirroring, log shipping, and the ability to take backups of the transaction logs.
- **Network bandwidth and latency** Low bandwidth and latency can prohibit the successful implementation of most high-availability options.
- **FILESTREAM requirements** *FILESTREAM* can't be used with database snapshots or database mirroring. It can cause issues with other items such as transactional replication, log shipping, and backups.
- **Software limitations** Certain versions of SQL Server are required for some high-availability options. Make sure that the required version is available.

Now that you've looked at what's needed to gather requirements and identify limitations, you can concentrate on some of the high-availability options in SQL Server.

#### **NOTE SINGLE-LABEL DOMAIN LIMITATIONS**

SharePoint 2013 doesn't support installations on single-label domains. Many features don't work correctly, including user profile import. For more information, refer to the Knowledge Base article at <http://support.microsoft.com/kb/2379373>.

## Planning for SQL Server mirroring

SQL Server mirroring provides an almost complete redundancy of the data used by SQL Server—*almost* because it depends on how SQL Server mirroring is configured. Another benefit is that SharePoint has fully supported database mirroring since version 2010, especially concerning content databases. Database mirroring is also fairly simple to set up and provides the following benefits:

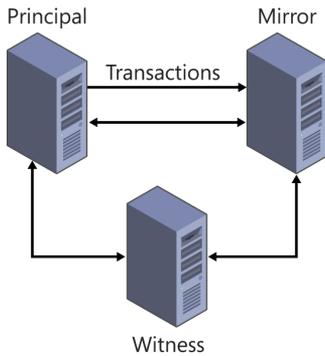
- It increases data protection. Because a full copy of the database is available, mirroring greatly increases the likelihood of data survival.
- Database availability improves because SharePoint can use either database.
- Availability during upgrades means that you can upgrade one database at a time in a mirror.

As the benefits show, database mirroring provides a valid high-availability solution. It's also one of the most commonly used methods of providing high availability. The main drawbacks are cost (because it requires, at minimum, another SQL Server) and disk space (because it's a full duplicate, the storage requirements are now doubled). Assuming that you have the spare SQL Server (it doesn't have to be dedicated to this function) and the disk space, you still need to understand some terms so that you have a better understanding of how SQL Server mirroring works:

- **Principal database** This read/write database provides transaction log data to the mirrored database.
- **Principal server** The principal database resides on this server.
- **Mirror database** This database is a copy of the principal database.
- **Mirror server** The mirror database resides on this server.
- **High-performance mode** Because the database mirroring session operates asynchronously, data might be lost. This mode supports only forced switching.
- **High-safety mode** The database mirroring session operates synchronously and optionally uses a witness.
- **Forced switching** This manual process switches the database functionality to the mirrored database until the principal database can be recovered. SharePoint can be set up to do this automatically for content databases.
- **Witness** This SQL Server instance watches the principal and mirror to determine when a failover is required. It doesn't need to store data.
- **Transaction safety** This database setting determines whether a database mirroring session operates asynchronously or synchronously. The only options are FULL and OFF.
- **Manual failover** This occurs when a database owner forces the mirror to act at the principal, even though the principal is fully functional (as in a database upgrade situation).

- **Automatic failover** This occurs when a witness sees that the principal isn't functioning properly and causes the mirror server to become the primary server.

Looking at the terminology can help you determine what's happening in Figure 1-2. All three servers are communicating with each other, but only transaction log data is being sent from the principal server to the mirror server.



**FIGURE 1-2** A high-safety database mirroring session with a witness

If you haven't dealt with database mirroring before, this can seem like a lot to take in at one time, but with a little practice it will soon make sense. Before you start actually setting up a mirror, you need to consider some significant requirements for database mirroring.

## Determining mirroring requirements

Planning for high availability generally involves working with multiple groups. Normally, the person responsible for planning the SharePoint installation isn't the same person overseeing the SQL Server instance. The database administrator(s) involved must be aware of the following mirroring requirements:

- Full recovery model is required.
- Recommend latency of less than 1 millisecond is required.
- Recommend network bandwidth of at least 1 GB per second is required.
- Transaction logs can grow very large, so you need a plan for truncating them.
- The principal and mirror database must run the same version and same edition of SQL Server, although the witness can be any version, even SQL Server Express.
- Each mirroring session needs at least two threads, so make sure enough threads are available for SQL Server to run.

Latency and network bandwidth aren't strict requirements but can affect performance and lead to large losses of data if a catastrophic failure occurs on the principal database. You can adjust the number of threads in SQL Server, but each thread requires memory.

After you meet the preceding requirements, you can start creating some mirrors and testing them. You can set up mirrored pairs on existing SQL Server instances, but practicing in a production environment isn't recommended. Creating a mirror involves three main steps:

1. Back up the principal database and transaction log.
2. Restore the backed-up database on the mirror server.
3. Create the mirror connection.

The following section explains these steps in more detail.

## Setting up a database mirror

You can create a database mirror using T-SQL commands or PowerShell commands, but for the purposes of illustration and for the exam, the following steps use SQL Server Management Studio:

1. Connect to the server that will serve as the principal server with SQL Server Management Studio, regardless of which version is being used.
2. Expand the database in the Object Explorer pane and right-click the database that is to be mirrored.
3. Choose Tasks and then choose Backup to launch the Backup Wizard.
4. Create a full backup of the principal database (to tape or disk, but disk is usually the easiest) as well as the transaction log.
5. Close the connection to the principal server.
6. Connect to the mirror server with SQL Server Management Studio.
7. Right-click the databases folder in the right navigation pane and choose Restore Database.
8. Restore the backed-up principal database to the mirror server. In the Options panel, select Restore With No Recovery and click OK.
9. Right-click the restored database after it's finished being restored. Choose Tasks and then navigate to the Restore submenu.
10. Select Transaction Log to open the Restore Transaction Log window.
11. Choose the principal database's backed-up transaction log.
12. Select Options in the Select a Page navigation panel and choose Restore With No Recovery.
13. Click OK to restore the transaction log.
14. Right-click the mirror database in the databases folder and choose Properties.
15. Click Mirroring in the Select a Page navigation pane.

16. Click Configuration Security to launch the Configure Database Mirroring Security Wizard.
17. Click Next to move to the Choose A Witness Server page. Choose the appropriate option depending on whether you want a witness server.
18. After you finish choosing options, click Next to open the Principal Server Options page.
19. Make sure the values for the principal server are correct and click Next.
20. On the Mirroring Server Instance page, select the mirror server if it isn't already there and click Next.
21. Click Connect to start the connection.
22. When the connection is established, click Next to open the Service Accounts page.
23. Enter the service account to use or leave blank to use the same service account.
24. Click Next to open the Complete the Wizard page.
25. Verify the settings and click Finish to save the settings.
26. On the Database Properties window, click Start Mirroring to begin the mirroring process.
27. After verifying that mirroring has started, click OK to close the dialog box. Both the mirrored and principal database should show that they are being mirrored in the database folder by having the word *Synchronized* next to them.

A database mirror should have been established if you followed these steps correctly and set up the SQL Server computers correctly.

Unless a witness is established, the mirror must be manually promoted to be the primary server. The primary database can be mirrored to an additional two SQL servers for a total of four copies. The additional demands on the memory and network need to be taken into consideration if this method is pursued. Each time a transaction occurs, it must be written to disk and then packaged up and sent across the network to be written again to another set of disks.

#### **REAL WORLD GEOGRAPHICALLY DISPERSED MIRRORING**

I wanted to establish a database mirror for a production database located on the west coast of the United States with a SQL Server node located in London. It had a relatively low bandwidth and latency was relatively high, but I decided to try database mirroring because it would provide a geographically diverse location for the data and read-only queries. I did a test on a small test database and could mirror it, but the production database was bigger than 100 GB. The mirror was never successfully established due to the latency and bandwidth limitations of the connection. If something similar happens to you, log shipping might be an option, depending on the business requirements.

## Using permissions for database mirroring

SQL Server uses Transmission Control Protocol (TCP) to communicate between the servers involved in the mirroring process. TCP handles the moving of the transaction data as well as the monitoring of server health if a witness is being used. SQL Server database mirroring supports NT LAN Manager (NTLM), Kerberos, and certificates. Authentication occurs at the port level when the session is opened. Unless the network is secure, some sort of security should be used. SQL Server supports two types of encryption: Advanced Encryption Security (AES) and RC4.

### **MORE INFO** DATABASE MIRRORING TRANSPORT SECURITY

See <http://go.microsoft.com/fwlink/p/?LinkId=83569> for more information on Database Mirroring Transport Security.

Regarding permissions, SharePoint requires additional steps for database mirroring to work. Setting up a database mirror doesn't add the permissions required on the master and msdb databases. The Central Administration application pool account (as well as members of the farm administrators group) requires SQL Server logins and the following two fixed server roles be added to both databases:

- dbcreator
- securityadmin

Also, any application pool accounts, service application accounts, and default access accounts should have SQL Server logins on the mirrored database server. In essence, whatever SQL Server permissions are required by SharePoint on the principal server will also be required on the mirror in case of a failover. Creating a script to transfer these permissions is considered a best practice.

### **IMPORTANT** PERMISSIONS FOR DATABASE MIRRORING

A successfully created database mirror without the appropriate SharePoint permissions can result in an unsuccessful failover.

## Planning for SQL Server clustering

Planning for a SQL Server cluster involves considerably more planning than for SQL Server mirroring. With mirroring, you can use existing SQL Server computers and choose which database or database set you want to mirror. This allows selective high availability and reduced resource requirements. With a cluster, it is all or none.

A cluster is made of two or more SQL Server nodes that act as a single SQL Server instance. Therefore, the resource requirements are the same for each cluster node, except for the data storage, which is shared among all nodes. To get a better understanding of a cluster, look at what makes up a SQL Server cluster instance:

- A network name for the cluster instance
- One or more IP address for the cluster instance
- A combination of one or more disks in a Microsoft Cluster Service group
- At least one SQL Server instance with SQL Server Agent, Full-text Search service, Replication, and the SQL Server service
- Registry keys that provide checkpoints across the nodes
- Specific DLLs that handle the failover process

A SQL Server failover cluster is referred to by a single name. You don't connect to the individual servers but to the cluster instance name. On the network, the cluster appears as a single server. You must connect to the instance name rather than try to connect to a SQL Server's IP address or machine name. Internally, one node does all the interactions with applications and synchs up the other nodes. However, any of the nodes can become the primary node.

SQL Server clustering has specific hardware requirements, especially if the clusters are geographically dispersed. Latency and network bandwidth are limitations that can cause a cluster to fail. You can make a cluster work with unapproved hardware, but doing so isn't supported.

**MORE INFO WINDOWS CLUSTERING AND GEOGRAPHICALLY DISPERSED SITES**

See <http://go.microsoft.com/fwlink/?LinkId=116970> for more information on Windows clustering and geographically dispersed sites.

Storage Area Network (SAN) solutions work with SQL Server clusters, but some SAN solutions are incompatible with clustering. Check with the SAN manufacturer and with Microsoft to determine whether the available SAN solution is compatible.

## Planning for SQL Server AlwaysOn

AlwaysOn is the new high-availability option available with SQL Server 2012. It's intended to provide a comprehensive disaster recovery and high-availability package to minimize downtime in organizations. AlwaysOn comes in two varieties:

- **AlwaysOn High Availability Groups** replace database mirroring.
- **AlwaysOn Failover Cluster Instances** replace database clustering.

Choosing which solution is best is similar to choosing between database mirroring and database clustering. The requirements are similar to their predecessors but provide additional enhancements. Because each solution is so different, this chapter looks at each one independently. AlwaysOn has had some improvements that might be significant enough to warrant an upgrade. Both AlwaysOn solutions rely on the Windows Server Failover Clustering (WSFC) infrastructure, which, combined with SQL Server, works to provide a robust high-availability platform.

## AlwaysOn High Availability Groups

The AlwaysOn High Availability Groups technology is intended to replace database mirroring. It has all the benefits of database mirroring plus many more. Because many of the issues with hardware limitations have been addressed, it's a far more robust product than what was available in SQL Server 2008. Database mirroring is still available in SQL Server 2012, but it's being phased out over the next few releases. Some new features of AlwaysOn High Availability Groups are as follows:

- It can fail over a whole group of databases instead of just one at a time.
- It uses a virtual name to provide faster failover.
- Multiple secondary servers are supported, up to four.
- Asynchronous active log synchronization is available for solutions that require high latency.
- Built-in encryption and compression provide more secure and faster data transfer.
- Automatic page repair is available for corrupted pages.
- It provides a flexible failover policy with more control.
- It uses two synchronous secondary servers.

As you can see, quite a few features have been added to database mirroring. From this version forward, it's the recommended solution for high availability.



---

### **EXAM TIP**

**A database can use AlwaysOn High Availability Groups or database mirroring, but not both.**

---

Using AlwaysOn High Availability Groups is a much more involved solution to implement and requires considerable configuration both on SQL Server as well as the Windows Server that it's installed on. You need to know the prerequisites involved when deciding whether AlwaysOn should be part of the SharePoint farm's high-availability solution. The prerequisites for configuring AlwaysOn High Availability Groups are as follows:

- A server with AlwaysOn installed can't be a domain controller.
- It can't be installed on a server with WOW64 (Windows 32-bit on a 64-bit server).
- Each server in the group has to be a node in a WSFC cluster.
- The WSFC cluster must contain sufficient nodes to support the availability group configurations.
- All applicable Windows hotfixes have been installed on every node in the WSFC cluster.

### **MORE INFO** ALWAYS ON HIGH AVAILABILITY GROUPS

See [http://msdn.microsoft.com/library/ff878487\(v=sql.110\).aspx](http://msdn.microsoft.com/library/ff878487(v=sql.110).aspx) for more information on prerequisites, restrictions, and recommendations for AlwaysOn High Availability Groups.

WSFC is an advanced topic, but a general understanding of it is important to grasp the concepts of AlwaysOn High Availability Groups as well as AlwaysOn Failover Cluster Instances. The term *clustering* describes several different technologies used in planning and configuring SharePoint farms. Clearly stating which type of clustering is being used helps avoid confusion during the planning process.

WSFC is a Windows Server technology that supports high availability and disaster recovery for applications that sit on top of it, such as SQL Server and Exchange Server. Whenever a node in a cluster fails or is turned off for maintenance, it can automatically (or manually) transfer those services to another node in the cluster. In SQL Server, an AlwaysOn high-availability group becomes a WSFC cluster resource. WSFC then monitors the health of each node and initiates the failover processes. Although knowing every detail of how to configure a WSFC instance isn't necessary, understanding the terms helps with an overall comprehension of high-availability options. Some terms used with WSFC are as follows:

- **WSFC cluster** A group of independent servers that work together to provide high availability of services and applications such as SQL Server, Exchange, and/or SharePoint.
- **Failover cluster instance** A Windows Server service that manages an IP resource, a network name resource, and the application and/or services to be managed by the cluster. Rather than connect to an IP or computer name, an application connects to the IP resource or network name resource that crosses over all servers in the cluster.
- **Node** A server that's part of the cluster.
- **IP resource** An IP address (or IPv6 address) that refers to the cluster as a single entity.
- **Network name resource** Used in place of a server name or network name so that applications and servers can refer to the cluster as a whole.
- **Resource group** The set of resources typically required to manage an application or service as a single entity. Failover and fallback are done as a resource group.
- **Cluster resource** An object that can be "owned" by a node. It can be made online, offline, or transferred by a node but not shared by nodes.
- **Preferred owner** The node in a cluster that a resource group prefers to run on.
- **Possible owner** A node that a resource group can run on if the preferred owner is unavailable.

You can see, in these terms, many similarities between "typical" SQL Server clustering and WSFC. AlwaysOn High Availability Groups might seem to have more in common with traditional clustering than with database mirroring. In reality, the technology has some aspects of both, plus additional enhancements. One way to think about AlwaysOn high availability is that it works with WSFC to tie several resources together—database name, IP resource, network name resource, and SQL Server instance—so that they can act as a single entity regardless of what node is actively providing them.

### **MORE INFO** WINDOWS SERVER FAILOVER CLUSTERING WITH SQL SERVER

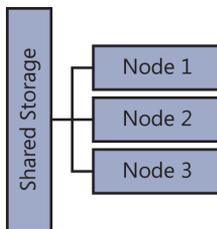
See <http://msdn.microsoft.com/en-US/library/hh270278.aspx> for more information on using WSFC with SQL Server.

## AlwaysOn Failover Cluster Instances

The AlwaysOn Failover Cluster Instances (AlwaysOn FCI) technology replaces the SQL Server clustering technology available in earlier versions. As with any clustering technology, you need to consider many hardware and software issues before deploying. As with SQL Server clustering specifically, the entire instance fails over rather than just a single database or group of databases—something you need to consider when choosing a high-availability option. The following enhancements are available with SQL Server 2012 AlwaysOn failover clustering:

- Localized tempdb for improved performance
- Multi-site clustering across subnets to improve site protection (subnets are generally in the same physical location)
- Flexible failover policy
- Improved diagnostics

As with clustering, an AlwaysOn FCI relies on a single shared storage area, such as a WSFC group, a SAN, or a file share. This way, any one node can immediately take over for another node using the same set of data. Figure 1-3 shows how this can be represented.



**FIGURE 1-3** A visual representation of how each node in a cluster uses the same shared storage

Because AlwaysOn FCI uses shared storage, this represents a potential single source of failure. The underlying storage implementation is responsible for providing data redundancy, such as mirrored disks in physically different locations. An AlwaysOn failover cluster can work with an AlwaysOn high-availability group to provide both instance and database high availability.

### **IMPORTANT** AUTOMATIC FAILOVER

If an AlwaysOn FCI is in an AlwaysOn high-availability group with another cluster node, automatic failover isn't supported.

An AlwaysOn FCI runs in a WSFC resource group. Every node within the cluster maintains its own synchronized copy of the configuration data and checkpoint registry keys. Only one of the nodes “owns” the resource group at any time, but any node can be the owner. In case of a failure or an upgrade, another node takes over ownership. Failure can be the result of hardware failure, operating system failure, or an application. The SQL Server binaries are stored locally on each node for improved performance, but the services are managed by the WSFC and not started automatically.

A virtual network name (VNN) for the entire node provides a single connection point for applications to connect to the cluster. The VNN points to a single node in the cluster. In the case of a failover, the VNN points to a new node. The application is unaware of a node failure, and no modification is required on the client application for continued uptime.

AlwaysOn failover clustering can use multiple subnets for its nodes. Because a single IP address is required to provide unattended failover, a virtual IP address needs to be assigned to each subnet in the failover cluster. During a subnet failure, the VNN is updated in the DNS to point to the new subnet. Clients require no changes to point to the new subnet, but a slight amount of downtime will occur depending on the DNS configuration.

#### **MORE INFO** ALWAYS ON FAILOVER CLUSTER INSTANCES

See [http://technet.microsoft.com/en-us/library/ms189134\(v=SQL.110\)](http://technet.microsoft.com/en-us/library/ms189134(v=SQL.110)) for more information on AlwaysOn FCIs.

## Planning for SQL Server log shipping

SQL Server log shipping can be confused with database mirroring because both have similar functionality: The data from a primary server is sent to one or more secondary servers. Before looking at using SQL Server log shipping, you should look at the primary differences:

- Failover is always manual for log shipping.
- An infinite number of servers can receive log shipments.
- Failover duration is slow (many minutes) compared to database mirroring (seconds).
- In log shipping, transaction logs are backed up and then applied to secondary servers, rather than use TCP to transport changes.
- Role change is manual.
- Log shipping uses both committed and uncommitted transactions, whereas database mirroring only uses committed.

Here, you can see that the main points of concern in high availability are the necessity for a manual change of roles and the amount of time required for the secondary server to become the primary.

Because database mirroring and log shipping are so similar, knowing the terminology is important so that you can distinguish between the two. The following terms are used when discussing log shipping:

- **Primary server** This server contains the primary database.
- **Primary database** This database is backed up and shipped to the secondary databases. Also, all log shipping configuration is done on this database.
- **Secondary server** This server contains the secondary database.
- **Secondary database** This database receives the transaction logs and provides a copy of the database that can be used if the primary is unavailable. It has two states: RECOVERING and STANDBY. STANDBY provides a limited read-only version.
- **Monitor server** This optional SQL Server instance monitors the health of the log shipping process and provides backup information such as time of backups and alerts.
- **Backup job** This SQL Server Agent job, named Log Shipping Backup on the primary server, performs the backing up and notifying of the monitor.
- **Copy job** This SQL Server Agent job, named Log Shipping Copy on the primary server, transfers the backups to the secondary server(s) and notifies the monitor.
- **Restore job** This SQL Server Agent job, named Log Shipping Restore on the secondary server, restores the backup to the secondary, performs cleanup, and notifies the monitor.
- **Alert job** This SQL Server Agent job, named Log Shipping Alert, exists on the monitor. It sends alerts when a backup or restore job doesn't happen within a specified threshold.

**NOTE REMOVING A MONITOR SERVER**

If a monitor server is configured, you must remove log shipping from the primary before you can remove the monitor server.

You can always use log shipping to provide an additional copy of a database. If any databases need to be archived periodically, need a very high level of disaster recovery, or require remote read-only access, log shipping is an excellent option because it's relatively easy to set up compared to the AlwaysOn options. Before configuring log shipping, however, make sure that the following prerequisites are met:

- The full or bulk-logged recovery model is used for the primary database.
- An existing share for the backups must be available to the primary and secondary servers.

To get a better understanding of high-availability options, look at the steps on how to configure log shipping. You could use T-SQL to configure log shipping, but the following steps show only the method available via SQL Server Management Studio:

1. Open SQL Server Management Studio and connect to the database that will serve as the primary.
2. Right-click the database that you're configuring for log shipping and choose Properties.
3. Under the Select A Page section, click Transaction Log Shipping.
4. The right side of the dialog box shows the transaction log shipping options. Select the check box next to Enable This As A Primary Database In A Log Shipping Configuration, as shown in Figure 1-4.



**FIGURE 1-4** Selecting Enable This As A Primary Database In A Log Shipping Configuration so that the rest of the settings can be configured

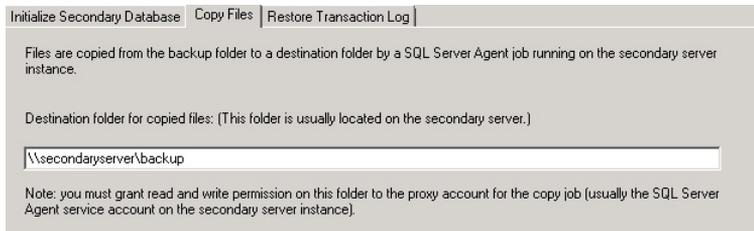
5. Click Backup Settings to display the Transaction Log Backup Settings dialog box.
6. Enter a network path (such as `\\fileserver\backup`) of a location for the transaction logs and choose a local path if the folder is on the same server as the SQL Server. The SQL Server service account needs read and write permissions to the folder, and the SQL Server Agent service accounts of all secondary servers need read access.
7. Change the settings for Delete Files Older Than and Alert If No Backup Occurs Within to fit the requirements, or leave them as they are for now.
8. The backup job is given a default Job Name and schedule. Change the schedule, if you want, and click OK after reviewing all settings to create a SQL Server Agent job.

#### **MORE INFO BACKUP COMPRESSION**

Backup compression is available if SQL Server 2012 (or SQL Server 2008 Enterprise) is used. See <http://msdn.microsoft.com/en-us/library/bb964719.aspx> for more information on backup compression in SQL Server 2012.

9. Click Add in the Secondary Server Instances And Databases section.
10. Click Connect in the Secondary Database Settings dialog box to connect to the SQL Server instance that will serve as the secondary database.
11. In the Secondary Database field, choose an existing database or type the name of a database to create.

12. On the Initialize Secondary Database tab, choose one of the options. The default option generates a full backup of the primary database and restores it into the secondary database. It also creates the database if it doesn't exist. If Restore Options is chosen, folders other than the default can be chosen for the secondary database location of the data and log files.
13. On the Copy Files tab, enter the destination folder for the copied transaction logs (see Figure 1-5).



**FIGURE 1-5** The Copy Files tab, on which you enter a location for the transaction log files

14. On the Restore Transaction Log tab, choose either No Recovery Mode or Standby Mode. Choosing Standby Mode disconnects anybody connected to the database when the restore of the transaction logs occurs. You also can choose to delay the backup (to counteract logical errors) and to have alerts sent if the restore doesn't occur within a given timeframe.
15. Click OK to close the Secondary Database Settings dialog box.
16. To add more secondary server instances, click the Add button.
17. If a monitor will be used to monitor the health of the primary server and the success of the transaction log shipments, select Use A Monitor Server Instance and configure the settings on the Settings tab. You can choose to impersonate the proxy account or use a specified account, as well as specify how long to keep the history of the jobs.

#### **IMPORTANT CHOOSING A MONITOR**

If you don't choose a monitor at this point, you must remove the transaction log configuration before you can add a monitor.

18. Click OK to start the configuration process.

#### **IMPORTANT ALTERNATE BACKUP METHODS**

If another job or maintenance plan is used to back up the transaction logs of the log shipping primary database, SQL Server Management Studio can't restore the transaction logs on secondary servers.

Failing over to a secondary database isn't as straightforward in log shipping as it is in database mirroring or one of the AlwaysOn technologies because of the time lag between the primary and secondary servers syncing. Typically, the secondary database needs to be synchronized with the latest transaction logs before it assumes the duties of the primary server (unless a logical error occurs in which part of the database is corrupt and getting an earlier version of the data is necessary). Some transaction logs might not have even been copied to the secondary server and will need to be recopied (if possible). Failing over requires some additional manual steps, as follows:

1. Make sure that all backup files are available on the secondary server.
2. Apply transaction logs in order, starting with the oldest unapplied backup to the most recent backup.
3. Assuming that the primary server is still available, back up the transaction tail *WITH NORECOVERY* to leave it in a usable state.
4. Apply the transaction log tail to the secondary server to bring it up to date, if available.
5. Recover the secondary server and make it available to users.
6. Point applications to the secondary server, which can also be set up to be the primary server in log shipping.

#### **MORE INFO FAIL OVER TO A LOG SHIPPING SECONDARY**

See <http://msdn.microsoft.com/en-us/library/ms191233.aspx> for more information on how to fail over to a log-shipping secondary.

Figuring out when logical errors occurred in the database requires some effort. As long as a backup and the transaction logs are available, getting to a point before the logical error occurred should be possible. This is an important component in many disaster recovery plans because it minimizes the amount of lost data.

## **Database mirroring with log shipping**

One major benefit to SQL Server log shipping is that it can be combined with database mirroring. This allows for redundancy not only in database storage but also in methods of maintaining that redundancy. For example, something happening to the database mirror (such as a corrupted DLL or latency issues) wouldn't stop log shipping from occurring. The net benefit is that you could still maintain up-to-date copies of your data. An additional benefit is that you can delay shipping logs in case an error occurs in the database itself that you don't want replicated. Just follow some basic guidelines:

- If you have only one copy, stick with database mirroring if the network bandwidth and latency allow.

- If you need more than one copy, do database mirroring for the first copy and transaction log shipping for the rest of them.
- If you need to guard against logical errors and database corruption, set up the log shipping to be delayed.

The general method of setting up database mirroring with log shipping is that the database mirror is established first, and then log shipping is established. It's not a requirement, however. The principal database in the mirror set should be the database that does the log shipping, but configuring the mirrored database for log shipping is also necessary. If the primary is offline for some reason, the mirror can continue the log shipping.

Setting up log shipping on the mirror with the exact same configuration as the principle is necessary. Because the mirror is in a restoring state, while the principle is active, the mirror won't interfere with the principle log shipping. Shipping logs to the primary or the mirror can cause failure and data inconsistency. All log shipping should go to servers not in the database mirror pair. Also, all servers should have access to where the transaction logs are backed up. Preferably this area is a network location but not physically on the same server as the principal server. Keeping the physical location separate provides for additional disaster recovery protection.

## Planning for storage redundancy

Data redundancy is a cornerstone of any disaster recovery center. The amount of redundancy planning required depends on the data's importance, the budget allowed, and the speed at which the data must be recovered. A multitude of disasters can occur, such as the following real-world examples:

- Hardware fails (that is, hard-disk failure)
- Database corruption
- Natural disaster (fire, earthquake, flood)
- Unnatural disaster (vandalism, terrorism, theft)
- Upgrade of database fails, leaving it in an unrecoverable state
- Cyber attack

Data redundancy can help prevent a disaster from turning into a catastrophe, in which the data is permanently destroyed. Several methods are available for providing storage redundancy. This chapter goes over some of the most common methods, but any storage plan requires a large amount of resources to provide true data redundancy.

## Database mirroring for data redundancy

Typically, you think of database mirroring as just providing a single copy of a database. This by itself can help prevent some disasters, such as hardware failure. A mirror in another physical location also can help recover from many of the other disasters (for example, that both locations would catch fire is highly unlikely). This provides a high level of redundancy as well

as high availability. In some cases, a single copy of the database might not be sufficient. SQL Server 2008 and above allows for a database to be mirrored at least three times. This provides for four copies of the data in potentially four different physical locations—assuming that they have enough bandwidth to support the mirroring sessions.

As mentioned in the log shipping section, database mirroring along with log shipping can provide for a virtually unlimited number of copies of the database in various stages of recovery, depending on the network's capabilities and the servers involved. If the primary SQL Server node can't keep up with the processing of the database mirroring and the replication of the transaction log data, it will start to fail. The nature of the failure depends on the particular installation, but you're always constrained by the hardware limitations.

## Database clustering for data redundancy

Using database clustering for data redundancy is also an option but doesn't provide the same level of redundancy as database mirroring (or the improved Microsoft SQL Server 2012 AlwaysOn version). Again, putting this solution into place requires much more restrictive hardware and software requirements. SQL Server supports up to 16 cluster nodes, depending on the SQL Server version and the hardware being used. The problem with clusters is that they all use common storage. This provides a single point of failure, depending on how the hardware is set up. Almost every cluster implementation will use a disk redundancy method such as RAID 0 (which is mirroring of the data disks) but you are still looking at physical location vulnerability unless the disks are being mirrored to a different location.

## Backups for data redundancy

Backups don't provide high availability but can protect from database corruption. The more often you back up the data, the more current the latest uncorrupted version of the database. The issue with database mirroring and clustering is that any corruption is also copied. Database backups are essential to providing a comprehensive database redundancy plan.

## Planning for login replication

Logins aren't replicated with the data unless the entire SQL Server instance is replicated. In the cases of database mirroring, log shipping and AlwaysOn high-availability group replication doesn't provide the SharePoint logins necessary for a successful failover. If a large number of servers are involved, maintaining the necessary logins can become a challenge. Even maintaining consistent logins between two servers can prove to be difficult. Nothing is more frustrating than performing an emergency failover and not having the logins in place to make the SharePoint instance function correctly. This isn't an issue with SQL Server clustering and AlwaysOn FCI because every node in the cluster is a copy of the whole server, not just the databases involved in replication.

You can create a script of SQL Server logins on the primary server and then run it on all secondary servers (mirror, secondary server, or member of an AlwaysOn high-availability group). Because scripting logins is a fairly involved process, you need to take great care when

creating a script because different versions of SQL Server need slightly different scripts. Also, any errors in the script can result in a non-functional failover database. To test the logins properly on the secondary database, you need to initiate a manual failover.

#### **MORE INFO** TRANSFERRING LOGINS BETWEEN SQL SERVER INSTANCES

Because a login script is fairly difficult to write, Microsoft has provided one in the Knowledge Base article at <http://support.microsoft.com/kb/918992>.



### **Thought experiment**

#### **Choosing a high-availability database option**

In this thought experiment, apply what you've learned about this objective. You can find answers to these questions in the "Answers" section at the end of this chapter.

You are told to design a high-availability solution for a group of content databases labeled as critical in the SharePoint farm. You are limited to using SQL Server 2008 R2 because of licensing and training concerns. The criterion is that if a content database fails, a failover must be available within 5 minutes. Two SQL Server instances are available for your use and are on the same network.

Develop a high-availability plan that meets these criteria.

## **Objective summary**

- Gathering requirements and understanding the technical as well as the non-technical limitations are the first steps in developing a comprehensive high-availability solution.
- Database mirroring provides redundant data as well as high availability of database servers.
- Database clustering provides high availability of servers but not redundancy of databases—at least, not without additional efforts.
- You can combine log shipping with database mirroring to provide additional copies of the data in different stages.
- Using AlwaysOn High Availability Groups—the enhanced version of SQL Server database mirroring available in SQL Server 2012—is the preferred way to provide high availability.
- The AlwaysOn Failover Cluster Instances technology is the enhanced version of SQL Server clustering available in SQL Server 2012.
- When developing a high-availability SharePoint plan, you have to consider additional database permission requirements.

## Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the “Answers” section at the end of this chapter.

1. You’ve been requested to have downtime of an hour or less during the calendar year in an environment that runs 24 hours a day, seven days a week. What RTO should you strive for?
  - A. 99.9 percent
  - B. 99.99 percent
  - C. 99.999 percent
  - D. 1 percent
2. Which of the following Active Directory DNS names does SharePoint 2013 support?
  - A. contoso.com
  - B. contoso
  - C. Both A and B
  - D. None of the above
3. Which of the following is allowed with the simple recovery model?
  - A. Transaction log backups
  - B. Log shipping
  - C. Database mirroring
  - D. Clustering
4. You have a small but critical content database that needs high availability. Which option would provide the most resource-efficient, high-availability option?
  - A. Clustering
  - B. Log shipping
  - C. Database mirroring
  - D. They are all the same
5. AlwaysOn High Availability Groups has been chosen as the high-availability option for your SharePoint farm installation. Which of the following needs additional configuration to make this possible?
  - A. One Windows Server node that’s part of the group
  - B. SQL Server
  - C. All the Windows Server instances that are part of the group
  - D. SharePoint 2013

## Objective 1.2: Plan for SharePoint high availability

Having a high-availability plan for SQL Server is essential to disaster recovery and business continuity, but planning for SharePoint high availability is just as important. If SharePoint goes down, you still aren't functional no matter how many copies of the SQL Server database you have. The goal of this objective is to ensure that an adequate SharePoint high-availability plan is in place for your business needs.

### This objective covers to the following topics:

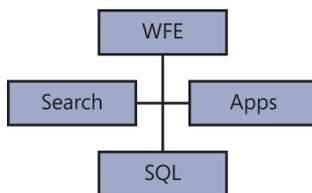
- Plan for service distribution
- Plan for service instance configuration
- Plan for physical server distribution
- Plan for network redundancy
- Plan for server load balancing
- Plan for SQL Server aliases

## Planning for service distribution

Unless you are dealing with a very small server farm, you need to distribute SharePoint services across several servers for performance as well as consider high availability. SharePoint has many different services roles that you can distribute across servers. Any SharePoint server can assume any of the roles with some configuration changes, giving SharePoint the ability to be an extremely robust and highly available system if it's configured correctly. A SharePoint server can provide these kinds of services:

- Web Front End (WFE)
- Search server
- Application server

Figure 1-6 shows how the services fit together and connect to SQL Server. Note that each service can have several servers—for example, a SharePoint 2013 farm could have five WFEs, three search servers, and two application servers. The number of different configurations depends totally on the resources available and the demands of your particular SharePoint environment.



**FIGURE 1-6** How each SharePoint Service connects with the database as well as the WFE

## Web Front End

The WFE's job is to deliver content via Microsoft Internet Information Services (IIS). A SharePoint farm can have an endless number of WFEs. For the load to be balanced across all the WFEs, additional configuration is needed. All the WFEs need to respond to the same vanity URL or IP address if they will provide the same content with the same URL. You can also leave one or more WFEs on standby in case of failure or maintenance. In any case, if no functioning WFE is available, end users will find SharePoint down. Ensuring that a WFE is available is the first step in any SharePoint high-availability plan.

## Search server

A SharePoint Search server crawls and indexes SharePoint content as well as returns query requests. These tasks don't all have to reside on the same server, and depending on how much content is being indexed, one search server might not be able to handle all the search functions. Separating the search server from the WFE provides the single biggest performance improvement (in a farm that has SQL Server on a separate server); otherwise, putting SQL Server on its own server would provide the biggest improvement. The search server is made up of a number of components that can be distributed across multiple servers to provide high availability:

- **Crawl component** crawls content (SharePoint or other content such as websites or file shares) and extracts crawled properties and metadata to send to the content processing component.
- **Content processing component** receives information from the crawl component, and then processes and sends it to the indexing component. It also interacts with the analytics processing component and is responsible for mapping crawled properties to managed properties.
- **Indexing component** receives information from the content processing component and writes it to the index. It also handles queries and sends back results to the Query processing component.
- **Query processing component** handles incoming query requests and sends them to indexing component for results. It also handles optimization of queries.
- **Analytics processing component** analyzes what users are querying on and how they interact with the results. This information is used to determine relevance, generate recommendations, and generate search reports.
- **Search administration component** manages administrative processes as well as changes to the search topology, such as adding or removing search components and servers.

You can distribute the search service components across multiple servers to provide high availability as well as improve performance. The Search Application Topology page in Central Administration shows which components are running on which servers (see Figure 1-7).

Search Application Topology						
Server Name	Admin	Crawler	Content Processing	Analytics Processing	Query Processing	Index Partition
SP13	✓	✓	✓	✓	✓	✓

**FIGURE 1-7** Each search component running on the server named SP13

The importance of each component varies in how it relates to high availability. Even if all search components were non-functional, SharePoint would still serve up content through the WFE. Of course, when a user tries to do a search, it comes back as the search service being unavailable, depending on which services are down. For example, the crawling component could be unavailable, but search would still return results; those results just become out of date until the crawling component is restored. To determine the high-availability plan for search services, you must first define what the high-availability goals are.

## Application server

A SharePoint 2013 farm can have one or more servers that perform the application server service. Any server in the SharePoint farm can be an application server, and you can have a SharePoint server with the sole functionality of a single service application. For example, if your organization relies heavily on Excel, services can benefit from having a SharePoint server dedicated to Excel Services because it's a resource-intensive application. Determining which applications will need the most resources can be difficult because doing so depends on usage. Because any SharePoint server can provide application services, a high-availability plan can target multiple servers as failover options.

## Planning for service instance configuration

A SharePoint Server can handle one or more service instances. Typically, the services moved from the WFE are the search service and application services such as Excel Services, but any service can be moved from the primary SharePoint instance. Generally speaking, search will be the most resource-intensive service to be considered for separation from the primary SharePoint server.

To configure the services available on a particular SharePoint 2013 server, follow these steps:

1. Open Central Administration.
2. Under System Settings, click Manage Services On Server.
3. The Services On Server page shows all the services running on the server. To configure a server, choose Change Server from the Server drop-down list in the upper-right corner.
4. On the Select Server page, select the desired server and click OK.
5. Back on the Services On Server page, choose Start or Stop for the appropriate services, as shown in Figure 1-8.

Service	Status	Action
App Management Service	Started	Stop
Business Data Connectivity Service	Started	Stop
Central Administration	Started	Stop
Claims to Windows Token Service	Started	Stop
Distributed Cache	Started	Stop
Document Conversions Launcher Service	Stopped	Start
Document Conversions Load Balancer Service	Stopped	Start
Lotus Notes Connector	Stopped	Start
Machine Translation Service	Started	Stop
Managed Metadata Web Service	Started	Stop
Microsoft SharePoint Foundation Incoming E-Mail	Started	Stop
Microsoft SharePoint Foundation Sandboxed Code Service	Started	Stop
Microsoft SharePoint Foundation Subscription Settings Service	Stopped	Start

**FIGURE 1-8** Part of the Services On Server page, showing some of the services that can be started or stopped

The distribution of search services and application services depends on the organization's individual business requirements. A properly prepared high-availability plan allows for redistribution of services depending on the analysis of server demand. You can always add another SharePoint server in the future to offload demand, but many physical and financial requirements need to be taken into consideration during the planning stages. For example, if the server room has no space to add another SharePoint server, that should be noted in the high-availability plan.

## Planning for physical server distribution

Physical distribution of servers occurs for two primary reasons:

- **For high availability during a disaster** If a disaster occurs in one building, physical distribution can keep services available.
- **For performance** If the SharePoint farm is servicing users in a wide, geographically dispersed region with varying bandwidth capabilities, physical server distribution can help improve performance. This is particularly true for WFEs. Putting a WFE locally in each region can greatly enhance performance from an end user's perspective.

During the planning phase, you can do the following to prepare for disaster recovery and plan for performance:

- Diagram the physical location of all servers, including those that host any virtual machines.
- During the pilot phase, record server usage (memory and CPU) so that distribution of services can be adjusted.
- Modify the architecture based on performance benchmarks and physical separation of critical services.
- After deployment, modify architecture (physical and virtual servers) as needed based on continuous measurements.

When you're planning for physical server distribution, taking a look at the underlying physical storage of data is important. Data is typically stored on a Redundant Array of Inexpensive Disks (RAID) or Storage Area Network (SAN). These two technologies represent the vast majority of disk configurations used in high-availability systems. The configuration of these technologies can greatly affect the high availability of the SharePoint farm (as well as the SQL Server implementation). To better understand how these technologies affect the SharePoint implementation, take a deeper look into how each one of them works.

## SAN

A Storage Area Network is a set of disks and tapes that can be viewed as a single entity. Accessed at the file block level, a SAN appears as a locally attached drive rather than as a remote storage unit. The disks and tapes that are part of the SAN can be in physically distant locations, as opposed to traditional network attached storage (NAS), which usually must be located within a few meters of each other. Knowing the terminology associated with a SAN helps with potential exam questions as well as when discussions about physical storage planning occur. Several SAN-related terminology items are as follows:

- **SAN Fabric** refers to the hardware that connects the servers to the storage disks and tapes.
- **LUN** The logical unit number represents a logical group of disks and/or tapes. A LUN can be used as a single drive or as multiple drives.
- **Fiber Channel** This high-speed networking terminology is often used by SAN implementations.

Implementing a SAN is vendor-specific, so exact specifications are beyond the scope of this book. Generally, a SAN is usually much more expensive to implement but gives the best performance for physically separated storage. In fact, SAN solutions have been successfully implemented on storage devices that are many miles apart. Successful implementations over such distances requires cooperation with a network provider that can provide sufficient network bandwidth and the technologies for transport required (such as Fiber Channel) to make SAN a viable option.

## RAID

Redundant Array of Inexpensive Disks (RAID) is a storage system that can be part of an individual server or a standalone set of disks that can be accessed by a number of different servers. You can set up RAID in a number of different varieties (referred to as *levels*), but typically only a few are relevant to storage planning. The RAID levels are as follows:

- **RAID 0** Referred to as *disk striping* because it writes blocks of data across multiple disks. This provides performance improvements by allowing data to be written to multiple disks at the same time but doesn't provide any data redundancy.
- **RAID 1** Often referred to as *disk mirroring* and requires at least two disks. Every bit of data written to one disk is written to the other. The disks don't have to reside in the physical location but usually must be close for latency reasons.

- **RAID 5** Known as *data striping with parity*. Data is written over a minimum of three disks with parity information that can help rebuild information if a disk is lost. A RAID 5 disk array can lose one disk and still keep performing with no data loss (at reduced functionality) until a disk can be replaced.
- **RAID 10** Sometimes called RAID 1+0. This level provides the benefits of both RAID 1 and RAID 0. It's composed of two sets of disks, each in the RAID 0 configuration. This RAID configuration is one of the most commonly implemented today.

To have any sort of physical separation of data, either RAID 1 or RAID 10 must be used. The ability to physically separate the disks in these two RAID options means that the drives can be placed in different racks, maybe even in different rooms. RAID 5 provides some protection in that a single disk can be lost and no data will be lost.



---

**EXAM TIP**

Familiarity with RAID 0, RAID 1, RAID 5, and RAID 10 is expected, as well as why you would choose one type of RAID storage over the others.

---

## Planning for network redundancy

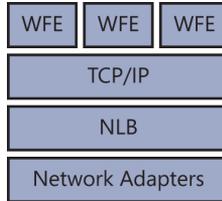
Network redundancy means that communication pathways exist, even if a particular link goes down. Because many real-world examples are known of when links go down, planning for them can help keep the SharePoint farm up and running. The following are examples of issues that arise where network redundancy is necessary to maintain service:

- Ethernet port has hardware failure.
- DNS server crashes (hard drive, operating system, software, and so on).
- Router has a hardware failure.
- Switch has a hardware failure.
- IP address is accidentally reused.
- Denial-of-service (DOS) attack is made against an IP address.

The exam doesn't expect you to be a certified network engineer, but you should be aware of these issues and how to plan for them. You should also be prepared to offer suggestions on how network redundancy can help to improve the overall uptime of the SharePoint farm. Many technologies are available to help improve network redundancy, but most of them are beyond the scope of this book. However, one that falls within the realm of this exam is Network Load Balancing (NLB).

NLB is most often discussed when WFEs are discussed. With this method, a single IP address or vanity URL can be used to connect to a set of WFEs (a user connects to only one server at a time, no matter which server). Noting that Central Administration can also be load balanced is important, although it's typically installed on only one server (best practices are that it's running on more than one SharePoint server even if it's not load balanced). NLB is configured outside of SharePoint on each server in the NLB group.

Because NLB operates at the network driver level, the TCP/IP stack isn't aware of it. Figure 1-9 shows how the communication layers are stacked to help get a better understanding of how the layers communicate with each other.



**FIGURE 1-9** Communication stack showing how NLB sits below TCP/IP, indicating that the WFE isn't aware of it

Figure 1-9 clearly shows that the WFEs aren't aware of the network load balancing occurring because items in the stack are aware only of layers immediately below or immediately above.

Now that you have a graphical understanding of how NLB works in the communication stack, you can focus on how it's implemented from a practical standpoint. The following steps, for Windows Server 2008 R2, show how to configure NLB:

1. Install Network Load Balancing on each server (it's an optional component of Windows Server) by going to Start | Administrative Tools | Server Manager.
2. In the Features Summary section of the Server Manager main window, choose Add Features.
3. In the Add Features Wizard, select the Network Load Balancing check box and click Next.
4. On the Confirm Installation Selections page, click Install.
5. After NLB is installed, you can configure it by opening the Network Load Balancing Manager (Start | Administrative Tools | Network Load Balancing Manager).
6. Right-click Network Load Balancing Clusters and select New Cluster (see Figure 1-10) to launch the New Cluster Wizard.



**FIGURE 1-10** Selecting New Cluster in the Network Load Balancing Manager

7. On the Connect page, enter the host name of the first server to be part of the cluster (the name of the server being configured) and click Connect to list all available network addresses. (If you have only one network adapter and only one IP address is configured for that adapter, only one choice will be listed.)
8. Select the network IP address to add and click Next.
9. On the Hosts Parameters page, select a Priority (unique host identifier). The host with the lowest priority number takes precedence when no port rules are in effect (start with number 1 and work downward unless a specific server is set as the default).
10. If you need to add IP addresses, click the Add button. Otherwise, click Next.
11. On the Cluster IP Addresses page, enter the shared IP address (or addresses). This IP address is used by end users to connect to SharePoint. Then click Next.
12. On the Cluster Parameters page, select Unicast (if your servers have only one NIC card in them, choose Multicast). This way, the cluster's Media Access Control (MAC) address is used for each network adapter, rather than the individual MAC addresses. Click Next.
13. On the Port Rules page, click Edit and add the appropriate port rules, if necessary. Click Finish when done.
14. To add more servers to the cluster, go back to the Network Load Balancing Manager, right-click Network Load Balancing Cluster, and choose Connect To Existing. Fill in the details as was done for the first server and repeat for each server that is to be part of the cluster.

#### **IMPORTANT NLB AND DHCP**

Network Load Balancing can't use Dynamic Host Configuration Protocol (DHCP). IP addresses must be static.

NLB must be installed on every server that's part of the NLB cluster before it's added to via the Network Load Balancing Manager tool. After NLB is configured, if an individual server goes down, the load is automatically redistributed among the remaining servers. NLB determines whether a server is up via a "heartbeat" that each server sends out. If the Network Load Balance Manager doesn't receive a "heartbeat" for 5 seconds (by default), it assumes that the server is unresponsive and routes requests to the other servers in the NLB cluster. This is referred to as a *convergence*, and the server with the lowest priority number takes over as the default. If the server that was down (unresponsive heartbeat) comes back online during the convergence, it becomes part of the group. If it doesn't, it must be added back in via the NLB Manager tool. The convergence takes only a few seconds, so the total outage (assuming that a user was accessing the downed server) should be about 7–10 seconds. For other users, the cluster is down for only a few seconds.

You can add or remove a server from the NLB cluster at any time. This greatly improves uptime of the SharePoint server farm and allows for maintenance of servers with minimal downtime. NLB provides a number of additional features in Windows Server 2008 R2, including the following:

- An NLB cluster can have up to 32 servers.
- You can add hosts without bringing the cluster down.
- You can bring a host down for maintenance without downtime.
- A downed host can recover within 10 seconds.
- You can define port rules for individual websites, giving great flexibility on how the load is distributed.
- You can perform remote administration from anywhere on the network.
- No additional hardware is required.
- NLB allows for multiple IP addresses.

#### **MORE INFO** NLB OVERVIEW

To learn about more NLB features, see <http://technet.microsoft.com/en-us/library/cc725691>.

## Planning for SQL Server load balancing

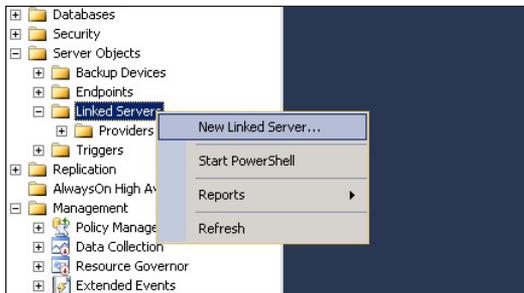
Using SQL Server load balancing with Active/Active clustering isn't supported. *Active/Active* refers to two separate SQL Server instances using each other as failovers. This is different in that a single SQL Server instance uses one or more cluster nodes as failovers in case the primary node fails. Database mirroring or AlwaysOn high availability still relies on a single SQL Server instance to receive transactions and then push them out to secondary servers so that both methods rely on a single SQL Server instance to interact with for transactions. Log shipping also has a single server as a bottleneck for inputs because the secondary servers only receive content and can't be used for write transactions unless a failover is initiated. Even if clustering is set up, it still relies on a single server to receive the transactions and then replicate them to the other servers.

Planning for SQL Server load balancing is an advanced topic and requires considerable planning and resources. Generally, giving the primary SQL Server node more memory and network bandwidth resources gives more of a benefit than trying to load balance several SQL Server instances. One option for balancing some of the load is to use linked servers. This allows for a database that resides on one SQL Server instance to appear as it is on another SQL Server instance. The following section describes this option in more detail.

## Linked servers

Linked servers are most useful in situations where the data can be separated into distinct databases. This is perfect for SharePoint, which can put sites into separate databases. A web application can have any number of databases, which, by using linked servers, can be distributed across multiple SQL Server nodes. Most resources involved in a database transaction are related to writing and reading data from the disk. Having multiple SQL Server instances allows for the distribution of this time-expensive operation. Connecting a linked server is fairly straightforward, but looking at the steps involved is worth the effort for a better understanding of how linked servers work:

1. Open Microsoft SQL Server Management Studio on the database server that SharePoint connects to.
2. In the Object Explorer window, expand the Server Objects folder.
3. Right-click Linked Servers and select New Linked Server as shown in Figure 1-11.



**FIGURE 1-11** Creating a new linked server in SQL Server Management Studio

4. In the New Linked Server dialog box, select SQL Server under Server Type.
5. In the Linked Server text box, enter the name of the server to be linked.
6. Back in the Select A Page section on the right side of the screen, select Security.
7. In the Local Server Login To Remote Server Login Mappings section, add the login for the account under which the web application runs.
8. Click OK.

The linked server should now be connected and available for use. Limit linked servers to content databases that receive relatively little use, such as archived records. Queries against linked servers require distributed transactions, which are relatively expensive in memory and network bandwidth. The balance between disk operations and storage requirements need to be weighed against the distributed transaction costs.

### **MORE INFO CREATE LINKED SERVERS**

See <http://msdn.microsoft.com/en-us/library/ff772782.aspx> for more information on creating linked servers.

## Federated database servers

For SharePoint implementations that require high availability of SQL Server databases, federated database servers represent a viable solution for data that can benefit from horizontal partitioning. This highly advanced database subject requires significant planning and in-depth knowledge of SQL Server. Each SQL Server node in a federation group is managed independently, but all nodes cooperate to provide data as though from a single server.

### **MORE INFO FEDERATED DATABASE SERVERS**

Although the federated databases present themselves as a single server instance, they have internal differences. See [http://technet.microsoft.com/en-us/library/ms190381\(v=sql.105\).aspx](http://technet.microsoft.com/en-us/library/ms190381(v=sql.105).aspx) for more information.

## Planning for SQL Server aliases

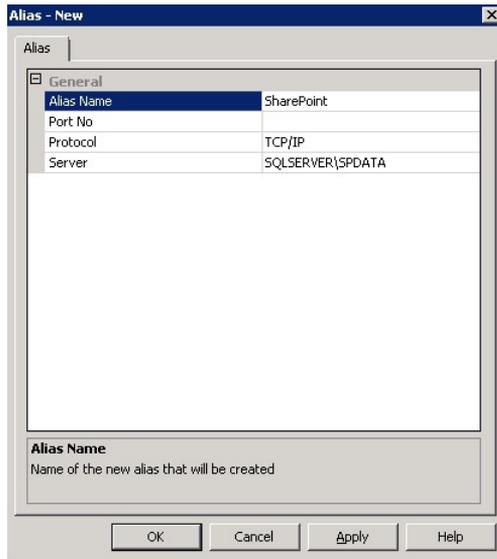
A SQL Server alias allows for SharePoint to switch from one SQL Server instance to another should a disaster occur. SharePoint does this by connecting to the alias rather than directly to the SQL Server instance. If that SQL Server instance goes down, SharePoint can be pointed to a mirrored instance and be back up by just changing the details of the alias.

### **NOTE LOCATION OF SQL SERVER CLIENT CONNECTIVITY COMPONENTS**

SQL Server aliasing is done on the client computer (SharePoint Server node), not on SQL Server itself. Each SharePoint server needs the SQL Server client connectivity components installed on it.

You can create a SQL Server alias in just a few steps. The following steps create an alias on a SharePoint 2013 server with Microsoft SQL Server 2012 client connectivity components installed on it:

1. From the Start menu of the SQL Server instance, navigate to Microsoft SQL Server 2012 | Configuration Tools | SQL Server Configuration Manager.
2. Expand SQL Native Client 11.0 Configuration.
3. Right-click Aliases and select New Alias.
4. Fill in the values for Alias Name and Server, as shown in Figure 1-12. Leave the Protocol as TCP/IP and the Port No as blank.



**FIGURE 1-12** Alias dialog box from the SQL Server Configuration Manager

5. Click OK.
6. Repeat the preceding steps for SQL Native Client 11.0 Configuration (32bit). This is for *thunking* (converting 32-bit instructions to 64-bit instructions and vice versa) when communicating with SQL Server.
7. Repeat steps 1-6 for each SharePoint server.

After you configure each the SharePoint 2013 server with the appropriate alias, you can install the SharePoint 2013 farm. Where it asks for the database server, enter the alias rather than the actual server name. If you've already installed your SharePoint farm, you need to run the SharePoint Products Configuration Wizard and choose *Disconnect From This Server Farm* on the *Modify Server Farm Settings* page. After you disconnect, you can rerun the SharePoint Products Configuration Wizard and connect to an existing SharePoint farm by using the alias instead. Clicking *Retrieve Database Names* should bring back the existing configuration database. You also need to supply the farm passphrase and a port number (or use the one that's generated). Finally, all respective SharePoint services need to be restarted.

**IMPORTANT REMOVING THE USER PROFILE SYNCHRONIZATION SERVICE**

The Forefront Identity Manager (FIM) Services on the SharePoint server responsible for the User Profile Synchronization becomes disabled if you detach it from the existing SharePoint farm. You need to set the start mode for the FIM services manually from *Disabled* to *Automatic*.

If the SQL Server Configuration Management client tools aren't available on the SharePoint server, you can create an alias using the SQL Server Client Network Utility as follows:

1. Run the following application:  
`c:\windows\system32\clconfig.exe`
2. On the Alias tab, click Add.
3. Select TCP/IP from the Network Libraries section.
4. Type the alias in the Server Alias text box.
5. Enter the Server Name in the Connection Parameters section.
6. Leave the Dynamically Determine Port check box selected and click OK.
7. Run the 64-bit version (`c:\windows\syswow64\clconfig.exe`) and repeat steps 1–6.
8. Repeat all these steps for each SharePoint server.

When these steps are complete, the aliases are complete and a farm can be installed or reconnected as described earlier.

If you're using database mirroring, AlwaysOn High Availability Groups, or log shipping, you should install SharePoint using a SQL Server alias to enable switching to a new SQL Server instance. The only high-availability options that don't benefit from an alias are database clustering and AlwaysOn Failover Cluster Instances. This is because if one node goes down in a cluster, another node takes over, but the SQL Server instance name stays the same. This doesn't prohibit an alias from being used in a clustering environment; in fact, it's encouraged in case the cluster itself fails and a backup needs to be restored to a new SQL Server while the cluster is restored.



### ***Thought experiment***

#### **Balancing a network load across multiple WFEs with high-availability databases**

In this thought experiment, apply what you've learned about this objective. You can find answers to these questions in the "Answers" section at the end of this chapter.

You have been told to create a SharePoint farm that should have multiple Web Front Ends and connect to a single mirrored SQL Server instance. You want to develop a solution that distributes the load across the WFEs and allows for switching over to the mirrored database if the primary SQL Server node fails or is down for maintenance. With the WFEs, the primary goal is to have high availability of front-end servers; the secondary goal is to have the processing load distributed across the servers, giving end users more responsive interactions with the SharePoint farm.

What two technologies would you implement to make this happen?

## Objective summary

- SQL Server aliases provide a quick way to point SharePoint to a new database server.
- SharePoint services can be distributed across any number of SharePoint servers.
- You can use Network Load Balancing to distribute the network load over several Web Front Ends to provide performance improvements and high availability.
- Linked servers provide a way to distribute SQL Server data for databases with low usage.
- SharePoint services fall under three main categories: Web Front Ends, Search, and Application server.

## Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the “Answers” section at the end of this chapter.

1. You want to create SQL Server aliases as part of your high-availability plan. On which of the following will you need to configure this?
  - A. All the SQL Server instances that are part of the SharePoint farm
  - B. SharePoint servers that access SQL Server data
  - C. Both A and B
  - D. Neither A nor B
2. A SharePoint server can be which of the following roles?
  - A. Web Front End
  - B. Search server
  - C. Application
  - D. Any or all of the above
3. You want to distribute the network load on your SharePoint Web Front Ends (WFEs). On which of the following will you need to configure this?
  - A. Central Administration
  - B. NLB on each WFE
  - C. In the IIS Manager
  - D. All of the above
4. Where does Network Load Balancing (NLB) sit in the communication stack?
  - A. Right below the WFE layer
  - B. Below the network adapter layer
  - C. Right below the TCP/IP layer
  - D. Above the WFE layer

5. Which of the following would provide data redundancy at the hardware level?
- A. RAID 0
  - B. RAID 1
  - C. RAID 5
  - D. RAID 10

## Objective 1.3: Plan backup and restore

---

Backup and restore are essential to any disaster recovery plan. You need to consider many factors when developing a comprehensive backup and restore plan. RTO and RPO are the drivers that should determine the methods and frequency of backups, but in a SharePoint farm, making sure everything required for a full recovery is properly backed up is also essential. This objective goes over the essential components of the SharePoint installation, backup and recovery methods, and the frequency that each component should be backed up.

### This objective covers the following topics:

- Establish a SharePoint backup schedule
- Establish a SQL Server backup schedule
- Plan a non-production farm content refresh
- Plan for a farm configuration recovery
- Plan for service application recovery
- Plan for content recovery

## Establishing a SharePoint backup schedule

SharePoint needs to be backed up independently of SQL Server. You can back up SharePoint in a couple different ways, but to understand what a SharePoint backup versus a SQL Server backup means, you need to look at the SharePoint components that need to be backed up:

- Farm configuration
- IIS Settings
- Files in the web application
- Farm solutions
- The entire SharePoint hive
- Server topology (which servers are performing which functions)
- IP Settings
- Certificates
- Any Dynamic Link Libraries (DLLs) stored in the Global Assembly Cache (GAC)

The farm configuration is stored in the central administration database but should also be backed up independently of the normal SQL Server backups. This provides several benefits, including marking a farm configuration as a “golden” copy so that if corruption happens to the configuration, you’re not trying to figure out when the last uncorrupted copy was backed up. Also, the backup can be kept local, which makes restoration easier for the SharePoint administrator. In many organizations, the person responsible for the SharePoint farm doesn’t have the rights (and perhaps training required) to restore a SQL Server database. A SharePoint farm backup allows that person to restore a farm configuration independently.

When deciding on a schedule, the recovery model is important to consider. By default, SharePoint content databases have the full recovery model, which provides a greater level of recovery but takes up more hard drive space. To better understand the recovery model, look at the options in Table 1-1.

**TABLE 1-1** Available recovery model options

Recovery model	Description	Recovery options
Simple	Automatically reclaims log space so that log space requirements are minimal.	Changes since the most recent backup aren’t protected. Only the most recent backup can be recovered.
Full	Requires log backups. Log files can be used to recover to data up until failure.	You can recover to a point in time assuming that backups are complete.
Bulk Logged	Between Full and Simple. Requires log backups. Reduces log space by using minimal logging requirements.	You can use this to recover to the end of any backup. Point-in-time recovery isn’t supported.

Most databases use either the Simple or Full recovery model. If you are using the Simple recovery model, backing up the data more frequently is important so that the amount of data lost is minimized. The main issue with the Simple recovery model is that it’s not available if high-availability options such as database mirroring or AlwaysOn High Availability Groups are being used.

## Using Central Administration backup

The primary tool used to back up a SharePoint 2103 farm is the Backup Tool available in Central Administration. You can use this easy-to-use graphical interface tool to back up the farm configuration as well as well as content databases. This tool has several types of backup types available:

- The complete farm with data and configuration
- Differential
- Configuration only
- Component level

#### **NOTE USING DIFFERENTIAL BACKUPS**

Differential backups are only for content databases. To restore, you need the full backup as well as the differential backup. Also, if you have large content databases (more than 100 GB each), the backup process could take a long time.

Central Administration backup provides a simple interface, but you need to consider some exceptions when using it. Although it's a great tool to use to back up farm configuration information, it should be used with caution as part of your content database backup strategy. Central Administration Backup can write to only a file location, not to tape, so enough file space must be available to handle the backup. Because it doesn't automatically get rid of old backups, they can add up quickly, depending on the type of backup. Finally, Central Administration has no built-in way to schedule backups. The only reason to use Central Administration Backup for content databases would be for archiving purposes. However, it's a valid backup strategy for having a locally available backup of the configuration data.

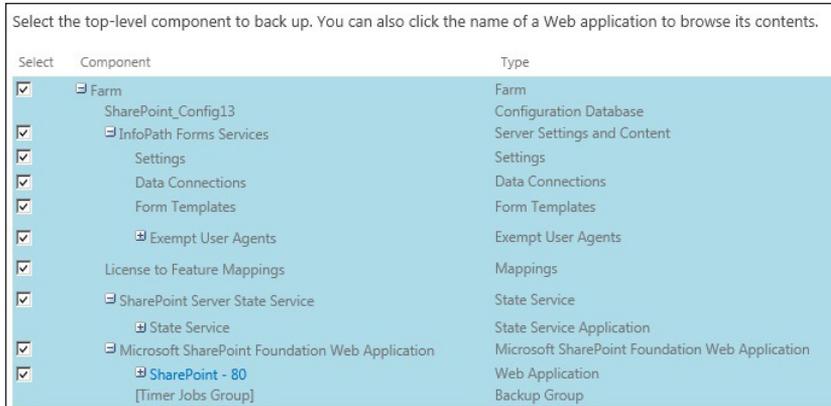
Configuration-only backups will be the primary function of the Central Administration Backup tool because no other built-in tool provides backups of configuration data with such completeness. This doesn't mean that it backs up all the configuration information necessary to restore a SharePoint server. Several important items aren't backed up by the Backup tool:

- **IIS settings not set by SharePoint** The IIS Microsoft Management Console (MMC) snap-in provides a backup mechanism that you can use to back up IIS settings. Although some overlap will occur, having this backup available is important in case settings need to be restored.
- **SSL Certificates** These need to be saved as individual files.
- **Anything in the SharePoint hive** The entire SharePoint hive needs to be saved (at least the templates folder) periodically, after you make major changes such as adding features or change any existing files, such as *docicons.xml*. The entire *Inetpubs* directory also can be saved. This takes off any files related to web application folders. You should document which services are running on each server and keep this information in a location not on your SharePoint server farm.
- **Files in the Global Assembly Cache** Files put into the GAC (DLLs in farm-level solutions) should be saved in a directory under a subfile that gives the version. All these files should be stored in a location that's not part of the SharePoint farm and in a different physical location from the database servers and the SharePoint farms. If a farm needs to be restored, redeploying the farm solutions should restore the DLLs.
- **Files in the web application folder** These should be manually copied to a safe location and documented.
- **Which servers are running which services** This information should be noted manually in a configuration document.

These settings should be backed up separately for each server in the SharePoint farm.

Backing up just the SharePoint configuration is straightforward with the Central Administration user interface. Follow these steps:

1. Navigate to Central Administration using a farm administrator account.
2. Click Backup And Restore.
3. Under Farm Backup And Restore, click Perform A Backup.
4. Select the very first check box that has the word Farm (see Figure 1-13).



**FIGURE 1-13** Selecting all available components in a farm backup

5. Click Next.
6. For Backup Type, leave Full selected.
7. In the Back Up Only Configuration Settings section, select the Back Up Only Configuration Settings option.
8. For the Backup File Location section, enter a UNC path (for example, `\\server\back-up`) in the Backup Location box.
9. Double-check your settings and click Start Backup to begin the backup process. The next page displays the status of the progress as soon as the timer job is started. When everything in the Progress column says that it's completed, the backup is done.

#### **IMPORTANT** ALLOWING SPACE FOR BACKUPS

When the Central Administration user interface is used, it calculates the space required to finish the backups and throws an error if not enough space is available. It tends to over-calculate because of the backup size of the truncated log files if you are using the simple recovery method.

SharePoint creates a folder in the backup location (such as *spbr0000*) each time a backup is performed. When you do a restore, you need to know which folder to choose to restore the proper backup. In the specified backup folder is a *spbrotoc.xml* file that specifies the backups available and when they occurred. Within each backup folder are two files that explain the contents of the folder—*spbackup.log* and *spbackup.xml*—in case you need to examine the contents of what was backed up.

You can also use PowerShell to perform SharePoint backups in a similar way as Central Administration because it provides the equivalent functionality. It's useful if you will be repeating the same type of backup or if you plan to automate any backup function. Being familiar with the syntax of the PowerShell backup commands is important:

```
Backup-SPFarm -BackupMethod <Full|Differential> -Directory <UNC Path> [-Configuration Only] [-Force] [-Item <named component>]
```

If you save scripts that you use regularly, you can run them whenever needed rather than have to rewrite them every single time.

SharePoint also provides an object model that programmers can use to perform backups. This will be more involved than would benefit most organizations, but it's available for more complex farm implementations as well as third-party organizations that want to develop tools to provide comprehensive backup solutions.

## Establishing a SharePoint backup schedule

A SharePoint backup schedule varies from a SQL server backup schedule. It should be driven by change, not by time. Whenever a significant environmental change occurs, you should perform a backup. Sometimes this might happen multiple times in a day, and sometimes weeks or months might pass without a need to back up SharePoint. This means that a manual process must be involved to start the backup process.

The following are some examples of when a full SharePoint backup should be performed:

- Right after installing and/or upgrading SharePoint and all the configurations are completed
- Before and after the addition or removal of a server to the farm
- Before and after significant services are added or removed (for example, adding or removing Project Server)
- Before and after installing any cumulative update (CU) or service pack
- Before attempting any upgrade

These backups can't be scheduled and should be performed by a SharePoint administrator. A snapshot of the server provides additional protection if you are using a virtual machine (VM); a backup of the entire server via disk or tape if you aren't using a VM provides additional protection and options for recovery. These backups do require significantly more resources than just backing up the pertinent information required for SharePoint.

A scheduled backup of the farm is also a good idea and should be part of high-availability planning. Central Administration doesn't allow for scheduled backups, so you need to use another method. For example, you can create a PowerShell script and then set up a Windows timer job to automate it.

## Establishing a SQL Server backup schedule

SQL Server is where SharePoint stores the bulk of the content that needs to be backed up. If you can recover a content database, you can build a new SharePoint server and attach to it, thus recovering the data within. Determining a SQL Server backup depends on a number of factors, including some that will be out of your hands but must still be considered:

- Technologies available, such as Microsoft Data Protection Manager (DPM)
- Size of the content databases to be backed up
- Relative importance of the data
- How long to keep backups
- Resources available to do backups
- Timing of backups
- Time to restore backups

A typical scenario would be doing a full backup once a week and doing a differential every day during the week. This scenario would generally be used in an environment that backs up the databases to tape or to a shared storage area such as a SAN. (Tape has the benefit of being cheaper, but a SAN provides a quicker recovery time.) As part of the backup schedule, the service level agreement (SLA) should be determined. Using differentials saves money but requires more time to restore because the most recent full backup must be restored before you can restore a differential.

If a VM is being used for SQL Server, backups of the individual content databases should still be used because restoring a single content database from a VM is impossible. The backup VM could be used to bring up a temporary instance, and then a single content database could be saved from it and restored to the production database.

## Planning a non-production environment content refresh

An up-to-date reproduction of your production environment is a great asset in development and testing. Every cumulative update and service pack should be tested in a non-production environment that matches as closely as possible the production environment. In many situations, however, this is prohibitively expensive. Not every single SharePoint server in the production farm needs to be reproduced in the non-production environment, but if at all possible, the reproduction should include all the production environment content. This is to allow for an environment that's as close to the real thing as possible without having a perfect replica.

If you are dealing with VMs, the process is much easier than if you are dealing with physical servers. You can simply take a snapshot of the servers involved and copy them over into the non-production environment replacing any out of date servers.

A non-production environment consists of at least the following items:

- Domain controller
- At least one SharePoint server
- SQL Server database(s)

With these three items, you can create a non-production environment with a high degree of certainty that your solutions, cumulative updates, or service packs behave similarly to the production environment.

#### **IMPORTANT MAINTAINING A CONSISTENT ENVIRONMENT**

**All servers should have the exact same operating systems that exist in the production environment. This means the same versions, updates, service packs, and so on.**

Keeping content refreshed is tricky if development is being done in the non-production environment (as is typical). Often, development is being stored in the content database. Items such as all client-side scripts, changes to master pages, workflows created out of the box (OTB) or with SharePoint Designer, changes to .aspx pages, and so on are all stored in the content database. A simple database restore would wipe all this work out unless everything had been re-created in the production environment before the content refresh. In these situations, determining how and when to refresh the content is up to the individual organization.

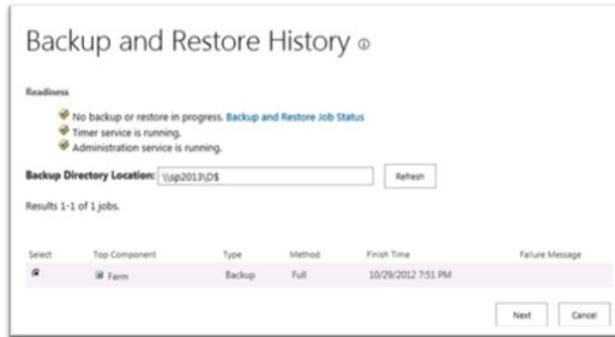
If you want a non-production environment with almost current (depending on the frequency of the updates but potentially less than an hour old) data all the time, you can use Content Deployment. SharePoint provides Content Deployment OTB. This can be used to keep a non-production environment current with scheduled data. Content Deployment can be found in Central Administration or can be run via PowerShell or the SharePoint object model. For the purposes of the exam, this book focuses on the Central Administration method. Later, this book explores in detail the content deployment methods and configuration.

## **Planning for farm configuration recovery**

Earlier, this chapter went over how to back up a farm configuration. Assuming that you backed up your configuration at that time, you can now recover it. Some settings have to be recovered manually, but you can start with recovering the farm configuration by using the built-in configuration restoration method in Central Administration. Because the Central Administration restore method uses a graphical UI that presents a visual representation of what's being restored, you will less likely restore something unintended.

A configuration backup must be done from a full backup and not a differential. (A differential backup is applicable only to content databases.) To better understand the process, you need to look at the steps involved:

1. In Central Administration, click Backup And Restore and then select Restore From A Backup.
2. The most recent backup path should be listed with the backup sets found. If the backup you need isn't in that location, enter the Backup Directory Location and click Refresh.
3. Select the backup that you want to restore, as shown in Figure 1-14, and then click Next.



**FIGURE 1-14** Selecting a specific backup to restore

4. SharePoint displays a list of all the components available for restore. Select the check boxes of the desired content to restore. If you are doing a full configuration restore from a configuration-only backup, choose Farm to select everything below it and then click Next.
5. The next page that appears depends on what you chose in step 4. To replace the current configuration, choose Same Configuration. A warning message appears, saying that all components will be overwritten. Click OK to continue.
6. Click the Start Restore button to begin the restore. A timer job is created, and you are redirected to the Backup And Restore Status page.
7. Click Refresh to view the progress. When everything in the progress column says *Completed*, the restore is done.

#### **IMPORTANT RESTORATION CONSIDERATIONS**

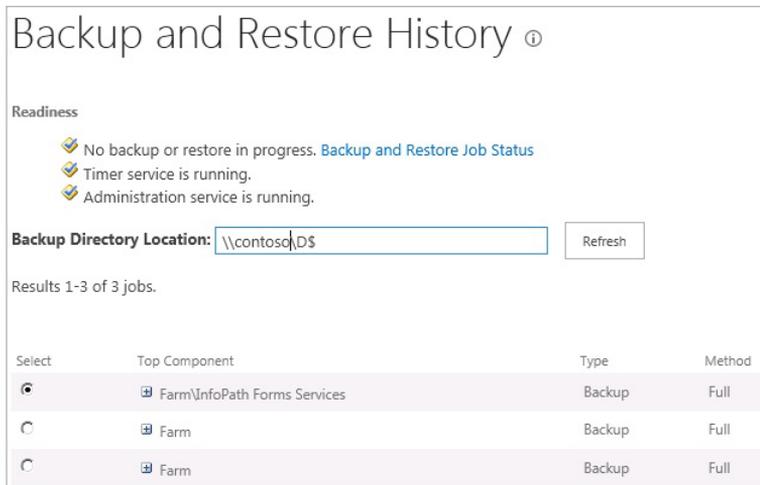
Double-check all your settings before clicking Start Restore. The restore could take a while, and during this time the SharePoint farm is unavailable. Don't try to restart or stop the restoration process after it starts—it could leave your farm in an unstable condition.

## Planning for service application recovery

You can use Central Administration to back up components. If you need to be able to recover an application, you should have backed them up as described earlier in this objective. If you have done a full farm backup, you can simply choose an individual component to restore (such as Excel Services) from the entire backup. Ensuring that you have backed-up settings since any major changes to a service application were made is important. If you have a current backup, the restoration process is fairly straightforward via Central Administration. Even if you back up service applications, you should store in a location that's not in SharePoint a documented list of which services are on which servers.

Central Administration provides individual component backup. For example, if you want to back up only the Machine Translation services, you could use the OTB backup-and-restore services to back up just this one component so that it can be recovered quickly. For the purposes of illustration, the following steps restore a single service application, a previously backed-up instance of the InfoPath Services:

1. Navigate to Central Administration and click the Backup And Restore link.
2. Click Restore From A Backup under the Farm Backup And Restore section.
3. Select Farm\InfoPath Form Services from the list of backup jobs, as shown in Figure 1-15, and then click Next.



**FIGURE 1-15** Choosing the InfoPath Form Services from the jobs in Backup and Restore History

4. Select InfoPath Forms Services on the Select Component to Restore page, and then click Next.
5. On the Select Restore Options page, select Restore Content And Configuration Settings under Data To Restore and choose Same Configuration under Type Of Restore.

6. Click Start Restore.
7. When the Backup And Restore Status page appears, click Refresh to check the status of the restore. When the restore is complete, everything under Status should show as completed.

## Planning for content recovery

Several viable options are available for recovering content. Each one has its pros and cons, but they can all provide a full recovery of all the content from a content database. This section covers the following options:

- Recovery via *Restore-SPSite*
- Content database recovery via SQL Server
- Recovery via Central Administration

Backing up a content database with SQL Server or backup-and-restore products such as Data Protection Manager (DPM) by Microsoft is a standard disaster-recovery protocol. When a restore is required, you can restore over the existing content database, in which case everything that has been changed since the backup is lost, or you can restore to a new database, in which case you can look at the backed-up data and choose what you want to recover. This feature is extremely useful if you just want to recover certain parts of the content database, such as a corrupted site or list.

### **NOTE USING STSADM**

*STSADM* is still available and can work to export and/or import websites. It's no longer the preferred method, though, because PowerShell is available.

Recovering a specific list or document library is best done using additional software because a list or document library template is limited to 25 MB by default. DPM enables you to back up and restore individual libraries and lists. If a library or list is below the limit, a list or library can be saved as a template (with content) to a temporary location by going through the galleries section of the site collection settings page. When it's saved to the temporary location, you can upload it to the production SharePoint farm and restore it in the proper location. If a library or list exceeds the SharePoint specified limit, another product needs to be used, such as Microsoft DPM.

### **MORE INFO USING MICROSOFT DPM**

See <http://technet.microsoft.com/en-us/library/cc424951.aspx> for more information on how to use Microsoft DPM to back up and recover lists and document libraries.

## Recovering a database using Restore-SPSite

If a content database has been restored to a temporary location, you can use the PowerShell commands *Backup-SPSite* and *Restore-SPSite* to recover an entire site collection. The *Restore-SPSite* syntax is as follows:

```
Restore-SPSite [-Identity] <String> -Path <String> [-AssignmentCollection  
<SPAssignmentCollection>] [-Confirm [<SwitchParameter>]] [-ContentDatabase  
<SPContentDatabasePipeBind>] [-Force <SwitchParameter>] [-GradualDelete  
<SwitchParameter>] [-HostHeaderWebApplication <String>] [-WhatIf [<SwitchParameter>]]
```

### **MORE INFO** RESTORE-SPSITE

See <http://technet.microsoft.com/en-us/library/ff607788.aspx> for more information on the *Restore-SPSite* PowerShell command.

If the backed-up site collection is being restored over an existing site collection, you must use the *-Force* option to overwrite the site collection and everything that has changed since the backup. Use caution when restoring large site collections; doing so could cause the SharePoint farm to be nonresponsive or to respond slowly. Also, the operation must have enough memory to finish. If not enough memory exists, the restore fails.

Depending on the size of the site collection, the restore operation could take hours through the PowerShell command. Thorough testing should be done to establish expectations and determine whether the SharePoint farm's capabilities meet the business requirements.

## Recovering a database via SQL Server

Recovering a database with SQL Server involves two options: You can restore a whole content database, replacing an existing one, or you can restore a content database to a new name. Restoring an existing content database with a different name enables you to extract certain pieces of data as well as review the data against the current set of data. If you try to attach the content database to the same SharePoint farm that's using the current content database, SharePoint won't allow two databases with the same ID. Therefore, when you mount the database, you need to assign it a new ID. This can be done with the following PowerShell command:

```
Mount-SPContentDatabase [-Name] <String> [-WebApplication] <SPWebApplicationPipeBind>  
[-AssignNewDatabaseId]
```

Example:

```
Mount-SPContentDatabase "Contoso Sales" -WebApplication http://contoso/sales  
-AssignNewDatabaseId
```

#### **MORE INFO MOUNT-SPCONTENTDATABASE**

See <http://technet.microsoft.com/en-us/library/ff607581.aspx> for more information on the PowerShell command *Mount-SPContentDatabase*.

## Recovering a database using Central Administration

Even with a different database ID, you shouldn't attach the database to the same web application as the one that contains the existing content database because the URLs will be the same for both databases. Create a new temporary web application to attach the database (the preferable solution) or attach it to a different existing web application that doesn't have the same tree structure.

You can use PowerShell to restore content from an unattached database backup. This way, you can avoid attaching with a new ID and still allow for data restoration. The PowerShell command for this is as follows (after the database is restored to SQL Server):

```
Get-SPContentDatabase -ConnectAsUnattachedDatabase -DatabaseName <DatabaseName>  
-DatabaseServer <DatabaseServer>
```

After the PowerShell command finishes, you can recover the data using Central Administration as follows:

1. Log onto SharePoint 2013 Central Administration with a farm administrator account that also has db\_owner rights on the SQL Server instance where the restored database resides.
2. Click the Backup And Restore link.
3. On the Backup And Restore page, click Recover Data From An Unattached Content Database in the Granular Backup section.
4. On the Unattached Content Database Data Recovery page, enter the name of the unattached database in the Database Name box and the name of the SQL Server instance it's on in the Database Server box.
5. Select the database authentication to use.
6. Select the Browse Content option and click Next.
7. On the Browse Content page, select the site collection, site, or list that needs to be restored and click Restore. The restoration should be complete at this point.

#### **MORE INFO RESTORING CONTENT FROM UNATTACHED CONTENT DATABASES**

See <http://technet.microsoft.com/en-us/library/hh269602.aspx> for more information about restoring content from unattached content databases in SharePoint 2013.

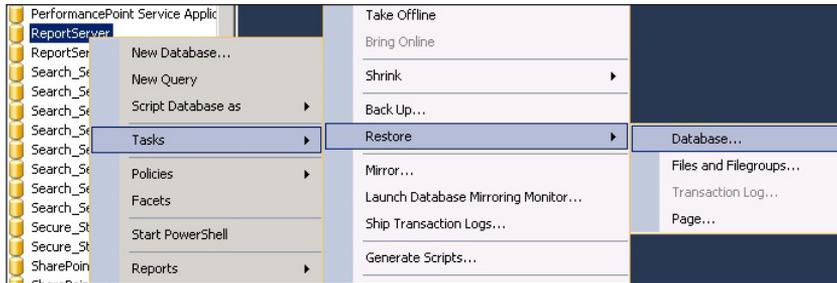


### EXAM TIP

The exam expects familiarity with restoring content from an unattached database. The PowerShell syntax involved with restoring content from an unattached database is one of the few commands you should memorize.

Restoring a content database over an existing one using SQL Server totally replaces the existing database. Before doing this, make sure that that's what you want, because all data between the last backup and the restore is lost. You should be familiar with the process of restoring a database and the steps involved. Assuming that you have used SQL Server to back up a database, you can restore a database using the following steps:

1. Open Microsoft SQL Server Management Studio.
2. Expand the Databases folder in the Object Explorer.
3. Right-click the database you want to restore.
4. Choose Tasks | Restore | Database, as shown in Figure 1-16.



**FIGURE 1-16** The Restore Database option in Microsoft SQL Server Management Studio

5. In the list of Backup Sets To Restore, select the backup set to restore and then click OK. The dialog box remains open and shows progress until the task completes or fails. Depending on the database size and the server speed, it could take minutes or hours.
6. After the restore finishes successfully, click OK in the dialog box that appears. The database is now ready to use.

### IMPORTANT EXCLUSIVE ACCESS FOR RESTORATION

Exclusive access to a database is necessary to be able to restore it. If this is a SharePoint database, you have to stop any processes that use it (for example, stop the web application that's using the database).

This restore option uses the defaults and is sufficient for most database restoration needs.

### **MORE INFO RESTORING CONTENT DATABASES IN SHAREPOINT 2013**

The other task options—Files and Filegroups, Transaction Log, and Page—are available on the Restore Database dialog box are more advanced restoration options not covered in this book, but you can find out more about them at <http://technet.microsoft.com/en-us/library/ee748604.aspx#proc3>.



### ***Thought experiment***

#### **Choosing the proper database recovery model**

In this thought experiment, apply what you've learned about this objective. You can find answers to these questions in the "Answers" section at the end of this chapter.

Your assignment is to decide which recovery model would be best for the content databases in the SharePoint farm. The SQL Server nodes in the farm run on an AlwaysOn FCI group. The decision for high availability is often out of the hands of the designer of the SharePoint farm and in this case was made by the database administrators at your organization. The content databases range in size but go up to 200 GB each, and the expected size of the databases occupies from 60 percent to 70 percent of the space available on the storage allocated to the SharePoint farm. The only decision you have concerning the content databases is which recovery model to choose: Simple or Full.

Which model would make the most sense based on these restrictions?

## **Objective summary**

- SQL Server and SharePoint backup schedules are completely separate entities that need individual planning.
- Testing the amount of time to recover and the acceptable RTO and RPO will drive backup schedules.
- Keeping a non-production environment as close to the production environment is crucial for testing any changes and for developing solutions with minimal impact.
- The configuration of a SharePoint farm occurs in several different areas that need to be backed up and maintained separately.
- You can restore SharePoint content using various methods, including SharePoint Central Administration, PowerShell, and SQL Server.
- You can use Central Administration to back up and restore individual service applications.

## Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the “Answers” section at the end of this chapter.

1. You need to back up the farm configuration. What are valid options?
  - A. Backup and Restore in Central Administration
  - B. PowerShell *Backup-SPFarm* method
  - C. *Stsadm* commands
  - D. All of the above
2. You are backing up SQL Server content databases, and a particular database needs point-in-time recovery. Which recovery model is required for the database?
  - A. Bulk logged
  - B. Simple
  - C. Full
  - D. Any of the above will work
3. A backup of a content database is restored to the production SQL Server. It needs to be attached to the SharePoint farm so that the backup data can be compared to production data. For that to happen, what needs to happen to the restored database?
  - A. It can simply be added as a content database in Central Administration.
  - B. It must be attached with a new database ID using a command-line function.
  - C. It can't be attached to the same SharePoint farm.
  - D. It can be attached with Central Administration but not to the same web application.
4. Which of the following isn't/aren't backed up by a configuration-only backup of the SharePoint farm via Central Administration?
  - A. SSL Certificates
  - B. DLLs in the Global Assembly Cache (GAC)
  - C. IIS-specific settings not managed through SharePoint
  - D. All of the above

5. Which of the following can be backed up as individual components using Central Administration?
- A. Global search settings
  - B. SharePoint Server State service
  - C. InfoPath Forms services
  - D. All of the above

## Chapter summary

---

- High availability for SQL Server data is the cornerstone for any SharePoint disaster recovery plan, and you can choose from a wide range of options.
- High availability for SharePoint data not stored in databases should be planned for separately as part of a comprehensive disaster recovery plan.
- Technologies such as Network Load Balancing (NLB) can help to ensure a high level of availability of SharePoint content.
- Resources and criticality of data dictate frequency of backups (both of databases and SharePoint-specific products).
- You can use Central Administration Backup to back up SharePoint-specific items such as service applications and settings.
- Some items, such as IIS settings, aren't backed up by SQL Server backups or Central Administration Backup and must be maintained separately.

# Answers

---

This section contains the solutions to the thought experiments and answers to the lesson review questions in this chapter.

## Objective 1.1: Thought experiment

In this scenario, database mirroring provides the best option. It provides high availability, and the 5 minute failover is well within the amount of time needed to fail over to the secondary database. Using AlwaysOn High Availability Groups isn't a good solution because SQL Server 2008 R2 is the available server and, based on the training and version, isn't a viable option. Clustering isn't an option because it doesn't provide high availability of the content databases, and log shipping isn't a good option due to the amount of time it takes to manually move over to a new database.

## Objective 1.1: Review

### 1. Correct answer: B

- A. Incorrect:** An RTO of 99.9 percent means more than 8 hours of allowed downtime.
- B. Correct:** An RTO of 99.99 percent means a maximum of 53 minutes downtime in a 24-hour-a-day, seven-days-a-week environment.
- C. Incorrect:** An RTO of 99.999 percent (or *five nines* as it's referred to sometimes) means a downtime of less than 6 minutes—an admirable but expensive goal.
- D. Incorrect:** An RTO of 1 percent means the farm would be down 99 percent of the time—obviously not a desirable goal.

### 2. Correct answer: A

- A. Correct:** SharePoint 2013 supports multiple label domains.
- B. Incorrect:** SharePoint 2013 doesn't support single-label domains.
- C. Incorrect:** Only A is correct.
- D. Incorrect:** A is the correct answer.

### 3. Correct answer: D

- A. Incorrect:** Transaction log backups can't use the simple recovery model.
- B. Incorrect:** Log shipping can't use the simple recovery model.
- C. Incorrect:** Database mirroring can't use the simple recovery model.
- D. Correct:** Clustering is unaware what recovery models are used on what databases.

**4. Correct answer: C**

- A. Incorrect:** Clustering can't be used to provide high availability for just a single content database. It's also the most expensive and complicated way to provide high availability in this situation.
- B. Incorrect:** Transaction log shipping is a valid option here but isn't as effective as database mirroring.
- C. Correct:** Database mirroring is an excellent choice for providing high availability for an individual database.
- D. Incorrect:** Not all high-availability options are equal.

**5. Correct answers: B and C**

- A. Incorrect:** WSFC must be configured on *all* the Windows Server instances running AlwaysOn High Availability Groups.
- B. Correct:** SQL Server must be configured as well as the Windows Server instance it's running on.
- C. Correct:** All Windows servers need to be configured that will be part of the AlwaysOn high-availability group.
- D. Incorrect:** SharePoint 2013 doesn't need to be aware of an AlwaysOn high-availability group because it's aware of it only as though it were a single SQL Server.

## Objective 1.2: Thought experiment

You need to deal with two issues in this thought experiment. The first is how to distribute the load across multiple WFEs. Network Load Balancing (NLB) is an excellent choice for this because it not only provides a way to distribute the processing load across several servers, but it also automatically detects when a Windows Internet Name Service (WINS) server becomes unavailable. This solves both the primary and secondary goals of the WFE requirement. The database failover goal can be solved by using SQL Server aliases. Because you know that SQL Server is being mirrored, you can modify the alias to point to the secondary or WINS server if the primary goes down or has been brought down for maintenance. Therefore, the combination of NLB and SQL Server aliases should provide a solution that meets the business goals outlined.

## Objective 1.2: Review

1. **Correct answer:** B
  - A. **Incorrect:** Servers running SQL Server don't need aliases created on them.
  - B. **Correct:** All SharePoint servers that connect to the database need to be configured to use SQL Server aliases.
  - C. **Incorrect:** Only the SharePoint servers need to be configured to use SQL Server aliases.
  - D. **Incorrect:** To use SQL Server aliases in a SharePoint farm, they must be configured on the SharePoint servers.
  
2. **Correct answer:** D
  - A. **Incorrect:** WFE is one role a SharePoint server can assume but not the only one.
  - B. **Incorrect:** Search server is a processor-intensive role, but the SharePoint server that assumes this role can do others as well.
  - C. **Incorrect:** Application server is also an intensive role, but a SharePoint server can also handle more than just this role.
  - D. **Correct:** A SharePoint server can take on any or all of the roles listed above.
  
3. **Correct answer:** B
  - A. **Incorrect:** Central Administration is unaware that NLB is being used, yet NLB is the recommended way to distribute the load across multiple WFEs.
  - B. **Correct:** Configuring NLB on each SharePoint server WFE is the recommended way to distribute the load.
  - C. **Incorrect:** IIS isn't involved in configuring an NLB cluster.
  - D. **Incorrect:** Because B is the only correct answer, D can't be the answer.
  
4. **Correct answer:** C
  - A. **Incorrect:** The WFE layer communicates with the TCP/IP layer.
  - B. **Incorrect:** The network layer is the lowest layer in the communication stack.
  - C. **Correct:** Network Load Balancing (NLB) is between the TCP/IP layer and the network adapter layer.
  - D. **Incorrect:** The layer above the WFE is the interaction between the user and the WFE.

**5. Correct answers:** B and D

- A. Incorrect:** RAID 0 is writing data across several disks at the same time for increased performance in disk operations.
- B. Correct:** RAID 1 is disk mirroring, in which two exact copies of the data exist.
- C. Incorrect:** RAID 5 is like RAID 0 but with data parity so that a disk can be lost, but it provides no data redundancy.
- D. Correct:** RAID 10 has the benefits of RAID 0 and RAID 1.

## Objective 1.3: Thought experiment

In this scenario, you are presented with a difficult decision. Because the SQL Server high-availability plan is out of your hands, you have only the choice of which recovery model to use for the SharePoint farm content databases. Clusters aren't aware of what recovery model individual databases use, so the only factor involved is how much disk space is being used and if point-in-time recovery is required. Because point-in-time recovery wasn't specified as a requirement, you can use the Simple recovery model to save space because the logs are truncated every time the database is fully backed up. Alerting the people involved that point-in-time recovery isn't possible in this scenario is always prudent. This is based on using SQL Server backups and not a third-party backup process.

## Objective 1.3: Review

**1. Correct answer:** D

- A. Incorrect:** Backup and restore in Central Administration is a valid option, but it's not the only one.
- B. Incorrect:** PowerShell is only one of the valid options for backing up the farm configuration.
- C. Incorrect:** *STSADM* is still available in SharePoint 2013, although it's recommended that commands be executed via PowerShell.
- D. Correct:** All the preceding options are valid ways to back up the farm configuration.

**2. Correct answer:** C

- A. Incorrect:** Bulk logged backups can restore only to last backup.
- B. Incorrect:** Simple backups can restore only the entire backup.
- C. Correct:** The Full recovery mode is the only option that allows for point-in-time restore.
- D. Incorrect:** Because C is the only correct answer, this one can't be correct.

**3. Correct answer: B**

- A. Incorrect:** An error occurs if you try to attach a content database with the same ID.
- B. Correct:** To attach a backed-up content database to the same SharePoint farm, it must have a new ID (for example, using the PowerShell command *Mount-SPContentDatabase* with the *AssignNewDatabaseId* parameter).
- C. Incorrect:** A content database can be attached as long as it has a new ID.
- D. Incorrect:** Whether it's within the same web application or not doesn't matter. A content database with the same ID can't be attached within the same SharePoint farm.

**4. Correct answer: D**

- A. Incorrect:** SSL certificates aren't backed up by the Central Administration tool, but it's not the only correct answer in the list.
- B. Incorrect:** Again, DLLs aren't backed up as part of the Central Administration farm backup (either configuration only or full).
- C. Incorrect:** IIS-only settings are backed up separately from the SharePoint farm.
- D. Correct:** All the preceding answers are correct and not backed up via the SharePoint Central Administration tool (either the configuration only or the full backup).

**5. Correct answer: D**

- A. Incorrect:** Global Search Settings can be backed up individually, but it's just one of the right answers.
- B. Incorrect:** SharePoint Server State Service is just one of the right answers.
- C. Incorrect:** InfoPath Forms Services is just one of the right answers.
- D. Correct:** All of the preceding options are correct.

# Index

## A

- Access apps, 231, 238
- Access Services 2010, 231–233
- Access Services Application Name text box, 232
- Access Services, configuring, 231–239
  - Access Services, 234–238
  - Access Services 2010, 231–233
- Active Directory Domain Services (AD DS) permissions, 77
- Active Directory Logon Data connections, 81
- Active Directory Resource connections, 81
- Active property (channels), 111
- activity feeds, social workloads, 77–81
- Activity Ranking analysis, 108
- Add Content Types page (Analysis Services), 328–329
  - adding
    - query suggestions (seraches), 91–92
- Additional Lock Information text box, 174
- Address setting, trusted file locations, 226
- AD DS (Active Directory Domain Services) permissions, 77
- Add-SPSolution command, 296
- Administrator's Toolkit, 163–164
- Advanced Encryption Security (AES), 10
- AES (Advanced Encryption Security), 10
- Alert job, log shipping, 16
- alerts
  - provisioning for Reporting Services, 313–315
- aliases (SQL Server), 34–36
- Alias property (channels), 111
- Allow Cross Domain Access option, configuring Excel Services application, 223
- Allow External Data setting, trusted file locations, 228
- Allow External Data Using REST setting, trusted file locations, 229
- Allow Non-Remotable Queries setting, configuring Access Services 2010, 233
- Allow Outer Joins setting, configuring Access Services 2010, 232
- Allow Triggers To Fire Others setting, configuring Access Services, 235
- AlwaysOn Failover Cluster Instances, 14–16
- AlwaysOn High Availability Groups technology, 12–14
- AlwaysOn (SQL Server), 11–15
- analysis of data, BI tools, 304–330
  - Excel Services data modeling, 326
  - PerformancePoint, 304–310
  - PowerPivot, 319–326
  - Power View, 327–329
  - Reporting Services, 310–319
- Analysis Services EffectiveUserName setting, configuring global settings, 225
- Analysis Services in Tabular mode, installation, 327–328
- analysis, upgrades
  - content database test results, 153–155
- analytics processing component (Search server), 25
- analytics reports, searches, 107–108
- Anchor Text Processing analysis, 107
- App Catalog
  - adding SharePoint apps to, 299–300
  - configuring settings, 211
  - creating, 210–211
  - manual addition of apps, 212
- App Domain setting, 214
- AppFabric Caching service, 80
- Application Discovery service application, 259
- Application Objects section (Manage Service Applications page), 233
- application pool, configuring for Access Services service application, 237–238
- application server services, 26

## App Management, creating and configuring

App Management, creating and configuring, 209–218

App Store, 210–212

DNS entries, 215–216

subscriptions, 212–214

wildcard certificates, 216–218

AppPoolConcurrentUpgradeSessionLimit property (throttling), 196

\$appPool -Name <ServiceName> -DatabaseName <ServiceDB> command, 213

App Prefix, 214

App Requests list, 211–212

apps

App Management, 209–218

App Store, 210–212

DNS entries, 215–216

subscriptions, 212–214

wildcard certificates, 216–218

deployment, 299–300

upgrades, 301–302

App Store, configuring, 210–212

app URLs, 214

assemblies

User-Defined Function Assemblies, 230

Assembly Location option, , adding user-defined function assemblies, 230

Assembly option, adding user-defined function assemblies, 230

AssignmentCollection parameter (Test-SPContentDatabase command), 154

attaching read-only content databases, 176

audiences, social permissions, 71–72

authentication

server-to-server, configuring, 248

web applications, 155–158

authoritative pages, searches, 93–94

automatic deletion, record disposition, 135

automatic failover

AlwaysOn FCI, 14

SQL Server mirroring, 7

automatic translation of content, 125–126

## B

backup and recovery planning

business continuity management, 38–51

Central Administration backup, 39–42

content recovery, 47–51

farm configuration recovery, 44–45

non-production environment content refresh, 43–44

service application recovery, 46–47

SharePoint backup schedules, 38–39, 42

SQL Server backup schedules, 43–44

backup compression, 17

backups

data redundancy, 21

log shipping, 16

schedules, 38–39, 42–43

Backup-SPFarm -Directory <path> -BackupMethod Full -Item <service name> command, 274

Backup Tool (Central Administration), 39–42

badges, Community Sites, 64–65

BCS (Business Connectivity Service), configuring, 265–272

importing BCS models, 265–270

model security, 270–271

BDC (Business Data Catalog) models, 265–270

BDC (Business Data Connectivity) service

application, 178–179

BI (business intelligence) tools, 304–330

Excel Services data modeling, 326

PerformancePoint, 304–310

PowerPivot, 319–326

Power View, 327–329

Reporting Services, 310–319

budget limitations, high-availability planning, 4

Bulk Crawled Property Settings, 106

Bulk Logged recovery model, 39

Business Connectivity Service (BCS), configuring, 265–272

importing BCS models, 265–270

model security, 270–271

business continuity

backup and recovery plan, 38–51

backup schedules, 38–39, 42

Central Administration backup, 39–42

content recovery, 47–51

farm configuration recovery, 44–45

non-production environment content refresh, 43–44

service application recovery, 46–47

SQL Server backup schedules, 43–44

high-availability plan, 24–37

network redundancy, 29–32

physical server distribution, 27–29

- service distribution, 24–26
- service instance configuration, 26–27
- SQL Server aliases, 34–36
- SQL Server load balancing, 32–34
- SQL Server high-availability plan, 1–22
  - AlwaysOn, 11–15
  - clustering, 10–11
  - gathering requirements, 2–5
  - login replication, 21
  - log shipping, 15–20
  - mirroring, 6–11
  - storage redundancy, 20–21
- Business Data Catalog (BDC) models, 265–270
- Business Data Connectivity (BDC) service application, 178–179
- Business Data Connectivity connections, 81
- business intelligence (BI) tools, 304–330
  - Excel Services data modeling, 326
  - PerformancePoint, 304–310
  - Power View, 327–329
  - PwerPivot, 319–326
  - Reporting Services, 310–319

## C

- Cache Timeout setting
  - configuring Access Services, 237
  - configuring Access Services 2010, 233
- Caching Of Unused Files setting, configuring global settings, 225
- Canonical Name (CNAME) records, 216
- capacity planning, 137–138
- Central Administration
  - attaching an App Catalog to a web application, 300
  - backing up Secure Store service, 273
  - backups, 39–42
  - changing Secure Store target application permissions, 278
  - changing trusted data source locations, 308–309
  - configuring My Site host, 66–68
  - configuring unattended service accounts, 306
  - consuming service applications, 261–262
  - creating PerformancePoint service application, 304–305
  - creating Project Server service application, 335–336
  - creating Project Server web app, 336–337

- creating SQL Server PowerPivot service application, 323–324
- database recovery, 49–51
- deploying farm solutions, 297–298
- exporting BDC models, 268
- importing BCS models, 266–267
- limiting locations allowed for trusted data sources, 309–310
- provisioning with Psconfig, 167–169
- publishing service applications, 260–261
- restoring Secure Store service application, 274
- setting sandboxed solution quotas, 292–293
- starting Project Server Application Service, 335
- Upgrade Status page, 172
- certifications (Microsoft), xvi
- changing
  - default result source, 95
- channels
  - configuring WCM, 111–113
- Choose A File dialog box, 299
- claims-based authentication
  - upgrade from classic-mode authentication, 156–159
- claims-based authentication directory service, 81
- classic-mode authentication
  - upgrading to claims-based authentication, 156–159
- cleaning environment, migration pre-check, 151
- Click Distance analysis, 107
- clustering (SQL Server), 10–11
  - data redundancy, 20
  - multi-subnet failover, 4
- cluster resource (WFSC), 13
- CMIS (Content Management Interoperability Services), 163
- CNAME (Canonical Name) records, 216
- Collocated mode (Distributed Cache), 80
- commands
  - Add-SPSolution, 296
  - \$appPool -Name <ServiceName> -DatabaseName <ServiceDB>, 213
  - Backup-SPFarm -Directory <path> -BackupMethod Full -Item <service name>, 274
  - databaserepair, 160
  - \$db.Repair(\$false), 160
  - \$db.Repair(\$true), 160
  - EnumAllWebs, 153
  - Export-SPBusinessDataCatalogModel, 267
  - Get-SPBusinessDataCatalogMetadataObject, 267

- Get-SPFarm | Select Id, 258
- Get-SPSiteUpgradeSessionInfo, 193
- Get-SPTopologyServiceApplication, 259
- Import-SPBusinessDataCatalogModel, 267
- Install-SPRSService, 312
- Install-SPRSServiceProxy, 312
- Install-SPSolution, 296
- Mount-SPContentDatabase, 172
- New-AuthServer, 340
- New-SPBusinessDataCatalogServiceApplicationProxy, 262
- New-SPEnterpriseSearchServiceApplicationProxy, 262
- New-SPMetadataServiceApplicationProxy, 262
- New-SPProfileServiceApplicationProxy, 262
- New-SPSecureStoreServiceApplicationProxy, 262
- Publish-SPServiceApplication -Identity <Service App GUID>, 261
- Receive-SPServiceApplicationConnectionInfo -FarmURL <Published URL>, 262
- Remove-SPSite, 159
- Remove-SPSiteUpgradeSessionInfo -Identity, 194
- Repair-SPSite -Identity, 185
- Request-SPUpgradeEvaluationSiteCollection, 192
- Restore-SPFarm -Directory <path> -Item <service name> -RecoveryMethod Overwrite, 275
- \$rootCert.Export("Cert")|Set-Content <Root Cert Path> -Encoding byte, 256
- \$rpa.ResourceMeasures["<resource name>"].Update(), 294
- \$rpa.ResourceMeasures | Select-Object Name,ResourcesPerPoint,AbsoluteLimit, 293
- \$rpa.Update(), 294
- Set-SPSite -Identity "<site collection"> -QuotaTemplate "<template name>", 293
- Start-SPServiceInstance, 169
- Stop-SPServiceInstance, 169
- \$stsCert.Export("Cert")|Set-Content <STS Cert Path> -Encoding byte, 256
- LocalLoginProvider.SigningCertificate, 256
- Test-SPContentDatabase, 152
- Test-SPSite -Identity, 185
- Uninstall-SPSolution, 297
- Update-SPSecureStoreApplicationServerKey -Passphrase <passphrase>, 275
- Update-SPSolution, 299
- Upgrade-SPSite, 194
- \$wa.MigrateUsers(\$true), 157
- \$wa.ProvisionGlobally(), 157
- \$wa.Update(), 156
- \$webApp.CompatibilityRange, 189
- \$webApp.GrantAccessToProcessIdentity("<domain account>"), 221
- \$webApp.SiteUpgradeThrottleSettings, 195
- \$webApp.Update(), 189
- \$webApp.UpgradeMaintenanceLink, 189
- \$webApp.UpgradeReminderDelay, 189
- Community Sites, 60–65
  - badges, 64–65
  - reputation settings, 63–64
- CompatibilityLevel property, 296
- Complete Certificate Request dialog box, 218
- comp tool, 161
- conditions, query rules, 100
- configuration issues, upgrades, 162–163
- Configuration Wizard, 74
- Configure Quotas And Locks option, 173
- configuring
  - AlwaysOn High Availability Groups, 12
  - App Catalog settings, 211
  - App Store, 210–212
  - BI (business intelligence) tools, 304–330
    - Excel Services data modeling, 326
    - PerformancePoint, 304–310
    - PowerPivot, 319–326
    - Power View, 327–329
    - Reporting Services, 310–319
  - content deployment jobs, 121
  - DNS entries, 215–216
  - ECM (Enterprise Content Management), 128–138
    - eDiscovery, 128–131
    - large document repositories, 137–138
    - record disposition and retention, 135–137
    - record management, 131–135
  - FIM (Forefront Identity Manager) services, 76
  - log shipping, 16–18
  - Machine Translation, 125–126
  - Machine Translation Services, 247–250
    - creating service application, 248
    - running services, 251
    - settings, 249–251
  - Microsoft Access, 231–239
    - Access Services, 234–238
    - Access Services 2010, 231–233
  - Microsoft Excel Services, 221–230
  - Microsoft PowerPoint Conversion Services, 246–247

- Microsoft Visio Services, 239–242
  - unattended service accounts, 239–240
  - Visio Graphics Service settings, 240–242
- Microsoft Word Automation Services, 242–245
  - creating service application, 243–244
  - file conversions, 245–246
  - modifying properties, 243–244
- My Site host, 66–68
- NLB (Network Load Balancing), 30–31
- parallel upgrades, 171–173
- read-only access for content, 173–175
- reputation settings, community sites, 63–64
- sandboxed solution management, 295–296
- searches
  - analytics reports, 107–108
  - end-user experience, 99–102
  - index freshness, 94–95
  - result relevancy, 87–94
  - results, 98
  - result sources, 95–99
  - search schema, 102–107
- server-to-server authentication, 248
- service applications, 209
  - App Management, 209–218
  - BCS (Business Connectivity Service), 265–272
    - federation of services, 253–263
    - productivity services, 220–251
    - Secure Store, 272–278
- service instances, 26–27
- site collections available for upgrading, 188–189
- site mailboxes, 133
- Trusted My Site Host Locations, 68–69
- upgrade farms, 176–178
- variation settings, 123
- Visio Graphics Service settings, 241–243
- WCM (Web Content Management)
  - channels, 111–113
  - content deployment, 119–122
  - product catalogs, 113–117
  - topic pages, 117–118
  - variations, 122–126
- web application authentication, 155–158
- wildcard certificates, 216–218
- work management, 332–342
  - Microsoft Exchange, 339–341
  - Microsoft Project Server, 332–339
- Confirm Pass Phrase text box, 273
- Connection Encryption option, configuring Excel Services application, 223
- Connection Lifetime setting, configuring global settings, 225
- Connection Permissions dialog box, 259
- connections
  - social workloads, 81–85
- Connect To Server dialog box, 314
- consuming farm, importing root certificates, 256
- consuming service applications
  - federation of services, 261–262
- consumption, product catalogs, 116
- content
  - deployment
    - WCM (Web Content Management), 119–122
    - processing component (Search server), 25
    - recovery, 47–51
    - refresh, non-production environments, 43–44
    - translation, 125–126
- content databases
  - configuring read-only access, 175
  - SharePoint migration pre-check, 152
  - testing prior to migration, 153–155
- Content Deployment, 44
- Content Deployment Source Feature, 120
- content evaluation, upgrades, 149–164
  - analyzing content database test results, 153–155
  - configuration issues, 162–163
  - configuring web application authentication, 155–158
  - migration pre-check tasks, 150–153
  - resolving missing files, 160–162
  - resolving orphan objects, 158–160
- Content Management Interoperability Services (CMIS), 163
- Content Organizer, 134–135
- Content Search Web Part (CSWP), 119–120
- continuous crawling, content sources, 94–95
- contributors, community sites, 62
- Conversion Processes setting, configuring Word Automation Services, 244
- Conversion Throughput setting, configuring Word Automation Services, 244
- Copy job, log shipping, 16
- crawl component (Search server), 25
- Crawl component (Search service application), 168
- crawled properties, 103

## crawling a content source

- crawling a content source, 94–95
- Create New Word Automation Services Application dialog box, 243
- Create Personal Site permission, 72
- Create PowerPivot Service Application task pane, 322
- Create Project dialog box, 163
- Create Upgrade Evaluation Site Collection option, 192
- creating
  - Access apps, 238
  - Access Service service application, 236
  - App Catalog, 210–211
  - audiences, 71–72
  - My Site host, 66
  - PerformancePoint service application, 304–305
  - PowerPivot service application, 322–323
  - Project Server service application, 335–336
  - Reporting Services service application, 312–313
  - result sources, 96
  - result types, 102
  - service applications, 209
    - App Management, 209–218
    - BCS (Business Connectivity Service), 265–272
    - federation of services, 253–263
    - productivity services, 220–251
    - Secure Store, 272–278
  - source variation, 124
  - SQL Server aliases, 34–35
  - Synchronization Connection with AD, 82–83
  - thesaurus, 88–89
  - unattended service accounts, 239–240
  - User Profile service applications, 74–75
  - Visio Graphics Service service application, 240
  - work management, 332–342
    - Microsoft Exchange, 339–341
    - Microsoft Project Server, 332–339
- credentials
  - managing Secure Store target application permissions, 278–279
- critical sites, upgrading site collections, 188
- cross-site publication, 119–123
- CSWP (Content Search Web Part), 119–120
- custom entities, searches, 89–91
- customizations
  - solution management, 291–302
    - deployment of apps, 299–300
    - deployment of farm solutions, 296–298
    - sandboxed solution quotas, 292–296
    - upgrading apps, 301–302

- upgrading farm solutions, 298–299
- customizations, evaluating before upgrades, 149–164
  - analyzing content database test results, 153–155
  - configuration issues, 162–163
  - configuring web application authentication, 155–158
  - migration pre-check tasks, 150–153
  - resolving missing files, 160–162
  - resolving orphan objects, 158–160
- customized sites, upgrading site collections, 188
- Custom Web App dialog box, 238

## D

- Dashboard Designer, 307
- data analysis, BI tools, 304–330
  - Excel Services data modeling, 326
  - PerformancePoint, 304–310
  - PowerPivot, 319–326
  - Power View, 327–329
  - Reporting Services, 310–319
- database-attach method (upgrade path), 151
- DatabaseCredentials parameter (Test-SPContentDatabase command), 154
- databaserepair command, 160
- databases
  - mirroring
    - log shipping, 19
    - SQL Server, 6–11
  - recovery
    - Central Administration, 49–51
    - Restore-SPSite command, 48–49
    - SQL Server, 48
- Database Server text box, 305
- databases, upgrades
  - analyzing test results, 153–155
- data connection libraries
  - configuring Microsoft Excel Services, 229–230
- data modeling
  - Excel Services, 326
- Data Model Settings
  - configuring Microsoft Excel Services, 230
- Data Protection Manager (DPM), 47
- data redundancy, 20–21
- data sources, configuring BCS models, 265
- Data-Tier Application Framework (Microsoft SQL Server 2012), 235

- dbcreator built-in server role, 234
- \$db.Repair(\$false) command, 160
- \$db.Repair(\$true) command, 160
- Dedicated mode (Distributed Cache), 80
- Deep Links analysis, 108
- Default Account, PowerPivot Configuration Tool, 321
- Default index update group, 103
- default result source, 95
- deletecorruption parameter, 160
- deleting sandboxed solutions, 295
- deployment
  - apps, 299–300
  - farm solutions, 296–298
- deployment of content
  - WCM (Web Content Management), 119–122
- Description option, adding user-defined function assemblies, 230
- Description property (channels), 111
- Description setting, trusted file locations, 226
- Design Manager, creating channels, 111
- Device Inclusion Rules property (channels), 111
- Diagnostic Studio, 162–163
- dialog boxes
  - Choose A File, 299
  - Complete Certificate Request, 218
  - Connection Permissions, 259
  - Connect To Server, 314
  - Create New Word Automation Services Application, 243
  - Create Project, 163
  - Custom Web App, 238
  - Generate New Key, 273
  - New Linked Server, 33
  - New Resource Record, 216
  - Permissions For User Profile Service Application, 72–73
  - Publish Service Application, 261
  - Request Certificate, 217
  - Secondary Database Settings, 17
  - Site Bindings, 218
  - SQL Server Installation Center, 311
  - Transaction Log Backup Settings, 17
- directory services, 81
- Disable Office 97-2003 Document Scanning? setting, configuring Machine Translation Service, 250
- disaster recovery planning
  - backup and restore plans, 38–51
  - Central Administration backup, 39–42
  - content recovery, 47–51
  - farm configuration recovery, 44–45
  - non-production environment content refresh, 43–44
  - service application recovery, 46–47
  - SharePoint backup schedules, 38–39, 42
  - SQL Server backup schedules, 43–44
- SharePoint
  - network redundancy, 29–32
  - physical server distribution, 27–29
  - service instance configuration, 26–27
  - SQL Server aliases, 34–36
  - SQL Server load balancing, 32–34
- SQL Server, 1–22
  - AlwaysOn, 11–15
  - clustering, 10–11
  - gathering requirements, 2–5
  - login replication, 21
  - log shipping, 15–20
  - mirroring, 6–11
  - storage redundancy, 20–21
- Discovery Management role group, 340
- Display Granular External Data Errors setting, trusted file locations, 228
- Display Name text box, 240
- display templates, searches, 101–102
- disposition of records
  - ECM (Enterprise Content Management), 135–137
- Distributed Cache
  - activity feeds, 80–81
- distribution, SharePoint services, 24–26
- DNS (Domain Name Services)
  - configuring entries, 215–216
- document libraries, migration pre-check, 151
- document repositories
  - ECM (Enterprise Content Management), 137–138
- document routing, 134–135
- Domain Name Services (DNS)
  - configuring entries, 215–216
- downloading
  - Microsoft SQL Server 2012 Feature Pack, 235
  - PowerPivot, 319
- DPM (Data Protection Manager), 47
- Drop Off Library (record centers), 134
- Duplicate Submissions setting, 134

## E

- earned badges, community sites, 64
- ECM (Enterprise Content Management), 128–138
  - eDiscovery
    - , 128–131
    - large document repositories, 137–138
    - record disposition and retention, 135–137
    - record management, 131–135
  - ECS (Excel Calculation Services) service, 221
  - eDiscovery, 128–131
    - configuring a connection to Exchange, 339–341
  - Edit permission, 272
  - Email Notifications section (My Site host), 68
  - Email Settings (Manage Reporting Services Application page), 318–319
  - Embedded Font Support setting, configuring Word Automation Services, 244
  - Enable Assembly option, adding user-defined function assemblies, 230
  - Enable Contained Database setting, configuring Access Services, 235
  - Enabled File Extensions setting, configuring Machine Translation Service, 249
  - Enable This As A Primary Database In A Log Shipping Configuration option, 17
  - enabling
    - Distributed Cache, 80
    - monitoring of offensive content, 62
    - Standard Collection feature, 61
  - encryption
    - AES (Advanced Encryption Security), 10
    - RC4, 10
  - encryption key, exporting
    - User Profile Synchronization service application, 177–178
  - end-user experience, searches, 99–102
    - display templates, 101–102
    - query rules, 99–101
    - result types, 102
  - Enterprise Content Management (ECM), 128–138
    - eDiscovery, 128–131
    - large document repositories, 137–138
    - record disposition and retention, 135–137
    - record management, 131–135
  - EnumAllWebs command, 153
  - environment
    - ECM (Enterprise Content Management), 128–138
    - search workloads, 87–108
    - social workloads, 59–85
    - upgrades and migrations
      - evaluation of content and customizations, 149–164
      - planning process, 166–182
      - site collections, 184–197
    - WCM (Web Content Management), 110–126
  - error messages, read-only content, 174
  - evaluation of content, upgrades, 149–164
    - analyzing content database test results, 153–155
    - configuration issues, 162–163
    - configuring web application authentication, 155–158
    - migration pre-check tasks, 150–153
    - resolving missing files, 160–162
    - resolving orphan objects, 158–160
  - evaluation sites, upgrading site collections, 191–193
  - Everyone activity feed, 80
  - Everyone (permission level), 73
  - Excel Calculation Services (ECS) service, 221
  - Excel Services
    - data modeling, 326
  - Excel Services Application Data Model Settings page, 326
  - Excel Services, configuring, 221–230
    - data connection libraries, 229–230
    - Data Model Settings, 230
    - global settings
      - external data settings, 225
      - load balancing settings, 223
      - memory usage settings, 224
      - security settings, 223
      - session management settings, 224
      - workbook cache settings, 224–225
    - Trusted Data Providers, 229
    - Trusted File Locations, 225–229
    - User-Defined Function Assemblies, 230
  - Exchange Auto Discovery, 67
  - Exchange, configuring a connection to, 339–341
  - Exchange Server instance, discovering content, 129
  - exchanging trust certificates
    - federation of services, 255–256
  - excluding query suggestions, 92
  - Execute permission, 272
  - Execution account (Manage Reporting Services Application page), 318
  - existing sites

- adding Community Sites feature, 60–61
- exporting
  - BDC models, 268
  - root certificates, 256
- exporting, encryption key
  - User Profile Synchronization service application, 177–178
- Export-SPBusinessDataCatalogModel command, 267
- ExtendedCheck parameter (Test-SPContentDatabase command), 154
- External Data Cache Lifetime setting, trusted file locations, 228
- External Data setting, configuring Visio Graphics Service, 242
- external data settings
  - configuring Microsoft Excel Services, 225
  - trusted file locations, 228–229
- extractors, custom entities, 89–91

## F

- failover cluster instances, 13
- failover cluster (SQL Server), 11
- farm configuration
  - backups, 39
  - recovery, 44–45
- Farm Configuration Wizard, 246
- farm solutions
  - deployment, 296–298
  - upgrading, 298–299
- FAST Search Center, removal before migration, 153
- federated database servers, SQL Server load balancing, 34
- federation of services, configuring, 253–263
  - consuming service applications, 261–262
  - exchanging trust certificates, 255–256
  - permissions, 258–260
  - planning services to federate, 254–256
  - publishing service applications, 260–262
  - trust management, 256–257
- Feed Cache, 80
- File Access Method option, configuring Excel Services application, 223
- file conversions
  - configuring Word Automation Services, 245–246
- FILESTREAM requirements, high-availability planning, 5

- filters
  - synchronization connections, 83–84
- FIM (Forefront Identity Manager) services, 76
- Firewall settings, configuring Access Services, 235
- Folder Partitioning setting, 134
- Follow People and Edit Profile permission, 72
- forced switching, database mirroring, 6
- Forefront Identity Manager (FIM) services, 76
- Forefront Identity Manager Service, 170
- Forefront Identity Manager Synchronization Service, 76, 170
- forms-based authentication directory service, 81
- forward lookup zones, 215–216
- Foundation Subscription Setting service, 212–214
- Frequency With Which To Start Translations setting, configuring Machine Translation Service, 250
- Full recovery model, 39
- full trust solutions. *See* farm solutions

## G

- GAC (Global Assembly Cache), 160
- GAC (Global Assembly Cache) files
  - Central Administration Backup tool, 40
- gathering requirements
  - SQL Server high availability, 2–5
- Generate New Key dialog box, 273
- generating master key
  - Secure Store service application, 272–274
- geographically dispersed mirroring, 9
- Get-SPBusinessDataCatalogMetadataObject command, 267
- Get-SPFarm | Select Id command, 258
- Get-SPSiteUpgradeSessionInfo command, 193
- Get-SPTopologyServiceApplication command, 259
- gifted badges, community sites, 64
- Global Assembly Cache (GAC), 160
- Global Assembly Cache (GAC) files
  - Central Administration Backup tool, 40
- global settings
  - configuring Microsoft Excel Services, 222–225
    - external data settings, 225
    - load balancing settings, 223
    - memory usage settings, 224
    - security settings, 223
    - session management settings, 224

## GradualDelete parameter

- workbook cache settings, 224–225
- configuring Visio Graphics Services service application, 241
- GradualDelete parameter, 159
- Grant Create Child Objects permission, 77
- group mappings, 276

## H

- hash tag, keywords, 77–78
- health check, upgrading site collections, 184–188
- heat limitations, high-availability planning, 5
- hierarchy, product catalogs, 114
- high-availability planning
  - SharePoint, 24–37
    - network redundancy, 29–32
    - physical server distribution, 27–29
    - service distribution, 24–26
    - service instance configuration, 26–27
    - SQL Server aliases, 34–36
    - SQL Server load balancing, 32–34
  - SQL Server, 1–22
    - AlwaysOn, 11–15
    - clustering, 10–11
    - gathering requirements, 2–5
    - login replication, 21
    - log shipping, 15–20
    - mirroring, 6–11
    - storage redundancy, 20–21
- high-performance mode, SQL Server mirroring, 6
- hives, Central Administration Backup tool, 40

## I

- IBM Tivoli Directory Server (ITDS) connections, 82
- Identity parameter (Test-SPContentDatabase command), 154
- importing
  - BCS models, 265–270
  - query suggestion files, 92
- Import-SPBusinessDataCatalogModel command, 267
- incremental profile synchronization, 84–85
- index freshness, searches, 94–95
- indexing component (Search server), 25
- Index Partition and Query component (Search service application), 168

- individual mappings, 276
- inetpub folder, 161
- in-place holds, 130
- in-place upgrades, 151
- installation
  - Analysis Services in Tabular mode, 327–328
  - language packets, 176
  - Project Server, 333–335
  - Reporting Services on SQL Server, 310–312
  - sequence for planning upgrades, 181–182
- Install-SPRSService command, 312
- Install-SPRSServiceProxy command, 312
- Install-SPSolution command, 296
- Interactive queue, personal site collections, 70
- IP resource (WFSC), 13
- ISS settings, Central Administration Backup tool, 40
- ITDS (IBM Tivoli Directory Server) connections, 82

## J

- Job Monitoring setting, configuring Word Automation Services, 244
- Just Me (permission level), 73

## K

- Key component (synonyms), 88
- Key Management link (Manage Reporting Services Application page), 317–318
- keywords, hash tag, 77–78

## L

- labels, variations, 123–124
- Language component (synonyms), 88
- Language Options section (My Site host), 67
- language packs
  - installing/uninstalling, 176
  - pre-upgrade site collection health check, 188
- large document repositories, ECM, 137–138
- levels (RAID), 28
- Likes activity feed, 80
- limitations
  - high-availability plans, 4–5
- linked servers, SQL Server load balancing, 33–34

- Link index update group, 103
- List And Queries section, configuring Access Services 2010, 232
- lists, migration pre-check, 151
- load balancing
  - configuring Microsoft Excel Services, 223
  - Load Balancing service application, 259
  - SQL Server, 32–34
- load testing
  - document repositories, 137
- Load Testing Toolkit, 163
- Local DB (Microsoft SQL Server 2012), 235
- Local setting, configuring global settings, 223
- Local SharePoint Results, 95
- location settings, trusted file locations, 226
- Location Type setting, trusted file locations, 226
- login replication
  - high-availability plan for SQL Server, 21
- log shipping (SQL Server), 15–20

## M

- Machine Translation, 125–126
- Machine Translation Services, configuring, 247–250
  - creating service application, 248
  - running services, 251
  - settings, 249–251
- Manage Content Deployment Path And Jobs page, 122
- Managed Metadata service application, 178
- Manage Excel Services Application settings page, 222
- Manage Jobs link (Manage Reporting Services Application page), 317
- management
  - App Catalog, 300
  - business continuity
    - backup and recovery plan, 38–51
    - high-availability plan for SharePoint, 24–37
    - high-availability plan for SQL Server, 1–22
  - records, ECM (Enterprise Content Management), 131–135
  - Secure Store master key, 274–275
  - Secure Store target application permissions, 278–279
  - solutions, 291–302
    - deployment of apps, 299–300
    - deployment of farm solutions, 296–298
    - sandboxed solution quotas, 292–296
    - upgrading apps, 301–302
    - upgrading farm solutions, 298–299
  - trust, federation of services, 256–257
  - work management, 332–342
    - Microsoft Exchange, 339–341
    - Microsoft Project Server, 332–339
- Manage PerformancePoint Services page, 306
- Manage Project Web Apps page, 337
- Manage Query Rules page, 99
- Manage Reporting Services Application page, 315
  - Email Settings, 318–319
  - Execution account, 318
  - Key Management link, 317–318
  - Manage Jobs link, 317
  - Provision Subscriptions and Alerts link, 319
- manpower limitations, high-availability planning, 5
- manual failover, SQL Server mirroring, 6
- master key
  - Secure Store service application, 272–275
- Maximum Application Log Size setting, configuring Access Services 2010, 233
- Maximum Cache Age setting, configuring Visio Graphics Service, 241
- Maximum Cache Size setting, configuring Visio Graphics Service, 241
- Maximum Calculated Columns Per Query setting, configuring Access Services 2010, 232
- Maximum Character Count For Microsoft Word Documents setting, configuring Machine Translation Service, 249
- Maximum Chart Or Image Size setting, trusted file locations, 227
- Maximum Chart Render Duration setting, trusted file locations, 227
- Maximum Columns Per Query setting, configuring Access Services 2010, 232
- Maximum Concurrent Queries Per Session setting, trusted file locations, 229
- Maximum Conversion Attempts setting, configuring Word Automation Services, 244
- Maximum File Size For Binary Files In KB setting, configuring Machine Translation Service, 249
- Maximum File Size For Text Files In KB setting, configuring Machine Translation Service, 249
- Maximum Memory Usage setting, configuring Word Automation Services, 244

## Max Number Of Items Which Can Be Queued Within A 24-Hour Period

- Maximum Number Of Items Which Can Be Queued Within A 24-Hour Period Per Site Subscription setting, configuring Machine Translation Service, 250
- Maximum Number Of Items Which Can Be Queued Within A 24-Hour Period setting, configuring Machine Translation Service, 250
- Maximum Number Of Synchronous Translation Requests setting, configuring Machine Translation Service, 250
- Maximum Order By Clauses Per Query setting, configuring Access Services 2010, 232
- Maximum Private Bytes setting
  - configuring Access Services, 238
  - configuring global settings, 224
- Maximum Recalc Duration setting, configuring Visio Graphics Service, 241
- Maximum Records Per Table setting, configuring Access Services 2010, 233
- Maximum Request Duration setting
  - configuring Access Services, 237
  - configuring Access Services 2010, 233
  - trusted file locations, 227
- Maximum Rows Per Query setting, configuring Access Services 2010, 232
- Maximum Session Memory setting, configuring Access Services 2010, 233
- Maximum Sessions Per Anonymous User setting
  - configuring Access Services, 237
  - configuring Access Services 2010, 233
- Maximum Sessions Per User setting
  - configuring Access Services, 237
  - configuring Access Services 2010, 233
- Maximum Size Of Workbook Cache setting, configuring global settings, 224
- Maximum Sources Per Query setting, configuring Access Services 2010, 232
- Maximum Synchronous Conversion Requests setting, configuring Word Automation Services, 244
- Maximum Template Size setting, configuring Access Services 2010, 233
- Maximum Translation Attempts setting, configuring Machine Translation Service, 250
- Maximum Unused Object Age setting, configuring global settings, 224
- Maximum Web Drawing Size setting, configuring Visio Graphics Service, 241
- Maximum Workbook Size setting, trusted file locations, 227
- measuring upgrade performance, 179–181
- Members text box, 240
- Memory Cache Threshold setting, configuring global settings, 224
- memory usage settings, configuring Microsoft Excel Services, 224
- Memory Utilization section
  - Global Settings, 224
  - Manage Service Applications page, 233
- Memory Utilization setting, configuring Access Services 2010, 233
- Mentions activity feed, 80
- Metadata service applications, 61
- Microsoft Access Services, configuring, 231–239
  - Access Services, 234–238
  - Access Services 2010, 231–233
- Microsoft certifications, xvi
- Microsoft Excel Services
  - data modeling, 326
- Microsoft Excel Services, creating and configuring, 221–230
  - data connection libraries, 229–230
  - Data Model Settings, 230
  - global settings, 222–225
    - external data settings, 225
    - load balancing settings, 223
    - memory usage settings, 224
    - security settings, 223
    - session management settings, 224
    - workbook cache settings, 224–225
  - Trusted Data Providers, 229
  - Trusted File Locations, 225–229
  - User-Defined Function Assemblies, 230
- Microsoft Exchange, configuring a connection to, 339–341
- Microsoft Identity Integration Server Key Management Utility wizard, 177
- Microsoft PowerPoint Conversion Services, configuring, 246–247
- Microsoft Project Server, configuring a connection to, 332–339
- Microsoft SQL Server 2012 Feature Pack, 235
- Microsoft SQL Server Management Studio
  - Restore Database option, 50
- Microsoft System CLR Types (Microsoft SQL Server 2012), 235
- Microsoft Visio Services, configuring, 239–242
  - unattended service accounts, 239–240

- Visio Graphics Service settings, 240–242
- Microsoft Word Automation Services, configuring, 242–245
  - creating service application, 243–244
  - file conversions, 245–246
  - modifying properties, 243–244
- migrated scopes, 97
- migrations
  - evaluation of content and customizations, 149–164
    - analyzing content database test results, 153–155
    - configuration issues, 162–163
    - configuring web application authentication, 155–158
  - pre-check tasks, 150–153
  - resolving missing files, 160–162
  - resolving orphan objects, 158–160
- planning process, 166–182
  - configuring parallel upgrades, 171–173
  - configuring read-only access for content, 173–175
  - configuring upgrade farms, 176–178
  - installation sequence, 181–182
  - measuring upgrade performance, 179–181
  - removal of servers in rotation, 166–171
- upgrading site collections, 184–197
  - analyzing and resolving health check results, 186–189
  - configuring available site collection modes, 188–189
  - evaluation sites, 191–193
  - health check, 184–185
  - queues and throttles, 193–198
  - self-service upgrades, 190–191
- Minimum Cache Age setting, configuring Visio Graphics Service, 241
- mirroring (SQL Server), 6–11
- mirror servers, 6
- missing files, resolving after upgrades, 160–162
- Mixed Mode Authentication setting, configuring Access Services, 235
- moderators, community sites, 62
- modifying
  - social permissions, 73
- monitoring
  - offensive content, 62
  - Reporting Services resource consumption, 319
- monitor servers
  - log shipping, 16

- Most Popular Items reports, 108
- Mount-SPContentDatabase command, 172
- moving
  - UPS service, 169–170
- MUI references, pre-upgrade site collection health check, 188–189
- multi-subnet failover clustering, 4
- My Site Cleanup section (My Site host), 68
- My Site host, 66–68
- My Sites feature, 65–71
  - My Site host, 66–68
  - personal site collections (SkyDrive), 69–71
  - Trusted My Site Host Locations, 68–69

## N

- Named Pipes setting, configuring Access Services, 235
- Name property (channels), 111
- Native Client (Microsoft SQL Server 2012), 235
- .NET 4.0, upgrading farm solutions, 298
- network bandwidth limitations, high-availability planning, 5
- Network Load Balancing (NLB), 29–32
- network redundancy, 29–32
- New Application Database Server setting, configuring Access Services, 238
- New-AuthServer command, 340
- New Linked Server dialog box, 33
- New Resource Record dialog box, 216
- newsfeeds, 79
- Newsfeed section (My Site host), 68
- New-SPBusinessDataCatalogServiceApplicationProxy command, 262
- New-SPEnterpriseSearchServiceApplicationProxy command, 262
- New-SPMetadataServiceApplicationProxy command, 262
- New-SPProfileServiceApplicationProxy command, 262
- New-SPSecureStoreServiceApplicationProxy command, 262
- New Workbook Session Timeout setting, trusted file locations, 227
- New Zone Wizard, 215
- NLB (Network Load Balancing), 29–32
- nodes (WSFC), 13
- Non-interactive queue
  - personal site collections, 70

## non-production environments

- non-production environments
  - content refresh, 43–44
- Novell eDirectory connections, 82
- Number Of Translations To Start setting, configuring Machine Translation Service, 250

## O

- object models
  - Central Administration Backup tool, 42
- OData data sources, 266
- ODBC DSN trusted provider, 229
- ODBC trusted provider, 229
- offensive content, monitoring, 62
- OLE DB trusted provider, 229
- OPC (Open Packing Conventions) standard, 239
- Open Packing Conventions (OPC) standard, 239
- order of installation, planning upgrades, 181–182
- orphan objects
  - resolving after upgrades, 158–160
- Outlook Social connector, 78
- Overview tab (Resource Monitor), 180–181
- overwriting query suggestion files, 92
- owners, community sites, 62

## P

- parallel upgrades, configuring, 171–173
- People index update group, 103
- PerformancePoint, configuring, 304–310
  - site collections, 309–310
  - trusted data sources, 308–309
  - unattended service accounts, 306–308
- PerformancePoint service application, creating, 304–305
- PerformancePoint Services service application, 178
- performance, upgrades, 179–181
- permissions
  - database mirroring, 10–11
  - federation of services, 258–260
  - managing site collection upgrades, 193
  - Replicating Directory Changes, 77
  - Secure Store target applications, 278–279
  - social, 71–74
  - Use Tag and Notes, 72
  - Write, 77

- Permissions For User Profile Service Application dialog box, 72–73
- personal site collections (SkyDrive), 69–71
- Personal Site Location section (My Site host), 67
- Personal Site Quota, 71
- phases, upgrade process, 171
- physical server distribution, 27–29
- planning
  - BI (business intelligence) infrastructure, 304–330
    - Excel Services data modeling, 326
    - PerformancePoint, 304–310
    - PowerPivot, 319–326
    - Power View, 327–329
    - Reporting Services, 310–319
  - business continuity management
    - backup and recovery, 38–51
    - high-availability plan for SharePoint, 24–37
    - high-availability plan for SQL Server, 1–22
  - environment
    - ECM (Enterprise Content Management), 128–138
    - search workloads, 87–108
    - social workloads, 59–85
    - WCM (Web Content Management), 110–126
  - federation of services, 254–256
  - upgrades, 166–182
    - configuring parallel upgrades, 171–173
    - configuring read-only access for content, 173–175
    - configuring upgrade farms, 176–178
    - installation sequence, 181–182
    - measuring upgrade performance, 179–181
    - removal of servers in rotation, 166–171
- political limitations, high-availability planning, 5
- Popularity Trends reports, 108
- Port value, PowerPivot Configuration Tool, 322
- possible owners (WFSC), 13
- PostUpdate web service, 301
- power availability, high-availability planning, 5
- PowerPivot, 319–326
- PowerPivot Configuration Tool window, 320
- PowerPivot for SharePoint 2013 Configuration tool, 320–322
- PowerPivot service application, creating, 322–323
- PowerPoint Automation Services. *See* PowerPoint Conversion Services
- PowerPoint Broadcast, removal before migration, 153
- PowerPoint Conversion Services, configuring, 246–247
- Power View, 327–329

- pre-check tasks, upgrades and migrations, 150–153
- preferred owners (WFSC), 13
- presentation of data, BI tools, 304–330
  - Excel Services data modeling, 326
  - PerformancePoint, 304–310
  - PowerPivot, 319–326
  - Power View, 327–329
  - Reporting Services, 310–319
- Preserving Context setting, 134
- pre-upgrade health checks, site collections, 184–185
- primary databases
  - log shipping, 16
- primary servers
  - log shipping, 16
- principal databases
  - SQL Server mirroring, 6
- principal servers
  - SQL Server mirroring, 6
- Privacy Settings section (My Site host), 68
- product catalogs
  - WCM (Web Content Management), 113–117
- Product content type, product catalogs, 114
- productivity services, configuring, 220–251
  - Machine Translation Services, 247–250
    - creating service application, 248
    - running services, 251
    - settings, 249–251
  - Microsoft Access, 231–239
    - Access Services, 234–238
    - Access Services 2010, 231–233
  - Microsoft Excel Services, 221–230
    - data connection libraries, 229–230
    - Data Model Settings, 230
    - Trusted Data Providers, 229
    - Trusted File Locations, 225–229
    - User-Defined Function Assemblies, 230
  - Microsoft PowerPoint Conversion Services, 246–247
  - Microsoft Visio Services, 239–242
    - unattended service accounts, 239–240
    - Visio Graphics Service settings, 240–242
  - Microsoft Word Automation Services, 242–245
    - creating service application, 243–244
    - file conversions, 245–246
    - modifying properties, 243–244
- Products Preview Upgrade worksheet, 150
- Profile Database (User Profile service application), 74
- profile properties, incremental profile synchronization, 84

- Project Server Application Service, 335–336
- Project Server, configuring a connection to, 332–339
  - creating web app, 336–339
  - installation, 333–335
  - Project Server Application Service, 335–336
- Project Server service application, creating, 335–336
- Propagate permissions, 272
- Propagate Permissions To All option, 271
- properties
  - channels, 111
  - CompatibilityLevel, 296
  - configuring Word Automation Services, 243–244
  - product catalogs, 115
  - SPSiteUpgradeThrottleSettings, 196
- provisioning
  - subscriptions and alerts, Reporting Services, 313–315
- provisioning Central Administration, Psconfig, 167–169
- Provision Subscriptions and Alerts link (Manage Reporting Services Application page), 319
- Psconfig tool, 167–169
- publishing farms, importing certificates, 256
- publishing service applications
  - federation of services, 260–262
- Publish Service Application dialog box, 261
- Publish-SPServiceApplication -Identity <Service App GUID> command, 261

## Q

- queries, searches, 91–93, 99–101
- query processing component (Search server), 25
- Query Timeout setting, configuring Access Services, 238
- queues, upgrading site collections, 193–198
- Quick Deploy Job settings, 122
- quotas
  - sandboxed solutions, 292–294

## R

- RAID (Redundant Array of Inexpensive Disks), 28
- RC4 encryption, 10
- read-only access for content, configuring, 173–175
- Read Permissions Level section (My Site host), 67

## Receive-SPServiceApplicationConnectionInfo -FarmURL <Published URL> command

- Receive-SPServiceApplicationConnectionInfo -FarmURL <Published URL> command, 262
- Recommendations analysis, 108
- Record Library (record centers), 134
- records
  - centers, 134–135
  - ECM (Enterprise Content Management)
    - disposition, 135–137
    - management, 131–135
    - retention, 135–137
- recovery models, 5
- recovery planning
  - business continuity management, 38–51
  - Central Administration backup, 39–42
  - content recovery, 47–51
  - farm configuration recovery, 44–45
  - non-production environment content
    - refresh, 43–44
  - service application recovery, 46–47
  - SharePoint backup schedules, 38–39, 42
  - SQL Server backup schedules, 43–44
- Recovery Point Objective (RPO), 3
- Recovery Time Objective (RTO), 2–3
- Recycle Threshold setting, configuring Machine Translation Service, 250
- Redirect Users To The Drop Off Library setting, 134
- redundancy
  - storage, 20–21
- Redundant Array of Inexpensive Disks (RAID), 28
- refreshing content
  - non-production environments, 43–44
- relevancy, search results
  - authoritative pages, 93–94
  - custom entities, 89–91
  - query suggestions, 91–93
  - thesaurus, 88–89
- removal of servers, planning upgrades, 166–171
  - moving UPS service, 169–170
  - provisioning Central Administration with Psconfig, 167–169
  - removing services, 169
- Remove-SPSite command, 159
- Remove-SPSiteUpgradeSessionInfo -Identity command, 194
- repairing PowerPivot, 323
- repair mode, pre-upgrade health check, 185
- Repair-SPSite -Identity command, 185
- Replicating Directory Changes permission, 77
- Reporting Services, 310–319
  - configuring, 315–319
  - configuring SharePoint for, 312–313
  - installation on SQL Server, 310–312
  - provisioning subscriptions and alerts, 313–315
- Reporting Services SharePoint Mode for SharePoint 2013, 231
- repositories for documents
  - ECM (Enterprise Content Management), 137–138
- Representational State Transfer (REST) services, 251
- reputation settings, Community Sites, 63–64
- Request Certificate dialog box, 217
- Request-SPUpgradeEvaluationSiteCollection command, 192
- required permissions
  - social workloads, 72–74
- requirements
  - Access Services, 234
  - configuring AlwaysOn High Availability Groups, 12
  - Machine Translation service, 247
  - SQL Server mirroring, 7
- Reset All Pages In This Site To Site Definition Version option, 186
- Reset To Site Definition page, 186
- resolving
  - missing files, after upgrades, 160–162
  - orphan objects, after upgrades, 158–160
- resource groups (WFSC), 13
- Resource Monitor, 180–181
- resource monitoring, Reporting Services, 319
- resource quotas, sandboxed solutions, 292–294
- Resource Sharing, configuring a connection to Exchange, 339–341
- Restore Database option (Microsoft SQL Server Management Studio), 50
- Restore job, log shipping, 16
- Restore-SPFarm -Directory <path> -Item <service name> -RecoveryMethod Overwrite command, 275
- Restore-SPSite command, database recovery, 48–49
- restoring
  - read-only content databases, 176
- REST (Representational State Transfer) services, 251
- results (searches)
  - relevancy, 87–94
    - authoritative pages, 93–94
    - custom entities, 89–91
    - query suggestions, 91–93

- thesaurus, 88–89
- sources, 95–99
- types, 102
- retention of records, ECM (Enterprise Content Management), 135–137
- \$rootCert.Export("Cert")\Set-Content <Root Cert Path>-Encoding byte command, 256
- root certificates, 255
- Round Robin With Health Check setting, configuring global settings, 223
- \$rpa.ResourceMeasures["<resource name>"].Update() command, 294
- \$rpa.ResourceMeasures | Select-Object Name,Resource sPerPoint,AbsoluteLimit command, 293
- \$rpa.Update() command, 294
- RPO (Recovery Point Objective), 3
- RTO (Recovery Time Objective), 2–3
- rules, health checks, 184
- running
  - Machine Translation Services, 250
- Run Site Collection Health Checks page, 185

## S

- Sandboxed Code Service, 296
- sandboxed solution quotas, management and configuration, 292–296
- SAN (Storage Area Network), 11, 28–29
- scheduled backups
  - SharePoint, 38–39, 42
  - SQL Server, 43–44
- scopes, 97
- SCW (Security Configuration Wizard), 163
- search administration component (Search server), 25
- Search Administration service, 168
- Search Clicks analysis, 107
- search properties, product catalogs, 115
- Search Reports analysis, 108
- search schema, 102–107
- Search servers services, 25–26
- Search service application
  - components, 168
  - planning upgrades, 178
- Search service applications, 62
- search workloads, 87–108
  - analytics reports
    - , 107–108
  - end-user experience
    - , 99–102
  - index freshness
    - , 94–95
  - result relevancy
    - , 87–94
  - result sources
    - , 95–99
  - search schema
    - , 102–107
- Secondary Database Settings dialog box, 17
- secondary databases, log shipping, 16
- secondary servers, log shipping, 16
- Secure Store service application
  - configuring, 176
  - upgrading, 178
- Secure Store service application, configuring, 272–278
  - generating master key, 272–274
  - managing master key, 274–275
  - managing target application permissions, 278–279
  - target applications, 276–278
- Secure Store Service Association setting, configuring global settings, 225
- security
  - Business Connectivity Service (BCS), 270–271
  - securityadmin built-in server role, 234
  - Security Configuration Wizard (SCW), 163
  - Security index update group, 103
  - security settings
    - configuring Microsoft Excel Services, 223
  - Security Token Service (STS) certificates, 255
  - Security Trimming Option (My Site host), 67
  - Selectable in Clients permission, 272
  - self-service upgrades, site collections, 190–191
  - Sending to Another Site setting, 134
  - Send Warning E-mail When Usage Per Day Reaches text box, 293
  - ServerInstance parameter (Test-SPContentDatabase command), 154
  - server removal, planning upgrades, 166–171
    - moving UPS service, 169–170
    - provisioning Central Administration with Psconfig, 167–169
    - removing services, 169
- servers
  - work management, 332–342
    - Microsoft Exchange, 339–341
    - Microsoft Project Server, 332–339

## server-to-server authentication, configuring

- server-to-server authentication, configuring, 248
- service application recovery, 46–47
- service applications, 209
  - App Management, 209–218
    - App Store, 210–212
    - DNS entries, 215–216
    - subscriptions, 212–214
    - wildcard certificates, 216–218
  - BCS (Business Connectivity Service), 265–272
    - configuring model security, 270–271
    - importing BCS models, 265–270
  - Business Data Connectivity (BDC), 178–179
    - configuring and upgrading, 176–178
  - federation of services, 253–263
    - consuming service applications, 261–262
    - exchanging trust certificates, 255–256
    - permissions, 258–260
    - planning services to federate, 254–256
    - publishing service applications, 260–262
    - trust management, 256–257
  - Managed Metadata, 178
  - PerformancePoint, 304–305
  - PerformancePoint Services, 178
  - PowerPivot, 322–323
  - productivity services, 220–251
    - Machine Translation Services, 247–250
    - Microsoft Access, 231–239
    - Microsoft Excel Services, 221–230
    - Microsoft PowerPoint Conversion Services, 246–247
    - Microsoft Visio Services, 239–242
    - Microsoft Word Automation Services, 242–245
  - Project Server, 335–336
  - Reporting Services, 312
  - Search, 178
  - Secure Store, 178, 272–278
    - generating master key, 272–274
    - managing master key, 274–275
    - managing target application permissions, 278–279
    - target applications, 276–278
  - upgrading to SharePoint 2013, 151
  - User Profile, 178
  - User Profile Synchronization, 177–178
- service distribution, 24–26
- service instance configuration, 26–27
- session management settings
  - configuring Access Services 2010, 233
  - configuring Microsoft Excel Services, 224
  - trusted file locations, 227
- Session Timeout setting, trusted file locations, 227
- Set Credentials For Secure Store Target Application page, 279
- Set Permissions, 272
- Set-SPSite -Identity "<site collection" -QuotaTemplate "<template name>" command, 293
- settings
  - configuring PowerPivot, 323–326
  - configuring Reporting Services, 315–316
- setup
  - database mirrors, 8–9
  - eDiscovery, 129
- SharePoint 14 folder, 161
- SharePoint 15 folder, 161
- SharePoint Diagnostic Studio, 162–163
- SharePoint Products And Technologies Configuration Wizard Now check box, 334
- SharePoint Products Configuration Wizard, 35, 162–163, 333
- Short Session Timeout setting, trusted file locations, 227
- ShowLocation parameter (Test-SPContentDatabase command), 154
- ShowRowCounts parameter (Test-SPContentDatabase command), 154
- side-by-side folders, 298
- Silverlight, 239
- Simple recovery model, 39
- single-label domain limitations, high-availability planning, 5
- Site Bindings dialog box, 218
- site closure, record management, 132
- site collections
  - App Catalog, 210–212
  - configuring PerformancePoint, 309–310
  - configuring read-only access, 173–174
  - creating for Access Services, 237
  - sandboxed solutions quotas, 292–294
  - upgrading, 184–197
    - analyzing and resolving health check results, 186–189
    - configuring available site collection modes, 188–189
    - evaluation sites, 191–193
    - health check, 184–185
    - queues and throttles, 193–198

- self-service upgrades, 190–191
- site columns, product catalogs, 114
- Site Content Types page, 187
- site mailboxes, record management, 133
- Site Master Page Settings
  - associating master pages and channels, 112
- Site Naming Format option (My Site host), 67
- Site Policy, 132
- site retention, record management, 132–133
- SiteUpgradeThrottleSettings property (SPSite objects), 195
- SkyDrive
  - personal site collections, 69–71
- snapshot option, SQL Server, 191
- Social Distance analysis, 108
- social permissions, 71–74
- Social Tagging Database (User Profile service application), 74
- Social Tags analysis, 107
- social workloads
  - environment, 59–85
    - activity feeds, 77–81
    - Community Sites, 60–65
    - My Sites feature, 65–71
    - social permissions, 71–74
    - synchronized connections, 81–85
    - user profiles, 74–77
- software limitations, high-availability planning, 5
- Solution Daily Resource Update timer, 292
- solution management, 291–302
  - deployment of apps, 299–300
  - deployment of farm solutions, 296–298
  - sandboxed solution quotas, 292–296
  - upgrading apps, 301–302
  - upgrading farm solutions, 298–299
- source variation, creating, 124
- space limitations
  - Central Administration backups, 41
  - high-availability planning, 5
- Specify Export File Name And Location text box, 177
- speed, parallel upgrades, 172
- SPSiteUpgradeThrottleSettings properties, 195, 196
- SQL Native Client 11.0 Configuration, 34
- SQL Server
  - aliases, 34–36
  - backup schedules, 43–44
  - components required to run Access Services, 235
  - configuring read-only access for content databases, 175–176
  - database recovery, 48
  - high-availability planning, 1–22
    - AlwaysOn, 11–15
    - clustering, 10–11
    - gathering requirements, 2–5
    - login replication, 21
    - log shipping, 15–20
    - mirroring, 6–11
    - storage redundancy, 20–21
  - load balancing, 32–34
  - Reporting Services installation, 310–312
  - snapshot option, 191
  - SQL Server 2008 R2 Reporting Services (SSRS), 231
  - SQL Server 2012, availability to Access Services, 234–235
  - SQL Server Agent, 313
  - SQL Server Analysis Services (SSAS), 230
  - SQL Server Client Network Utility, 36
  - SQL Server Installation Center dialog box, 311
  - SQL Server Installation Wizard, 311
  - SQL Server Management Studio
    - configuring log shipping, 16–18
    - setting up a database mirror, 8–9
  - SQL Server PowerPivot System service, 323
  - SSAS (SQL Server Analysis Services), 230
  - SSL Certificates
    - Central Administration Backup tool, 40
  - SSRS (SQL Server 2008 R2 Reporting Services), 231
  - Standard Collection feature, enabling, 61
  - starting
    - Access Services service, 236–237
    - Project Server Application Service, 335–336
    - Word Automation Service service application, 242
  - Start-SPServiceInstance command, 169
  - Stop-SPServiceInstance command, 169
  - Stop When Refresh On Open Fails setting, trusted file locations, 228
  - Storage Area Network (SAN), 11, 28–29
  - storage redundancy, 20–21
  - streams, file conversions, 245
  - \$stsCert.Export("Cert")|Set-Content <STS Cert Path> -Encoding byte command, 256
  - STS (Security Token Service) certificates, 255
  - subscriptions
    - provisioning for Reporting Services, 313–315

## subscriptions, configuring

- subscriptions, configuring, 212–214
- Subscription Services service application, 213
- SubwebCountLimit property (throttling), 196
- Sun Java System Directory Server connections, 82
- Supported File Formats setting, configuring Word Automation Services, 244
- synchronization connections
  - social workloads, 81–85
    - creating, 82–83
    - filters, 83–84
- Synchronization Database (User Profile service application), 74
- Synonym component (synonyms), 88
- synonyms, components, 88
- System Settings page, configuring Reporting Services, 315–316
- System Status Notifications, 191

## T

- tabs
  - Overview (Resource Monitor), 180–181
- Target Application ID text box, 240
- target applications
  - Secure Store service application, 276–279
- technical limitations, high-availability planning, 5
- Templates section (Manage Service Applications page), 233
- test farms, measuring upgrade performances, 179–181
- test mode, pre-upgrade health check, 185
- Test-SPContentDatabase command, 152
- Test-SPSite -Identity command, 185
- thesaurus, searches, 88–89
- throttle settings, properties, 196–197
- throttling, site collection upgrades, 193–198
- time limitations, high-availability planning, 5
- tools
  - BI (business intelligence), 304–330
    - Excel Services data modeling, 326
    - PerformancePoint, 304–310
    - PowerPivot, 319–326
    - Power View, 327–329
    - Reporting Services, 310–319
  - comp, 161
  - detecting and repairing orphan objects, 159

- PowerPivot for SharePoint 2013 Configuration, 320–322
- Psconfig, 167–169
- SharePoint Administrator's Toolkit, 163
- SharePoint Diagnostic Studio, 162–163
- SQL Server Client Network Utility, 36
- Windiff, 161
- topic pages, WCM (Web Content Management), 117–118
- Topology service, 259
- training limitations, high-availability planning, 5
- Transaction Log Backup Settings dialog box, 17
- transaction safety, SQL Server mirroring, 6
- Transact-SQL ScriptDom (Microsoft SQL Server 2012), 235
- translation of content, 125–126
- Translation Processes setting, configuring Machine Translation Service, 250
- trust certificate exchange
  - federation of services, 255–256
- Trust Children setting, trusted file locations, 226
- Trusted Data Provider Description setting, configuring Visio Graphics Service, 242
- Trusted Data Provider ID setting, configuring Visio Graphics Service, 242
- Trusted Data Providers
  - configuring Microsoft Excel Services, 229
  - configuring Visio Graphics Service service application, 242–243
- Trusted Data Provider Type setting, configuring Visio Graphics Service, 242
- trusted data sources
  - configuring PerformancePoint, 308–309
- Trusted File Locations productivity services, 225–229
- Trusted My Site Host Locations, configuring, 68–69
- trust management, federation of services, 256–257
- turning off self-service upgrades, 191
- turning on
  - Foundation Subscription Setting service, 212–214

## U

- ULS (Universal Logging Service), 164
- unattended service accounts
  - configuring PerformancePoint, 306–308

- unattended service accounts, configuring Microsoft Visio Services, 239–240
- Unattended Service Account setting, configuring global settings, 225
- uninstalling, language packets, 176
- Uninstall-SPSolution command, 297
- Universal Logging Service (ULS), 164
- Universal Resource Name (URN), 260
- unused items, removal before migration, 152
- Update-SPSecureStoreApplicationServerKey -Passphrase <passphrase> command, 275
- Update-SPSolution command, 299
- upgrade farms, configuring, 176–178
- upgrades
  - apps, 301–302
    - evaluation of content and customizations, 149–164
      - analyzing content database test results, 153–155
      - configuration issues, 162–163
      - configuring web application authentication, 155–158
      - pre-check tasks, 150–153
      - resolving missing files, 160–162
      - resolving orphan objects, 158–160
  - farm solutions, 298–299
  - planning process, 166–182
    - configuring parallel upgrades, 171–173
    - configuring read-only access for content, 173–175
    - configuring upgrade farms, 176–178
    - installation sequence, 181–182
    - measuring upgrade performance, 179–181
    - removal of servers in rotation, 166–171
  - site collections, 184–197
    - analyzing and resolving health check results, 186–189
    - configuring available site collection modes, 188–189
    - evaluation sites, 191–193
    - health check, 184–185
    - queues and throttles, 193–198
    - self-service upgrades, 190–191
- Upgrade-SPSite command, 194
- Upgrade Status page (Central Administration), 172
- UPS (User Profile Synchronization) service
  - exporting encryption key, 177–178
  - moving, 169–170
- URLS, apps, 214
- URN (Universal Resource Name), 260

- usage analytics, searches, 108
- Usage Counts analysis, 108
- Usage index update group, 103
- UsageStorageLimit property (throttling), 196
- User-Defined Function Assemblies, configuring Microsoft Excel Services, 230
- User Profile Replication Engine 2010, 163
- user profiles, 74–77
- User Profile service application, 61, 74–75
  - social permissions, 72
- User Profile service, configuring for upgrades, 177–178
- User Profile Synchronization service, 76–77
- User Profile Synchronization (UPS) service application
  - exporting encryption key, 177–178
  - moving, 169–170
- Use Tag and Notes permission, 72

## V

- variations
  - WCM (Web Content Management), 122–126
- versions
  - SQL Server, 3–4
- Visio Services, configuring, 239–242
  - unattended service accounts, 239–240
  - Visio Graphics Service settings, 240–242
- visitors, community sites, 62
- Volatile Function Cache Lifetime setting, trusted file locations, 228
- .vsdx extension (Visio files), 239

## W

- \$wa.MigrateUsers(\$true) command, 157
- \$wa.ProvisionGlobally() command, 157
- Warn Of Refresh setting, trusted file locations, 228
- \$wa.Update() command, 156
- WCM (Web Content Management), 110–126
  - channels, 111–113
  - content deployment, 119–122
  - product catalogs, 113–117
  - topic pages, 117–118
  - variations, 122–126
- Web Analytics, stopping before migration, 153
- \$webApp.CompatibilityRange command, 189

## **\$webApp.GrantAccessToProcessIdentity("<domain account>") command**

- \$webApp.GrantAccessToProcessIdentity("<domain account>") command, 221
- web application authentication, configuring, 155–158
- web application folder (Central Administration Backup tool), 40
- web apps, creating
  - Project Server, 336–339
- \$webApp.SiteUpgradeThrottleSettings command, 195
- \$webApp.Update() command, 189
- \$webApp.UpgradeMaintenanceLink command, 189
- \$webApp.UpgradeReminderDelay command, 189
- Web Content Management (WCM), 110–126
  - channels, 111–113
  - content deployment, 119–122
  - product catalogs, 113–117
  - topic pages, 117–118
  - variations, 122–126
- Web Front End (WFE) services, 25
- Web Proxy Server setting, configuring Machine Translation Service, 250
- WFE (Web Front End) services, 25
- wildcard certificates, configuring, 216–218
- wildcard characters
  - record management, 135
- Windiff tool, 161
- windows
  - PowerPivot Configuration Tool, 320
- Windows-based authentication directory service, 81
- Windows Server AppFabric, 80
- Windows Server Failover Clustering (WSFC) infrastructure, 11–13
- witnesses, SQL Server mirroring, 6
- wizards
  - Configuration
    - creating User Profile service application, 74
  - Farm Configuration, 246
  - Microsoft Identity Integration Server Key Management Utility, 177
  - New Zone, 215
  - SCW (Security Configuration Wizard), 163
  - SharePoint Products Configuration, 35, 162–163, 333
  - SQL Server Installation, 311
- Word 97-2003 Document Scanning setting, configuring
- Word Automation Services, 244
- Word Automation Services, configuring, 242–245
  - creating service application, 243–244
  - file conversions, 245–246

- modifying properties, 243–244
- Workbook Cache Location setting, configuring global settings, 224
- workbook cache settings
  - configuring Microsoft Excel Services, 224–225
- Workbook Calculation Mode setting, trusted file locations, 228
- Workbook Properties, trusted file locations, 227
- Workbook URL setting, configuring global settings, 223
- work management, creating and configuring, 332–342
  - Microsoft Exchange, 339–341
  - Microsoft Project Server, 332–339
- Write permission, 77
- WSFC (Windows Server Failover Clustering) infrastructure, 11–13

## **X**

- XML-based file format, Visio, 239

# About the Author



**MICHAEL DOYLE**, MCTS, MCSA, MCSE, and a senior architect at a major company, has been working with SharePoint almost exclusively for a decade. He is a frequent speaker and blogger. Michael is the author of *Microsoft SharePoint 2010: Customizing My Site* and co-author of *Microsoft SharePoint Foundation 2010 Inside Out*.



# What do you think of this book?

We want to hear from you!

To participate in a brief online survey, please visit:

[microsoft.com/learning/booksurvey](https://microsoft.com/learning/booksurvey)

Tell us how well this book meets your needs—what works effectively, and what we can do better. Your feedback will help us continually improve our books and learning resources for you.

Thank you in advance for your input!

**Microsoft**<sup>®</sup>  
Press