

# Configuring Windows Server Hybrid Advanced Services



Exam Ref AZ-801

Orin Thomas

FREE SAMPLE CHAPTER |



# Exam Ref AZ-801 Configuring Windows Server Hybrid Advanced Services

Orin Thomas

# Exam Ref AZ-801 Configuring Windows Server Hybrid Advanced Services

Published with the authorization of Microsoft Corporation by:  
Pearson Education, Inc.

Copyright © 2023 by Orin Thomas.

All rights reserved. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. For information regarding permissions, request forms, and the appropriate contacts within the Pearson Education Global Rights & Permissions Department, please visit [www.pearson.com/permissions](http://www.pearson.com/permissions).

No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained herein.

ISBN-13: 978-0-13-772949-4

ISBN-10: 0-13-772949-9

Library of Congress Control Number: 2022945387

ScoutAutomatedPrintCode

## TRADEMARKS

Microsoft and the trademarks listed at <http://www.microsoft.com> on the "Trademarks" webpage are trademarks of the Microsoft group of companies. All other marks are property of their respective owners.

## WARNING AND DISCLAIMER

Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. The information provided is on an "as is" basis. The author, the publisher, and Microsoft Corporation shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the programs accompanying it.

## SPECIAL SALES

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact our corporate sales department at [corpsales@pearsoned.com](mailto:corpsales@pearsoned.com) or (800) 382-3419.

For government sales inquiries, please contact [governmentsales@pearsoned.com](mailto:governmentsales@pearsoned.com).

For questions about sales outside the U.S., please contact [intlcs@pearson.com](mailto:intlcs@pearson.com).

## CREDITS

EDITOR-IN-CHIEF  
Brett Bartow

EXECUTIVE EDITOR  
Loretta Yates

DEVELOPMENT EDITOR  
Songlin Qiu

MANAGING EDITOR  
Sandra Schroeder

SENIOR PROJECT EDITOR  
Tracey Croom

TECHNICAL EDITOR  
Andrew Warren

COPY EDITOR  
Elizabeth Welch

INDEXER  
Tim Wright

PROOFREADER  
Barbara Mack

EDITORIAL ASSISTANT  
Cindy Teeters

COVER DESIGNER  
Twist Creative, Seattle

COMPOSITOR  
codeMantra

# Pearson's Commitment to Diversity, Equity, and Inclusion

Pearson is dedicated to creating bias-free content that reflects the diversity of all learners. We embrace the many dimensions of diversity, including but not limited to race, ethnicity, gender, socioeconomic status, ability, age, sexual orientation, and religious or political beliefs.

Education is a powerful force for equity and change in our world. It has the potential to deliver opportunities that improve lives and enable economic mobility. As we work with authors to create content for every product and service, we acknowledge our responsibility to demonstrate inclusivity and incorporate diverse scholarship so that everyone can achieve their potential through learning. As the world's leading learning company, we have a duty to help drive change and live up to our purpose to help more people create a better life for themselves and to create a better world.

Our ambition is to purposefully contribute to a world where:

- Everyone has an equitable and lifelong opportunity to succeed through learning.
- Our educational products and services are inclusive and represent the rich diversity of learners.
- Our educational content accurately reflects the histories and experiences of the learners we serve.
- Our educational content prompts deeper discussions with learners and motivates them to expand their own learning (and worldview).

While we work hard to present unbiased content, we want to hear from you about any concerns or needs with this Pearson product so that we can investigate and address them.

- Please contact us with concerns about any potential bias at <https://www.pearson.com/report-bias.html>.



# Contents at a glance

	<i>Introduction</i>	<i>xiii</i>
<b>CHAPTER 1</b>	<b>Secure Windows Server on-premises and hybrid infrastructure</b>	<b>1</b>
<b>CHAPTER 2</b>	<b>Implement and manage Windows Server High Availability</b>	<b>87</b>
<b>CHAPTER 3</b>	<b>Implement disaster recovery</b>	<b>119</b>
<b>CHAPTER 4</b>	<b>Migrate servers and workloads</b>	<b>157</b>
<b>CHAPTER 5</b>	<b>Monitor and troubleshoot Windows Server environments</b>	<b>197</b>
	<i>Index</i>	<i>251</i>



# Contents

## **Introduction**

**xiii**

- Organization of this book*
- Microsoft certifications*
- Quick access to online references*
- Errata, updates, & book support*
- Stay in touch*

## **Chapter 1 Secure Windows Server on-premises and hybrid infrastructure**

**1**

- Skill 1.1: Secure Windows Server operating system .....1
  - Configure and manage exploit protection 2
  - Configure and manage Windows Defender Application Control 5
  - Configure and manage Windows Defender for Endpoint 8
  - Configure and manage Windows Defender Credential Guard 9
  - Configure SmartScreen 11
  - Implement operating system security by using Group Policies 12
  - Configure Windows Server Auditing 21
- Skill 1.2: Secure a hybrid Active Directory infrastructure..... 23
  - Configure password policies 24
  - Enable password block lists 34
  - Manage protected users 34
  - Manage account security on a Read-Only Domain Controller (RODC) 36
  - Harden domain controllers 38
  - Configure authentication policies silos 41
  - Restrict access to domain controllers 43
  - Configure account security 44
  - Manage AD built-in administrative groups 46
  - Manage AD delegation 50
  - Implement and manage Microsoft Defender for Identity 51
  - Implement and manage Azure AD self-service password reset 53



Skill 1.3: Identify and remediate Windows Server security issues by using Azure services .....	55
Monitor on-premises servers and Azure IaaS virtual machines (VMs) by using Microsoft Sentinel	56
Identify and remediate security issues in on-premises servers and Azure IaaS VMs by using Microsoft Defender for Cloud (Azure Security Center)	57
Manage security updates using Windows Server Update Services	60
Manage security updates using Azure Automation Update Management	65
Onboard Windows Server using Azure Automanage	67
Skill 1.4: Secure Windows Server networking .....	68
Manage Windows Defender Firewall	68
Implement connection security rules	72
Implement domain isolation	75
Skill 1.5: Secure Windows Server storage .....	77
Manage Windows BitLocker Drive Encryption (BitLocker)	78
Manage and recover encrypted volumes	82
Enable storage encryption by using Azure Disk Encryption	82
Manage disk encryption keys for IaaS virtual machines	83
Chapter summary .....	85
Thought experiment.....	85
Thought experiment answers .....	86

## **Chapter 2 Implement and manage Windows Server High Availability 87**

Skill 2.1: Implement a Windows Server failover cluster.....	87
Implement a failover cluster on-premises, hybrid, or cloud-only	88
Create a Windows failover cluster	89
Stretch cluster across datacenter or Azure regions	91
Configure storage for failover clustering	93
Modify quorum options	93
Configure network adapters for failover clustering	95

Configure cluster workload options	96
Configure cluster sets	99
Create an Azure witness	100
Configure a floating IP address for the cluster	100
Implement load balancing for the failover cluster	100
Skill 2.2: Manage failover clustering .....	106
Implement Cluster-Aware Updating	106
Recover a failed cluster node	107
Upgrade a node to Windows Server 2022	108
Failover workloads between nodes	109
Manage failover clusters using Windows Admin Center	109
Skill 2.3: Implement and manage Storage Spaces Direct .....	110
Understand Storage Spaces Direct	111
Create failover cluster using Storage Spaces Direct	112
Configure Storage Spaces Direct	113
Upgrade Storage Spaces Direct node	114
Implement networking for Storage Spaces Direct	115
Chapter summary .....	116
Thought experiment .....	117
Thought experiment answers .....	117

## **Chapter 3 Implement disaster recovery 119**

Skill 3.1: Manage backup and recovery for Windows Server .....	119
Understand Azure Backup	120
Manage backups in Azure Recovery Services vault	121
Back up and restore files and folders to Azure Recovery Services vault	124
Install and manage Azure Backup Server	124
Back up and recover using Azure Backup Server	126
Create a backup policy	127
Configure backup for Azure Virtual Machines using the built-in backup agent	128
Restore a VM	137

Skill 3.2: Implement disaster recovery by using Azure Site Recovery . . . .	141
Understand Azure Site Recovery	142
Configure Site Recovery for on-premises VMs	142
Configure a recovery plan	144
Configure Site Recovery for Azure virtual machines	144
Configure Azure Site Recovery networking	145
Configure Azure Site Recovery policy	146
Skill 3.3: Protect virtual machines by using Hyper-V Replica . . . . .	148
Understand Hyper-V Replica	148
Configure Hyper-V hosts for replication	149
Manage Hyper-V Replica servers	150
Configure VM replication	152
Perform a failover	153
Chapter summary . . . . .	154
Thought experiment . . . . .	155
Thought experiment answers . . . . .	155

## **Chapter 4 Migrate servers and workloads 157**

Skill 4.1: Migrate on-premises storage to on-premises servers or Azure . . . . .	157
Transfer data and shares	158
Cut over to a new server by using Storage Migration Service	162
Use Storage Migration Service to migrate to Azure Virtual Machines	164
Migrate to Azure File Shares	164
Skill 4.2: Migrate on-premises servers to Azure. . . . .	166
Migrate by using Azure Migrate	167
Deploy and configure Azure Migrate appliance	168
Migrate VM workloads to Azure IaaS	169
Migrate physical workloads to Azure IaaS	171
Skill 4.3: Migrate workloads from previous versions to Windows Server 2022 . . . . .	172
Understand Windows Server Migration Tools	173

Migrate Internet Information Services (IIS)	174
Migrate Hyper-V hosts	175
Migrate Remote Desktop Services (RDS) host servers	178
Migrate Dynamic Host Configuration Protocol (DHCP)	179
Migrate print servers	182
Skill 4.4: Migrate IIS workloads to Azure	183
Migrate IIS workloads to Azure Web Apps	184
Migrate IIS workloads to containers	186
Skill 4.5: Migrate an AD DS infrastructure to Windows Server 2022	
AD DS	188
Upgrade an existing forest	189
Migrate AD DS objects, including users, groups and Group Policies, using Active Directory Migration Tool	193
Migrate to a new Active Directory forest	194
Chapter summary	195
Thought experiment	196
Thought experiment answers	196

## **Chapter 5 Monitor and troubleshoot Windows Server environments 197**

Skill 5.1: Monitor Windows Server by using Windows Server tools and Azure services	197
Monitor Windows Server by using Performance Monitor	198
Create and configure data collector sets	200
Monitor servers using Windows Admin Center	201
Monitor by using System Insights	202
Manage event logs	202
Deploy Azure Monitor and Log Analytics agents	206
Collect performance counters to Azure	207
Create alerts	208
Monitor Azure Virtual Machines by using Azure Diagnostics extension	208
Monitor Azure Virtual Machines performance by using VM Insights	209

Skill 5.2: Troubleshoot Windows Server on-premises and hybrid networking .....	210
Troubleshoot hybrid network connectivity	210
Troubleshoot Azure VPN	211
Troubleshoot on-premises connectivity	213
Skill 5.3: Troubleshoot Windows Server virtual machines in Azure .....	219
Troubleshoot deployment failures	220
Troubleshoot booting failures	221
Troubleshoot VM performance issues	224
Troubleshoot VM extension issues	225
Troubleshoot disk encryption issues	225
Troubleshoot storage	226
Troubleshoot VM connection issues	227
Skill 5.4: Troubleshoot Active Directory .....	229
Restore objects from AD Recycle Bin	230
Recover Active Directory database using Directory Services Restore Mode	231
Troubleshoot Active Directory replication	238
Recover SYSVOL	242
Troubleshoot hybrid authentication issues	242
Troubleshoot on-premises Active Directory	246
Chapter summary .....	249
Thought experiment .....	250
Thought experiment answers .....	250
 <i>Index</i>	 251

# Introduction

---

The AZ-801 exam deals with advanced topics that require candidates to have an excellent working knowledge of Microsoft Windows Server and Azure Hybrid functionality. The exam covers topics that even experienced Windows Server Hybrid administrators may rarely encounter unless they are consultants who manage hybrid cloud workloads on a regular basis. To be successful in taking this exam, candidates need to understand not only how to secure Windows Server and Active Directory, but also how to manage and configure high availability and disaster recovery, migrate workloads to newer versions of Windows Server and to Azure, as well as monitor and troubleshoot Windows Server workloads across on-premises, hybrid, and cloud infrastructure.

Candidates for this exam are information technology (IT) professionals who want to validate their advanced Windows Server Hybrid administration skills and knowledge. To pass, candidates require a thorough theoretical understanding as well as meaningful practical experience implementing the technologies involved.

This edition of this book covers Windows Server and the AZ-801 exam objectives as of mid-2022. As Windows Server hybrid technologies evolve, so do the AZ-801 exam objectives, so you should check carefully if any changes have occurred since this edition of the book was authored and study accordingly.

This book covers every major topic area found on the exam, but it does not cover every exam question. Only the Microsoft exam team has access to the exam questions, 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 “Need more review?” links you’ll find in the text to find more information and take the time to research and study the topic. Great information is available on Microsoft Docs and Microsoft Learn, and in blogs and forums.

## Organization of this book

---

This book is organized by the “Skills measured” list published for the exam. The “Skills measured” list is available for each exam on the Microsoft Learning website: [microsoft.com/learn](https://microsoft.com/learn). Each chapter in this book corresponds to a major topic area in the list, and the technical tasks in each topic area determine a chapter’s organization. If an exam covers six major topic areas, for example, the book will contain six chapters.

## 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-premises 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 [microsoft.com/learn](https://microsoft.com/learn).

Check back often to see what is new!

## Quick access to online references

---

Throughout this book are addresses to webpages that the author has recommended you visit for more information. Some of these addresses (also known as URLs) can be painstaking to type into a web browser, so we've compiled all of them into a single list that readers of the print edition can refer to while they read.

*Download the list at [MicrosoftPressStore.com/ExamRefAZ801/downloads](https://MicrosoftPressStore.com/ExamRefAZ801/downloads)*

The URLs are organized by chapter and heading. Every time you come across a URL in the book, find the hyperlink in the list to go directly to the webpage.

## Errata, updates, & book support

---

We've made every effort to ensure the accuracy of this book and its companion content. You can access updates to this book—in the form of a list of submitted errata and their related corrections—at:

*[MicrosoftPressStore.com/ExamRefAZ801/errata](https://MicrosoftPressStore.com/ExamRefAZ801/errata)*

If you discover an error that is not already listed, please submit it to us at the same page.

For additional book support and information, please visit [MicrosoftPressStore.com/Support](https://MicrosoftPressStore.com/Support).

Please note that product support for Microsoft software and hardware is not offered through the previous addresses. For help with Microsoft software or hardware, go to <http://support.microsoft.com>.

## Stay in touch

---

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





# About the author

**ORIN THOMAS** is a Principal Hybrid Cloud Advocate at Microsoft and has written more than 3 dozen books for Microsoft Press on topics including Windows Server, Windows Client, Azure, Office 365, System Center, Exchange Server, Security, and SQL Server. He holds a Master's degree in Computing Research, has authored Azure Architecture courses at Pluralsight, and has authored multiple Microsoft Official Curriculum courses on a variety of IT Pro topics. You can follow him on Twitter at <http://twitter.com/orinthomas> or find out more at <https://aka.ms/orin>.



# Implement and manage Windows Server High Availability

The primary high-availability technology available in Windows Server is failover clustering. Clustering allows you to ensure that your workloads remain available if server hardware or even a site fails. You can configure Windows Server failover clustering for on-premises, hybrid, and Azure hosted workloads. Understanding Windows Server high availability involves knowing the requirements for configuring a failover cluster and how to manage that cluster. This chapter also addresses the deployment, configuration, and management of the highly available Windows Server storage technology named Storage Spaces Direct.

## Skills covered in this chapter:

- Skill 2.1: Implement a Windows Server failover cluster
- Skill 2.2: Manage failover clustering
- Skill 2.3: Implement and manage Storage Spaces Direct

## Skill 2.1: Implement a Windows Server failover cluster

---

This objective deals with managing Windows Server failover clusters that are deployed on-premises, in a hybrid configuration, and in Azure. To master this objective, you'll need to understand Windows Server failover clustering requirements, possible cluster configurations, and the types of cluster workloads that you can deploy.

### This skill covers how to:

- Implement a failover cluster on-premises, hybrid, or cloud-only
- Create a Windows failover cluster
- Stretch cluster across datacenter or Azure regions
- Configure storage for failover clustering

- Modify quorum options
- Configure network adapters for failover clustering
- Configure cluster workload options
- Configure cluster sets
- Create an Azure witness
- Configure a floating IP address for the cluster
- Implement load balancing for the failover cluster

## Implement a failover cluster on-premises, hybrid, or cloud-only

The primary high-availability technology available in Windows Server is failover clustering. Failover clustering is a stateful, high-availability solution that allows an application or service to remain available to clients if a host server fails. You can use failover clustering to provide high availability to applications such as SQL Server and to scale out file servers and virtual machines (VMs). With clustering, you can ensure that your workloads remain available if server hardware or even a site fails.

Failover clustering is supported in both the Standard and Datacenter editions of Windows Server. In some earlier versions of the Windows Server operating system, you gained access to failover clustering only if you used the Enterprise edition. Windows Server supports up to 64 nodes in a failover cluster.

Generally, all servers in a cluster should run either a similar hardware configuration or should be similarly provisioned virtual machines. You should also use the same edition and installation option. For example, you should aim to have cluster nodes that run either the full GUI or the Server Core version of Windows Server, but you should avoid having cluster nodes that have a mix of computers running Server Core and the full GUI version. Avoiding this mix ensures that you use a similar update routine. A similar update routine is more difficult to maintain when you use different versions of Windows Server.

You should use the Datacenter edition of Windows Server when building clusters that host Hyper-V virtual machines because the virtual machine licensing scheme available with this edition provides the most VM licenses.

To be fully supported by Microsoft, cluster hardware should meet the Certified for Windows Server logo requirement. An easy way of accomplishing this is to purchase and deploy Azure Stack HCI, a prebuilt hyper-converged Windows Server installation available from select vendors. Even though it is called Azure Stack HCI and sounds as though it is far more of a cloud-based solution, it's primarily just an optimized Windows Server deployment on a certified configuration with all the relevant clustering and "Software-Defined Datacenter" features lit up.

## **NEED MORE REVIEW? FAILOVER CLUSTERING**

You can learn more about Windows Server failover clusters at <https://docs.microsoft.com/en-us/windows-server/failover-clustering/failover-clustering-overview>.

## **Create a Windows failover cluster**

Windows Server failover clusters have the following prerequisites:

- All cluster nodes should be running the same version and edition of Windows Server.
- You can add clustered storage during or after cluster creation.
- All servers in the cluster that are located in the same site should be members of the same Active Directory (AD) domain. If configuring a stretch cluster, nodes must be members of the same forest.
- The account used to create the cluster must be a domain user who has local administrator rights on all servers that will function as cluster nodes.
- The account used to create the cluster requires the Create Computer Objects permission in the organizational unit (OU) or container that will host the cluster-related Active Directory objects.

Recommended practice is to place the computer accounts for cluster nodes in the same OU and to use separate OUs for each cluster. Some organizations create child OUs for each separate cluster in a specially created parent Cluster OU.

You install failover clustering by installing the Failover Clustering feature, performing initial cluster configuration, running the cluster validation process, and then performing cluster creation. You can use Windows Admin Center, PowerShell, or the Server Manager console to perform these tasks. Once the cluster is deployed, you can manage your clusters using the Failover Clustering Remote Server Administration Tools (RSAT), PowerShell, or Windows Admin Center.

You can install the Failover Clustering feature and its associated PowerShell cmdlets on a node using the following PowerShell command:

```
Install-WindowsFeature -Name Failover-Clustering -IncludeManagementTools
```

## **Validating cluster configuration**

Cluster validation performs a check of a cluster's current or proposed configuration and allows you to determine whether you have the necessary pieces in place to create a cluster prior to attempting to perform this task. Although you can skip validation, recommended practice is to go through the process. This is because even though you may have created numerous clusters in the past it doesn't mean that the next time you go to create a cluster you accidentally overlook some small but critical detail.

The period prior to cluster deployment is not the only time that you can perform cluster validation. You should rerun cluster validation whenever you change or update a significant component of the cluster. This includes adding nodes, modifying storage hardware, updating

network adapters, updating firmware or drivers for network adapters, and updating multipathing software. Cluster validation performs tests in six categories:

- **Inventory** Inventory tests determine if the hardware, software, networking, and storage configuration support the deployment of a cluster.
- **Network** A detailed set of tests to validate cluster network settings.
- **Storage** A detailed set of tests to analyze shared cluster storage.
- **Storage Spaces Direct (S2D)** A detailed set of tests to analyze S2D configuration.
- **System Configuration** A set of tests on the current system configuration.
- **Cluster Configuration** This category of test only executes on deployed clusters to verify that best practices are being followed (for example, using multiple network adapters connected to different networks).

You can perform cluster validation from the Failover Cluster Management Tools that are part of the Remote Server Administration Tools, using Windows Admin Center to connect to an existing cluster, or by running the `Test-Cluster` PowerShell cmdlet.

#### **NEED MORE REVIEW? VALIDATE CLUSTERS**

You can learn more about validating Windows Server failover clusters at <https://techcommunity.microsoft.com/t5/failover-clustering/validating-a-cluster-with-zero-downtime/ba-p/371685>.

## **Prestage cluster computer objects**

During the cluster creation process, a computer object is created in Active Directory Domain Services (AD DS) that matches the cluster name. This AD DS object is called the *cluster name object*. As mentioned earlier in the chapter, the domain user account used to create the cluster must have the Create Computer Objects permission in order to create this object. It's possible to have an appropriately permissioned account pre-create a cluster name object. When this is done, the account used to then create the cluster using the constituent nodes does not require the Create Computer Objects permission.

## **Workgroup clusters**

*Workgroup clusters* are a special type of cluster where cluster nodes are not members of an Active Directory domain. Workgroup clusters are also known as Active Directory detached clusters. The following workloads are supported for workgroup clusters:

- **SQL Server** When deploying SQL Server on a workgroup cluster, you should use SQL Server Authentication for databases and SQL Server Always On Availability Groups.
- **File Server** A supported but not recommended configuration as Kerberos will not be available as an authentication protocol for SMB traffic.
- **Hyper-V** A supported but not recommended configuration. Hyper-V live migration is not supported, though it is possible to perform quick migration.

When creating a workgroup cluster, you first need to create a special account on all nodes that will participate in the cluster that has the following properties:

- The special account must have the same username and password on all cluster nodes.
- The special account must be added to the local Administrators group on each cluster node.
- The primary DNS suffix on each cluster node must be configured with the same value.
- When creating the cluster, ensure that the `AdministrativeAccessPoint` parameter when using the `New-Cluster` cmdlet is set to DNS. Ensure that the cluster name is present in the appropriate DNS zone, which depends on the primary DNS suffix, when running this command.
- You will need to run the following PowerShell command on each node to configure the `LocalAccountTokenFilterPolicy` registry setting to 1:

```
new-itemproperty -path HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System -Name LocalAccountTokenFilterPolicy -Value 1
```

To create a workgroup cluster, use the `New-Cluster` cmdlet with the `name` parameter listing the cluster name, the `node` parameters listing the nodes that you wish to join to the cluster where the nodes have been configured according to the prerequisites, and the `AdministrativeAccessPoint` parameter configured for DNS. For example, to create a new workgroup cluster named `workgrp1st` with member nodes `node1` and `node2`, run the following command on one of the nodes:

```
New-Cluster -name workgrp1st -node node1,node2 -AdministrativeAccessPoint DNS
```

## Stretch cluster across datacenter or Azure regions

Failover clusters can span multiple sites. From the perspective of a hybrid environment, site spanning can include spanning two separate on-premises locations, an on-premises location and an Azure datacenter, or having cluster nodes hosted in different Azure regions. When configuring a cluster that spans two sites, you should consider the following:

- Ensure that there are an equal number of nodes in each site.
- Allow each node to have a vote.
- Enable dynamic quorum. *Dynamic quorum* allows quorum to be recalculated when individual nodes leave the cluster one at a time. Dynamic quorum is enabled by default on Windows Server failover clusters.
- Use a file share witness. You should host the file share witness on a third site that has separate connectivity to the two sites that host the cluster nodes. When configured in this manner, the cluster retains quorum if one of the sites is lost. An alternative to a file share witness is an Azure Cloud Witness.



If you only have two sites and are unable to place a file share witness in an independent third site, you can manually edit the cluster configuration to reassign votes so that the cluster recalculates quorum.

Manually reassigning votes is also useful to avoid split-brain scenarios. *Split-brain scenarios* occur when a failure occurs in a multisite cluster and when both sides of the cluster believe they have quorum. Split-brain scenarios cause challenges when connectivity is restored and make it necessary to restart servers on one side of the multisite cluster to resolve the issue. You can manually reassign votes so that one side always retains quorum if intersite connectivity is lost. For example, by setting the Melbourne site with two votes and the Sydney site with one vote, the Melbourne site always retains quorum if intersite connectivity is lost.

You can use Storage Replica to enable stretch clusters with shared storage. Stretch clusters that use shared storage have the following requirements:

- Cluster nodes are all members of the same AD DS forest.
- Firewall rules allow ICMP, SMB (ports 445 and 5445 for SMB direct), and WS-MAN (port 5985) bidirectional traffic between all nodes that participate in the cluster.
- They must have two sets of shared storage that support persistent reservation. Each storage set must be able to support the creation of two virtual disks. One will be used for replicated data and the other for logs. These disks need to have the following properties:
  - They must be initialized as GUID Partition Table (GPT) and not Master Boot Record (MBR).
  - Data volumes must be of identical size and use the same sector size.
  - Log volumes must be of identical size and use the same sector size.
- Replicated storage cannot be located on the drive containing the Windows Server operating system.
- Premium SSD must be used for cluster nodes hosted as infrastructure-as-a-service (IaaS) VMs in Azure.
- Ensure that there is less than 5-millisecond round-trip latency if synchronous replication is being used. If asynchronous replication is being used, this requirement does not need to be met.
- Storage Replica–configured stretch clusters can use Storage Replica technology to replicate shared cluster storage between locations.

#### **NEED MORE REVIEW? STRETCH CLUSTER REPLICATION**

You can learn more about stretch cluster replication at <https://docs.microsoft.com/en-us/windows-server/storage/storage-replica/stretch-cluster-replication-using-shared-storage>.

## Configure storage for failover clustering

Storage for Windows Server failover clusters needs to be accessible to each node in the cluster. You can use serial-attached SCSI (SAS), iSCSI, Fibre Channel, or Fibre Channel over Ethernet (FCoE) to host shared storage for a Windows Server failover cluster.

You should configure disks used for failover clustering as follows:

- Volumes should be formatted using NTFS or ReFS.
- Use Master Boot Record (MBR) or GUID Partition Table (GPT).
- Avoid allowing different clusters access to the same storage device. This can be accomplished through LUN masking or zoning.
- Any multipath solution must be based on Microsoft Multipath I/O (MPIO).

Cluster Shared Volumes (CSV) is a technology that allows multiple cluster nodes to have concurrent access to a single physical or virtual storage device, also termed a logical unit number (LUN). CSV allows you to have virtual machines on the same shared storage run on different cluster nodes. CSV also has the following benefits:

- Support for scale-out file servers
- Support for BitLocker volume encryption
- SMB 3.0 and higher support
- Integration with Storage Spaces
- Online volume scan and repair

You can enable CSV only after you create a failover cluster and you have provided the shared storage available to each node that will be available to the CSV.

### **NEED MORE REVIEW? CLUSTER SHARED VOLUMES**

You can learn more about Cluster Shared Volumes at <https://docs.microsoft.com/en-us/windows-server/failover-clustering/failover-cluster-csvs>.

## Modify quorum options

A *cluster quorum mode* determines how many nodes and witnesses must fail before the cluster is in a failed state. Nodes are servers that participate in the cluster. Witnesses can be stored on shared storage, on file shares, in Windows Server, and even on a USB drive attached to a network switch; shared storage is the preferred method.

For unknown reasons, some people use Distributed File System (DFS) shares as file share witnesses when setting up their failover clusters. To stop this type of shenanigan from occurring in the future, Microsoft has configured Windows Server failover clustering so that it explicitly blocks the use of DFS namespaces when configuring a file share witness.

Microsoft recommends that you configure a cluster so that an odd number of total votes be spread across member nodes and the witness. This limits the chance of a tie during a quorum vote.

There are four cluster quorum modes:

- **Node Majority** This cluster quorum mode is recommended for clusters that have an odd number of nodes. When this quorum type is set, the cluster retains quorum when the number of available nodes exceeds the number of failed nodes. For example, if a cluster has five nodes and three are available, quorum is retained.
- **Node and Disk Majority** This cluster quorum mode is recommended when the cluster has an even number of nodes. A disk witness hosted on a shared storage disk, such as iSCSI or Fibre Channel, that is accessible to cluster nodes has a vote when determining quorum, as do the quorum nodes. The cluster retains quorum as long as the majority of voting entities remain online. For example, if you have a four-node cluster and a witness disk, a combination of three of those entities needs to remain online for the cluster to retain quorum. The cluster retains quorum if three nodes are online or if two nodes and the witness disk are online.
- **Node and File Share Majority** This configuration is similar to the Node and Disk Majority configuration, but the quorum is stored on a network share rather than on a shared storage disk. It is suitable for similar configurations to Node and Disk Majority. This method is not as reliable as Node and Disk Majority because file shares generally do not have the redundancy features of shared storage.
- **No Majority: Disk Only** This model can be used with clusters that have an odd number of nodes. It is only recommended for testing environments because the disk hosting the witness functions as a single point of failure. When you choose this model, as long as the disk hosting the witness and one node remain available, the cluster retains quorum. If the disk hosting the witness fails, quorum is lost, even if all the other nodes are available.

When you create a cluster, the cluster quorum is automatically configured for you. You might want to alter the quorum mode, however, if you change the number of nodes in your cluster. For example, you might want to alter the quorum mode if you change from a four-node to a five-node cluster. When you change the cluster quorum configuration, the Failover Cluster Manager provides you with a recommended configuration, but you can choose to override that configuration if you want.

You can also perform advanced quorum configuration to specify what nodes can participate in the quorum vote, which you can set on the Select Voting Configuration page of the Configure Cluster Quorum Wizard. When you do this, only the selected nodes' votes are used to calculate quorum. Also, it's possible that fewer nodes would need to fail to cause a cluster to fail than would otherwise be the case if all nodes participated in the quorum vote. This can be useful when configuring how multisite clusters calculate quorum when the connection between sites fails.

#### **NEED MORE REVIEW? CLUSTER QUORUM**

You can learn more about cluster quorum at <https://docs.microsoft.com/en-us/windows-server/failover-clustering/manage-cluster-quorum>.

## Configure network adapters for failover clustering

While you can create failover clusters with nodes that have a single network adapter, best practice is to have separate networks and network adapters for the following:

- A connection for cluster nodes to shared storage
- A private network for internal cluster communication
- A public network that clients use to access services hosted on the cluster

In scenarios where high availability is critical, you might have multiple redundant networks connected through several separate switches. If you have a cluster where everything is connected through one piece of network hardware, you can almost guarantee that piece of network hardware is the first thing that fails.

Failover clustering only supports IPv4- and IPv6-based protocols. You can use either IPv4 or IPv6 addresses that are dynamically or statically assigned, but you should not use a mix of dynamically and statically assigned IP addresses for nodes that are members of the same cluster. If you use a mixture of dynamically and statically assigned IP addresses, the Validate A Configuration Wizard generates an error.

Even if the Cluster Validation Wizard only gives you warnings when you perform the test, you cannot create a failover cluster unless each node is configured with a default gateway. The default gateway doesn't have to be a host that exists, but if you're having trouble in your virtual machine lab with creating a failover cluster, go back and check whether you've configured a default gateway for each node.

Ensure that TCP and UDP port 3343 is open on firewalls between cluster nodes. This port is used for cluster communication, and if communication on this port is disrupted, a node may appear to be in a failed state. Although it's possible to have single adapters and have cluster and client communication occur over the same adapters and networks, production deployments should use separate adapters and networks for cluster communication.

If there are multiple paths to physical storage on a server, you will need to enable multipath I/O (MPIO). Enabling MPIO aggregates multiple paths to a physical disk into a single logical path for data access. Enabling MPIO improves resiliency to failure. You should enable MPIO on all nodes that will participate in the cluster. You can enable MPIO with the following PowerShell command:

```
Install-WindowsFeature -ComputerName Node1 -Name MultiPath-IO
```

### **NEED MORE REVIEW? FAILOVER CLUSTER NETWORKING BASICS**

You can learn more about failover cluster networking basics at <https://techcommunity.microsoft.com/t5/itops-talk-blog/failover-clustering-networking-basics-and-fundamentals/ba-p/1472460>.

## Configure cluster workload options

Cluster workload options include workload startup priority, preferred owners, and failover settings. Configuring startup priority determines when the workload becomes available after a cluster disruption. You can choose between High, Medium, Low, and No Auto Start.

Workload failover settings allow you to configure the following:

- Maximum failures in a period that will trigger failover
- Period over which these failures are measured

Cluster preference settings allow you to configure the preferred owner for a specific cluster role, and you can also configure different preferred owners for different cluster roles. Where possible, the role is hosted on the preferred owner. You can configure a list of preferred owners so that if the most preferred owner isn't available, the next preferred owner hosts the role. Workloads will start on the preferred owner at the top of the list and will fail over to other nodes in list order unless specific failback settings are configured. For example, if you have configured the list as Node 4, Node 3, Node 2, and Node 1 if the workload is currently hosted on Node 3 and this node fails, the workload will fail over to Node 2 if it is available and then Node 1 before attempting to fail over to Node 4. You configure a role-preferred owner in the role's Properties dialog box. You can stop a workload from failing over or being moved to a specific node by ensuring that the node is not present on the workload's list of possible owners. You can configure this list using the `Set-ClusterOwnerNode` PowerShell cmdlet.

You configure whether the clustered role fails back to the preferred owner on the Failover tab of the cluster role's Properties dialog box. When configuring failback, you need to:

- Determine whether you want to prevent failback
- Determine whether you want to have failback occur automatically as soon as the preferred owner is in a healthy state
- Configure failback to occur within a certain number of hours of the preferred owner returning to a healthy state

Node quarantine settings allow you to configure a node so that it is unable to rejoin the cluster if the node fails a certain number of times within a certain period. Configuring node quarantine blocks workloads from being placed back on a quarantined node until the reason for the repeated failure can be dealt with by a server administrator.

### **NEED MORE REVIEW? PREFERRED OWNERS**

You can learn more about cluster preferred owners at <https://techcommunity.microsoft.com/t5/failover-clustering/preferred-owners-in-a-cluster/ba-p/371290>.

## File Server failover clusters

The traditional File Server cluster role allows one node in the cluster to serve files from a highly available file share that is hosted on cluster storage. In the event that the node that services client requests fails, the role fails over to another node and clients accessing files will use the

new node to perform those operations. Other than increasing resiliency of the workload, there is no performance benefit to increasing the number of nodes for a general-purpose file server workload.

## Scale-Out File Servers

A Scale-Out File Server (SoFS) is a new high-availability technology that allows you to share a single folder from multiple nodes of the same cluster. You can use SoFS to deploy file shares that are continuously available for file-based server application storage. This storage is suitable for hosting Hyper-V virtual machine files or Microsoft SQL Server databases with a level of reliability, availability, manageability, and performance that equates to what is possible with a storage area network.

Benefits of an SoFS deployment include:

- **Active-Active file share topology** SoFS allows the same folder to be accessed from multiple nodes of the same cluster. An SoFS file share remains online should one or more cluster nodes fail or be taken down for planned maintenance.
- **Scalable bandwidth** You can respond to a requirement for increased bandwidth by adding nodes.
- **Automatic rebalancing of clients** SMB client connects are tracked on a per-file share basis, with clients being redirected to the cluster node that has the best access to the storage device used by the file share.
- **CHKDSK with zero downtime** The Cluster Shared Volume File System, used with SoFS, allows CHKDSK operations to occur without affecting applications that have open handles on the file system.

You should consider SoFS file shares for the following scenarios:

- Storing Hyper-V configuration and live virtual disks
- Storing live SQL Server database files
- Storing shared IIS configuration data

SoFS has the following requirements:

- The storage configuration must be explicitly supported by failover clustering. This means that you must be able to successfully run the Cluster Validation Wizard before adding an SoFS.
- SoFS requires Cluster Shared Volumes.

Windows Server 2019 and Windows Server 2022 support an SoFS role called the Infrastructure File Server. An infrastructure SoFS uses a single namespace share for the Cluster Shared Volume drive. The benefit of the Infrastructure File Server role is that it allows the Hyper-V host to communicate using guaranteed continuous availability to the Infrastructure SoFS SMB server. A failover cluster can only support a single infrastructure SoFS instance. To create an infrastructure SoFS, run the following PowerShell command:

```
Add-ClusterScaleOutFileServerRole -Cluster ClusterName -Infrastructure -Name InfrastructureSoFSName
```

## Virtual machine failover clustering

One of the most common uses for failover clusters is hosting virtual machines. By deploying a workload such as SQL Server or Exchange on a highly available virtual machine, you can achieve high availability without the need for the application to be aware that it is now highly available. The virtual machine functions normally, provides services to clients on the network, and can switch between cluster nodes as necessary in the event that the individual cluster node hosting it requires maintenance or experiences some sort of failure. Building a Hyper-V failover cluster first involves creating a failover cluster and then adding the Hyper-V role to each node of the cluster.

You should use Cluster Shared Volumes to store virtual machines on a Hyper-V cluster because CSV allows multiple cluster nodes to manage a single shared storage device. This allows you to put multiple virtual machines on the same shared storage device but have those virtual machines hosted by different nodes in the failover cluster. Cluster Shared Volumes are mapped under the `C:\ClusterStorage` folder on cluster nodes.

When creating a new virtual machine on a failover cluster, first select which cluster node hosts the virtual machine. When creating a highly available virtual machine, specify the Cluster Shared Volume path as the location to store the virtual machine. If you have an existing machine that you want to make highly available, you can move the virtual machine to this path. As an alternative, you have the option to specify an SMB 3.0 file share as the storage location for the highly available virtual machine. Whether to select a Cluster Shared Volume or an SMB 3.0 file share depends on your organization's storage configuration.

After the virtual machine is created, you can control it by using the Failover Cluster Manager console. The Move option in the Actions pane allows you to select the cluster node to which you want to move the virtual machine.

In production environments, you should ensure that each Hyper-V host has an identical hardware configuration. However, in development environments, this is not always possible. If different processor types are used—for example, an Intel processor on one node and an AMD processor on another—you might have to perform a quick migration. Quick migration allows migration between nodes but does cause a disruption in client connectivity. You can allow migration between Hyper-V nodes with different processor types or versions by enabling the processor compatibility setting on the virtual machine.

VM Monitoring is a failover cluster feature that allows you to monitor the health state of applications that are hosted on a guest virtual machine. This monitoring allows the cluster to take remediation actions in the event that the application or service fails. You can configure VM Monitoring for any Windows service or event log event that occurs on a guest VM. To use VM Monitoring, you need to enable the Virtual Machine Monitoring rule on the guest VM. You can configure monitoring on a VM using the Failover Cluster Manager or the `Add-ClusterVMMonitoredItem` cmdlet. You can configure whether the failure of an application or service triggers a guest VM restart on the current node or the guest VM to fail over to another node by configuring the failover properties of the virtual machine.

## Configure cluster sets

Cluster sets are new features available in Windows Server 2019 and later, and they allow clusters to be *loosely federated*, allowing you to migrate workloads between clusters with minimal downtime. When you implement cluster sets, you can combine smaller clusters into larger virtual clusters. These virtual clusters support virtual machine fluidity and a unified storage namespace, meaning that virtual machines can be moved between clusters in a cluster set as easily as they are moved between nodes in a traditional failover cluster. While virtual machines can be live-migrated between clusters, Windows Server cluster sets do not allow virtual machines to be configured to automatically fail over between clusters in a cluster set.

Only clusters running at the Windows Server cluster functional level 10 or higher can participate in cluster sets. This means that clusters can have nodes running the Windows Server 2019 and Windows Server 2022 operating systems but that you cannot join clusters running Windows Server 2016 or earlier Windows Server operating systems to a cluster set. All clusters in a cluster set must be members of the same Active Directory forest. If you are going to perform live migration of virtual machines between clusters in a cluster set, you must ensure that the nodes share the same processor architecture.

Cluster sets improve the failover cluster life cycle. Rather than adding and removing nodes from a single cluster when the node hardware is to be retired, you can add a new cluster to the cluster set, migrate workloads across from the original cluster to the new cluster, and then retire that original cluster. Cluster sets are currently supported by Microsoft for up to 64 cluster nodes, which is the same number of nodes supported in an individual failover cluster. That being said, there is no specific limit to the number of nodes that may exist within a cluster set, so going beyond 64 nodes in a cluster set is possible should you want to try it.

Cluster sets consist of a management cluster and member clusters. The management cluster is the cluster set that holds the management role for the cluster set and also hosts the unified storage namespace for the cluster set Scale-Out File Server. The management cluster does not host workloads, such as virtual machines, and its role is to manage the relationship between other clusters in the cluster set and to host the storage namespace. The new role that the management cluster hosts is termed the *Cluster Set Master (CS-Master)*. Member clusters hold the Cluster Set Worker, or CS-Worker, role. The namespace for the cluster set is hosted by a Scale-Out File Server cluster role running on the management cluster. When creating the Scale-Out File Server role for the management cluster, you use the `Add-ClusterScaleoutFileServerRole` PowerShell cmdlet with the `Infrastructure` parameter.

Cluster sets support the configuration of fault domains and availability sets. Fault domains are collections of hardware and software where a single fault, such as a power outage, causes a set of cluster nodes to fail. If you are using stretch clusters to span datacenters, each datacenter could be configured as a separate fault domain. Availability sets allow you to configure workload redundancy across fault domains.



### **NEED MORE REVIEW? CLUSTER SETS**

You can learn more about cluster sets at <https://docs.microsoft.com/en-us/azure-stack/hci/deploy/cluster-set>.

## **Create an Azure witness**

A *Cloud Witness* has the Cluster Witness role hosted in Azure rather than at a third site. A Cloud Witness is suitable for multisite clusters. To configure a Cloud Witness, create a storage account in Azure, copy the storage access keys, note the endpoint URL links, and then use these links with the Configure Cluster Quorum Settings Wizard and specify a Cloud Witness.

A Cloud Witness should be blob storage in a general-purpose standard performance storage account. The Cloud Witness should use storage access keys for authentication rather than shared access signatures or Azure AD system–assigned managed identity. You should configure the Cloud Witness storage account to use locally redundant storage if the cluster is on-premises or hosted in Azure using a single availability zone. If the cluster uses multiple availability zones, you should choose zone-redundant storage.

### **NEED MORE REVIEW? CLOUD WITNESS**

You can learn more about Cloud Witness at <https://docs.microsoft.com/en-us/windows-server/failover-clustering/deploy-cloud-witness>.

## **Configure a floating IP address for the cluster**

Windows Server clusters allow you to define an IP address as a cluster resource and have the IP address fail over between cluster nodes. To do this, you need support for the following:

- Support for dynamic registration and deregistration of IP addresses (DHCP)
- Ability to update network address translation caches of hosts attached to the subnet on which the IP address resides

DHCP and ARP are likely to automatically handle these elements of floating IP address configuration as this technology has been available since the release of Windows NT 4.

### **NEED MORE REVIEW? CLUSTER FLOATING IP ADDRESS**

You can learn more about cluster floating IP address at <https://docs.microsoft.com/en-us/troubleshoot/windows-server/high-availability/cluster-information-ip-address-failover>.

## **Implement load balancing for the failover cluster**

There are two primary methods for load-balancing cluster workloads. The first is used to ensure that a Hyper-V cluster that hosts VMs balances VM workloads equitably across the available nodes. This ensures that you don't end up with a situation where one node has a consistent 95 percent processor utilization due to VM workloads while other nodes in the

# Index

## A

- account management
  - inactive accounts, 32
  - least privilege, 50
  - locked-out accounts, 31
  - lockout settings, 29
  - nonexpiring passwords, 30–31
  - PAWs (Privileged Access Workstations), 43–44
  - Protected Users group, 34–35
  - RBAC (role-based access control), 50–51
  - security, 44–45
- AD DS (Active Directory Domain Services), 23–24.
- See also* account management; Defender for Identity
  - account management
    - inactive accounts, 32
    - least privilege, 50
    - locked-out accounts, 31
    - lockout settings, 29
    - nonexpiring passwords, 30–31
    - PAWs (Privileged Access Workstations), 43–44
    - Protected Users group, 34–35
    - RBAC (role-based access control), 50–51
    - security, 44–45
  - authoritative restore, 234–236
  - Azure file shares, authentication, 164–166
  - built-in administrative groups
    - Administrators, 46–47
    - Domain Admins, 46
    - Enterprise Admins, 46
    - Schema Admins, 47
    - security groups, 47–49
  - cluster name object, 90
  - conflict resolution, 239–240
  - database optimization, 248
  - DNS, troubleshooting, 246–247
  - DSRM (Directory Services Restore Mode), 231–232
  - hybrid authentication issues, troubleshooting, 242
    - identity synchronization, 242
    - pass-through, 244–245
    - password hash synchronization, 243–244
    - Seamless SSO, 245–246
  - KCC (Knowledge Consistency Checker), 239
  - metadata cleanup, 248–249
  - migrating to Windows Server 2022, 188–189
    - demote existing domain controllers, 192
    - DNS migration, 191–192
    - migrate AD DS objects using ADMT, 193–194
    - migrate to a new AD forest, 194–195
    - transfer FSMO role holder, 190–191
    - upgrade an existing forest, 189
  - nonauthoritative restore, 236
  - password policies, 24, 25
    - account lockout settings, 29
    - balanced, 28
    - delegating password setting permissions, 25–26
    - determining password settings, 28
    - fine-grained, 26–27
    - locked-out accounts, 31
    - nonexpiring passwords, 30–31
    - PSO (Password Settings Object), 27–28
  - replication
    - monitoring, 241
    - multimaster, 239
    - RODC, 240–241
    - troubleshooting, 238
  - restoring objects from AD recycle bin, 230–231
  - RODC partial attribute set, 37–38
  - self-service password reset, 53–55
  - snapshots, 236–237
  - store and forward replication, 239
  - SYSVOL, recovering, 242
  - tombstone lifetime, 232–234
- adaptive application controls, 59
- adaptive network hardening, 60
- Add-ClusterNode cmdlet, 107
- Add-ClusterScaleOutFileServerRole cmdlet, 97
- Add-ClusterScaleoutFileServerRole cmdlet, 99

## Add-ClusterVMMonitoredItem cmdlet

- Add-ClusterVMMonitoredItem cmdlet, 98
  - Add-MpPreference cmdlet, 5, 8
  - Administrators group, 46–47
  - ADMT (Active Directory Migration Tool), 193–194
  - alerts
    - Azure Monitor, 208
    - Performance Monitor, 199
    - RSV (Azure Recovery Services vault), 123
  - architecture, Defender for Identity, 52
  - ASR (Azure Site Recovery), 142
    - recovery plans, 144
    - site recovery
      - for Azure VMs, 144–145
      - networking, 145–146
      - policy, 146–147
      - for on-premises VMs, 142–144
  - AT command, 40
  - audit policy(ies), 21
    - auditpol.exe, 22–23
    - expression-based, 21–22
    - file and folder, 22
    - settings, 21
  - authentication
    - exemptions, 74
    - identity synchronization, troubleshooting, 242–243
    - pass-through, troubleshooting, 244–245
    - password hash synchronization, troubleshooting, 243–244
    - policy silos, 41–42
    - Seamless SSO, troubleshooting, 245–246
  - authoritative restore, 234–236
  - automatic approval rules, WSUS, 64–65
  - Azure AD. *See also* AD DS (Active Directory Domain Services)
    - Password Protection, 34
    - self-service password reset, 53–55
  - Azure Automanage, 67
  - Azure Backup, 120–121. *See also* MARS (Microsoft Azure Recovery Services) agent
    - Azure File shares
      - backing up, 124
      - recovering, 124
    - backup policies, 127–128
    - CRR (Cross-Region Restore), 122
    - laaS VM
      - backup pre-check, 137
      - enabling backup, 136–137
      - restoring, 137–141
    - RSVs (Azure Recovery Services vaults), 121–122
      - deleting, 123
      - GRS (geo-redundant storage), 122
      - LRS (locally redundant storage), 122
      - monitoring, 122–123
      - storage redundancy, 122
      - ZRS (zone-redundant storage), 122
    - VMs
      - encrypted, restoring, 140
      - recover to new Azure VM, 138–139
      - restore and overwrite, 138
      - restore disks, 139–140
      - restoring individual files, 140–141
  - Azure Data Box, 166
  - Azure Disk Encryption, 82–84
  - Azure Migrate. *See also* migration
    - deployment
      - download a template VHD file from Azure portal, 168–169
      - script method, 169
    - Discovery and Assessment tool, 167–168
    - migrating on-premises servers to Azure, 167
    - Server Migration tool, 168
  - Azure Monitor
    - agents, 206
    - alerts, 208
    - Azure Diagnostics extension, 208–209
    - data collection rules, 207
    - installing agents, 206
    - network requirements, 207
  - Azure security baselines, 58
  - Azure Serial Console, troubleshooting laaS VM boot failures, 221
  - Azure Update Management, 65–66
- ## B
- backup and recovery, 119, 134. *See also* Azure Backup; disaster recovery; MARS (Microsoft Azure Recovery Services) agent; RSVs (Azure Recovery Services vaults)
    - AD DS (Active Directory Domain Services)
      - authoritative restore, 234–236
      - nonauthoritative restore, 236
      - restoring objects from AD recycle bin, 230–231
      - snapshots, 236–237
    - Azure Backup, 120–121
      - CRR (Cross-Region Restore), 122
      - RSVs (Azure Recovery Services vaults), 121–122

- backup policies, 127–128
  - on IaaS VMs, 136–137
- MABS (Microsoft Azure Backup Server), 124–125
  - backing up Hyper-V VMs, 126
    - installing, 125
    - recovering data, 126
  - shadow copies, 135–136
- Windows Server Backup, 128–129
  - backup locations, 129–130
  - MARS agent, 131–132
  - restore from backups, 130–131
  - restore to an alternative location, 131
  - role and application backup, 130
- balanced password policies, 28
- BitLocker, 78
  - Group Policy, 79
  - manage and recover encrypted volumes, 82
  - manage with PowerShell cmdlets, 81
  - Network Unlock, 80–81
  - recovery, 79–80
  - requirements, 78–79
- built-in administrative groups
  - Administrators, 46–47
  - Domain Admins, 46
  - Enterprise Admins, 46
  - Schema Admins, 47
  - security groups, 47–49

## C

- cache drives, S2D (Storage Spaces Direct), 113
- CAU (Cluster-Aware Updating), 106–107
- Certified for Windows Server logo requirement, 88
- CFG (Control Flow Guard), 3
- Cloud Witness, 100
- cluster sets, 99
- clustering, 87. *See also* failover clustering
- cmdlets
  - Add-ClusterNode, 107
  - Add-ClusterScaleOutFileServerRole, 97
  - Add-ClusterScaleoutFileServerRole, 99
  - Add-ClusterVMMonitoredItem, 98
  - Add-MpPreference, 5
  - BitLocker management, 81
  - Get-AZVM, 225
  - Get-ClusterNodeSupportedVersion, 108
  - Get-DnsServerStatistics, 218
  - Get-MpComputerStatus, 8
  - Get-MpPreference, 8
  - Get-MpThreat, 8
  - Get-MpThreatCatalog, 8
  - Get-MpThreatDetection, 8
  - Get-ProcessMitigation, 4
  - Install-ADServiceAccount, 20
  - Install-WindowsFeature, 89, 95, 174
  - Move-ClusterGroup, 109
  - Move-ClusterResource, 109
  - Move-ClusterSharedVolume, 109
  - Move-ClusterVirtualMachineRole, 109
  - New-ADServiceAccount, 19
  - New-Cluster, 91, 112
  - Remove-MpPreference, 5, 8
  - Remove-MPThreat, 8
  - Set-ADServiceAccount, 20
  - Set-ClusterOwnerNode, 96
  - Set-ExecutionPolicy, 243
  - Set-MpPreference, 4–5, 8
  - Set-ProcessMitigation cmdlet, 4
  - Start-MpScan, 8
  - Start-MpWDOScan, 8
  - start-service msdepsvc, 175
  - Test-Cluster, 90, 112
  - Uninstall-ADDSDomainController, 248
  - Uninstall-WindowsFeature, 8
  - Update-ClusterFunctionalLevel, 108
  - Update-MpSignature, 8
- commands, AT, 40
- connection security rules, 72–73
  - authentication exemptions, 74
  - tunnel, 73–74
- connectivity
  - hybrid network, troubleshooting, 210–211
  - on-premises, troubleshooting, 213–214
- containers, migrating IIS workloads to, 186–188
- Controlled Folder Access, 4–5
- counter alerts, Performance Monitor, 199
- CPU groups, 151
- creating
  - data collector sets, 200–201
  - Windows failover cluster, 89
- Credential Guard Readiness Tool, 7
- CRR (Cross-Region Restore), 122
- CSV (Cluster Shared Volumes), 93

## D

- dashboards, RSV (Azure Recovery Services vault), 122
- data collector sets, 200
  - creating, 200–201
  - scheduling, 201
- Datacenter edition, Windows Server, failover clustering, 88
- DAX (direct access), 114
- DCDiag, 246–247
- DCs (domain controllers), 36. *See also* RDOCs (read-only domain controllers)
  - hardening, 38–40
    - configure authentication policy silos, 41–42
    - KRBTGT account password, 40–41
    - restrict scheduled tasks, 40
  - restricting access, 43–44
- decommissioning an RODC, 38
- Defender for Identity, 51
  - architecture, 52
  - deployment, 52–53
- delegating password setting permissions, 25–26
- deleting, RSVs (Azure Recovery Services vaults), 123
- DEP (Data Execution Prevention), 3
- deployment
  - Defender for Identity, 52–53
  - MARS (Microsoft Azure Recovery Services) agent, 132–134
- DHCP (Dynamic Host Configuration Protocol)
  - migrating to Windows Server 2022, 181–182
  - preparing to migrate to Windows Server 2022, 180–195
  - troubleshooting, 219
  - verification and post-migration tasks, 182
- disabling, NTLM, 17–18
- disaster recovery, 119
  - ASR (Azure Site Recovery), 142
    - configure a recovery plan, 144
    - configure site recovery for Azure VMs, 144–145
    - site recovery for on-premises VMs, 142–144
  - Azure Backup, 120–121
    - RSVs (Azure Recovery Services vaults), 121–122
- Discovery and Assessment tool, 167–168
- disks, configuring for failover clustering, 93
- DNS (Domain Naming System)
  - troubleshooting, 214–215
    - cache locking, 217
    - DCDiag, 246–247
    - event logs, 216–217
    - netmask ordering, 218

- recursion, 217
- server tests, 216
- socket pool, 217
- zone level statistics, 218

Dockerfiles, 187–188

Domain Admins group, 46

domain isolation rules, 75–76

DSRM (Directory Services Restore Mode), 231–232

dynamic quorum, 91–92

## E

- encrypted VMs
  - restoring, 140
  - troubleshooting, 225–226
- Enterprise Admins group, 46
- Event Viewer, 202
  - event log filters, 202–203
  - event log views, 203–204
  - event subscriptions, 204–205
  - event-driven tasks, 205–206
- exploit protection, 2
  - application-level settings, 3–4
  - system-level settings, 3
- expression-based audit policies, 21–22

## F

- failover clustering, 87–88
  - CAU (Cluster-Aware Updating), 106–107
  - cluster sets, 99
  - configure workload options, 96
  - configuring network adapters, 95
  - create a Cloud witness, 100
  - CSV (Cluster Shared Volumes), 93
  - dynamic quorum, 91–92
  - file server, 96–97
    - Infrastructure, 97
    - SOFs (Scale-Out File Server), 93–97
  - floating IP address, 100
  - installing, 89
  - load balancing
    - network, 100, 102–105
    - VM, 101–102
  - manage using Windows Admin Center, 109–110
  - modify quorum options, 93–94
  - prerequisites, 89

- prestige cluster computer objects, 90
- recover a failed cluster node, 107
- S2D (Storage Spaces Direct), 112–113
  - networking, 115–116
  - upgrading, 114–115
- site spanning, 91–92
- split-brain scenarios, 92
- storage configuration, 93
- stretch clusters, 92
- upgrade a node to Windows Server 2022, 108–109
- validating, 89–90
- virtual machine, 98
- VM Monitoring, 98
- workgroup clusters, 90–91
- workloads between nodes, 109

fault domain, 107

file and folder auditing, 22

file integrity monitoring, 59

file server failover clusters, 96–97

- Infrastructure, 97
- SOFS (Scale-Out File Server), 93–97

fine-grained password policies, 26–27

firewall profiles, 68–69

floating IP address, failover clustering, 100

## G

Get-AZVM cmdlet, 225

Get-ClusterNodeSupportedVersion cmdlet, 108

Get-DnsServerStatistics cmdlet, 218

Get-MpComputerStatus cmdlet, 8

Get-MpPreference cmdlet, 8

Get-MpThreat cmdlet, 8

Get-MpThreatCatalog cmdlet, 8

Get-MpThreatDetection cmdlet, 8

Get-ProcessMitigation cmdlet, 4

Global Object Access Auditing, 22

gMSA (group-Managed Service Account), 19–20

GPOs (Group Policy Objects)

- password policies, 24
- security baseline, 15–16

Group Policy

- BitLocker, 79
- Protected Users group, 34–35
- Turn on Virtualization Based Security policy, 10
- user rights assignment policy, 12–15

groups, WSUS (Windows Server Update Services), 62–63

GRS (geo-redundant storage), 122

## H

hardening, 1

- DCs (domain controllers), 38–40
  - configure authentication policy silos, 41–42
  - KRBTGT account password, 40–41
  - restrict scheduled tasks, 40
  - restricting access, 43–44

hybrid network connectivity, troubleshooting, 210–211

Hyper-V Replica, 148–149

- configure VM replication, 152–153
- host configuration, 149
- perform a failover, 153–154
- server management, 150
  - deduplication, 151
  - smart paging, 150
  - Storage QoS, 151
  - storage tiering, 151–152
  - VM CPU groups, 151

Hyper-V VMs

- backing up, 126
- failover clustering, 98
- load balancing, 101–102
- migrating to Azure IaaS VMs, 169–171
- migrating to Windows Server 2022, 175–176
  - live migration, 176–177
  - VM exporting and import, 177–178

## I

IaaS virtual machines

- administrative connections
  - Azure AD account, 227–228
  - Azure Bastion, 228–229
  - JIT (just-in-time) access, 229
  - remote PowerShell, 228
  - Windows Admin Center, 229
- backup pre-check, 137
- boot failures, troubleshooting, 221
  - using Azure Serial Console, 221
  - using boot failure messages, 221–223
- deployment failures, troubleshooting, 220
- disk encryption issues, troubleshooting, 225–226
- enabling backup, 136–137
- extension issues, troubleshooting, 225
- manage disk encryption keys, 83–84
- migrating physical workloads, 171–172
- migrating VM workloads, 169–171

## IaaS virtual machines

- performance, troubleshooting, 224
- policies, 146–147
- repairing, 223
- site recovery, 144–145
- storage issues, troubleshooting, 226
- identity synchronization, troubleshooting, 242–243
- IIS (Internet Information Services)
  - migrating to Windows Server 2022, 174–175
  - migrating workloads to Azure Web Apps, 184–185
    - App Service Hybrid Connections, 185
    - Azure AD Application Proxy, 186
  - migrating workloads to containers, 186–188
- importing, VMs (virtual machines), 178
- inactive accounts, 32
- Infrastructure File Server, failover clustering, 97
- Install-ADServiceAccount cmdlet, 20
- installing
  - Azure Monitor agents, 206
  - failover cluster, 89
  - MABS (Microsoft Azure Backup Server), 125
  - WSMT (Windows Server Migration Tools), 174
- Install-WindowsFeature cmdlet, 89, 95, 174
- IPsec, WDFAS configuration, 71–72

## J-K

- JIT (just-in-time) access, 58, 229
- KCC (Knowledge Consistency Checker), 239
- KRBtgt account password, 40–41

## L

- LAPS (Local Administrator Password Solution), 32–34
- least privilege, 50
- load balancing
  - network, 100, 102–103
    - cluster operation modes, 103–104
    - filtering and affinity, 105
    - managing cluster nodes, 104–105
    - port rules, 104–105
    - prerequisites, 103
  - VM, 101–102
- Local Group Policy Object tool, 17
- lockout policies, 29, 31
- Log Analytics workspace, 56

- LRS (locally redundant storage), 122
- LSA (Local Security Authority) protection, 10

## M

- MABS (Microsoft Azure Backup Server), 124–125
  - backing up Hyper-V VMs, 126
  - installing, 125
  - recovering data, 126
- managing, Windows Defender, 8
- MARS (Microsoft Azure Recovery Services) agent, 131–132
  - backing up data, 134
  - deploying, 132–134
  - minimum retention limits, 132
  - restore from, 134–135
- Microsoft Defender for Cloud, 57
  - tools
    - adaptive application controls, 59
    - adaptive network hardening, 60
    - Azure security baselines, 58
    - file integrity monitoring, 59
    - JIT (just-in-time) access, 58
    - threat and vulnerability management, 58
- Microsoft Sentinel, 56–57
- migration, 157
  - AD DS infrastructure to Windows Server 2022
  - AD DS, 188–189
    - demote existing domain controllers, 192
    - DNS migration, 191–192
    - migrate AD DS objects using ADMT, 193–194
    - migrate to a new AD forest, 194–195
    - transfer FSMO role holder, 190–191
    - upgrade an existing forest, 189
  - IIS workloads to Azure Web Apps, 184–185
    - App Service Hybrid Connections, 185
    - Azure AD Application Proxy, 186
  - IIS workloads to containers, 186–188
  - on-premises servers to Azure, 166–167. *See also* Azure Migrate
    - migrate physical workloads to Azure IaaS, 171–172
    - migrate VM workloads to Azure IaaS, 169–171
      - using Azure Migrate, 167–169
  - on-premises storage to on-premises servers or Azure, 157
    - Azure Data Box, 166
    - Azure file shares, 164–166
    - copy data, 161–162

- cut over to a new server using Storage Migration Service, 162–164
- inventory source servers, 160
- Storage Migration Service Orchestrator, 159
- transfer data and shares, 158–159
- use Storage Migration Service to migrate to Azure virtual machines, 164
- from previous versions to Windows 2022, 172–173
  - DHCP (Dynamic Host Configuration Protocol), 179–182
  - IIS (Internet Information Services), 174–175
  - migrate Hyper-V hosts, 175–178
  - print servers, 182–183
  - RDS (Remote Desktop Services) host servers, 178–179
- mirror-accelerated parity, S2D (Storage Spaces Direct), 112
- monitoring
  - replication, 241
  - RSVs (Azure Recovery Services vaults), 122–123
  - servers, 201–202
  - using Azure Monitor
    - agents, 206
    - agents, installing, 206
    - alerts, 208
    - Azure Diagnostics extension, 208–209
    - data collection rules, 207
    - network requirements, 207
  - using data collector sets, 200–201
  - using Event Viewer, 202
    - event log filters, 202–203
    - event log views, 203–204
    - event subscriptions, 204–205
    - event-driven tasks, 205–206
  - using Performance Monitor, 198–199
  - using System Insights, 202
  - using VM Insights, 209–210
- Move-ClusterGroup cmdlet, 109
- Move-ClusterResource cmdlet, 109
- Move-ClusterSharedVolume cmdlet, 109
- Move-ClusterVirtualMachineRole cmdlet, 109
- multimaster replication, 239

## N

- nested resiliency, S2D (Storage Spaces Direct), 112
- network adapters, failover cluster configuration, 95
- network load balancing, 100
- Network Watcher, 210–211

- New-ADServiceAccount cmdlet, 19
- New-Cluster cmdlet, 91, 112
- NLB (network load balancing), 102–103
  - cluster operation modes, 103–104
  - filtering and affinity, 105
  - managing cluster nodes, 104–105
  - port rules, 104–105
  - prerequisites, 103
- nonauthoritative restore, 236
- nonexpiring passwords, 30–31
- ntdsutil.exe, 248
- NTPM, disabling, 17–18

## O

- OU (organizational unit), 89
- Overview dashboard, RSV (Azure Recovery Services vault), 122
- overwriting, VMs, 138

## P

- Package Inspector, 6–7
- pass-through authentication, troubleshooting, 244–245
- password hash synchronization, troubleshooting, 243–244
- password policy(ies), 24, 25. *See also* account management; LAPS (Local Administrator Password Solution)
  - account lockout settings, 29
  - Azure AD Password Protection, 34
  - balanced, 28
  - delegating password setting permissions, 25–26
  - determining password settings, 28
  - fine-grained, 26–27
  - KRBTGT account password, 40–41
  - locked-out accounts, 31
  - nonexpiring passwords, 30–31
  - password replication, 36–37
  - PSO (Password Settings Object), 27–28
- PAWs (Privileged Access Workstations), 43–44
- performance, troubleshooting on VMs, 224
- Performance Monitor, 198, 199
- point-to-site VPN, troubleshooting, 212–213
- Policy Analyzer tool, 15–17
- policy(ies)



## policy(ies)

- audit, 21
    - auditpol.exe, 22–23
    - expression-based, 21–22
    - file and folder, 22
    - settings, 21
  - authentication, 41–42
  - backup, 127–128
  - password, 24, 25
    - account lockout settings, 29
    - Azure AD Password Protection, 34
    - balanced, 28
    - delegating password setting permissions, 25–26
    - determining password settings, 28
    - fine-grained, 26–27
    - KRBTGT account, 40–41
    - locked-out accounts, 31
    - nonexpiring, 30–31
    - PSO (Password Settings Object), 27–28
  - site recovery, 146–147
  - Turn on Virtualization Based Security, 7, 10
  - user rights, 12–15
  - WSUS (Windows Server Update Services), 63
  - port rules, NLB (network load balancing), 104–105
  - PowerShell
    - Azure Migrate deployment script, 169
  - cmdlets
    - Add-ClusterNode, 107
    - Add-ClusterScaleOutFileServerRole, 97
    - Add-ClusterScaleoutFileServerRole, 99
    - Add-ClusterVMMonitoredItem, 98
    - Add-MpPreference, 5, 8
    - BitLocker management, 81
    - Get-AZVM, 225
    - Get-ClusterNodeSupportedVersion, 108
    - Get-DnsServerStatistics, 218
    - Get-MpComputerStatus, 8
    - Get-MpPreference, 8
    - Get-MpThreat, 8
    - Get-MpThreatCatalog, 8
    - Get-MpThreatDetection, 8
    - Get-ProcessMitigation, 4
    - Install-ADServiceAccount, 20
    - Install-WindowsFeature, 89, 95, 174
    - Move-ClusterGroup, 109
    - Move-ClusterResource, 109
    - Move-ClusterSharedVolume, 109
    - Move-ClusterVirtualMachineRole, 109
    - New-ADServiceAccount, 19
    - New-Cluster, 91, 112
    - Remove-MpPreference, 5, 8
    - Remove-MPThreat, 8
    - Set-ADServiceAccount, 20
    - Set-ClusterOwnerNode, 96
    - Set-ExecutionPolicy, 243
    - Set-MpPreference, 4–5, 8
    - Set-ProcessMitigation, 4
    - Start-MpScan, 8
    - Start-MpWDOScan, 8
    - start-service msdepsvc, 175
    - Test-Cluster, 90, 112
    - Uninstall-ADDSDomainController, 248
    - Uninstall-WindowsFeature, 8
    - Update-ClusterFunctionalLevel, 108
    - Update-MpSignature, 8
  - remote, 228
  - print servers, migrating to Windows Server 2022, 182–183
  - Protected Users group, 9, 34–35
  - PSO (Password Settings Object), 27–28
- ## Q-R
- quorum
    - dynamic, 91–92
    - failover cluster configuration, 93–94
  - RBAC (role-based access control), 50–51
  - RDS (Remote Desktop Services), migrating to Windows Server 2022, 178–179
  - recovering, Azure File shares, 124
  - recovery plans, 144. *See also* backup and recovery
  - remote PowerShell, 228
  - Remove-MpPreference cmdlet, 5, 8
  - Remove-MPThreat cmdlet, 8
  - repadmin, 241
  - repairing IaaS VMs, 223
  - replication
    - monitoring, 241
    - multimaster, 239
    - RODC, 240–241
    - store and forward, 239
    - troubleshooting, 238
  - requirements
    - failover clustering, 89
    - Windows Defender Credential Guard, 9–10
  - resetting passwords
    - Azure AD self-service password reset, 53–55
    - delegating password setting permissions, 25–26

- resiliency types, S2D (Storage Spaces Direct), 113
- RODCs (read-only domain controllers), 36
  - decommissioning, 38
  - local administrators, 38
  - partial attribute set, 37–38
  - password replication, 36–37
  - replication, 240–241
- RSVs (Azure Recovery Services vaults)
  - alerts, 123
  - CRR (Cross-Region Restore), 122
  - deleting, 123
  - GRS (geo-redundant storage), 122
  - LRS (locally redundant storage), 122
  - monitoring, 122–123
  - Overview dashboard, 122
  - ZRS (zone-redundant storage), 122
- rules
  - automatic approval, 64–65
  - Azure Monitor, 208
  - connection security, 72–73
    - authentication exemptions, 74
    - tunnel, 73–74
  - data collection, 207
  - domain isolation, 75–76
  - inbound, 69–70
  - outbound, 70–71
  - port, 104–105
- run profile, 107

## S

- S2D (Storage Spaces Direct), 110–111
  - cache drives, 113
  - DAX (direct access), 114
  - deployment options, 111–112
    - mirror-accelerated parity, 112
    - nested resiliency, 112
  - failover clustering, 112–113
    - upgrading, 114–115
  - networking, 115–116
  - properties, 111
  - resiliency types, 113
- Schema Admins group, 47
- SCT (Security Compliance Toolkit), 15
  - Local Group Policy Object tool, 17
  - Policy Analyzer tool, 15–17
- Seamless SSO, troubleshooting, 245–246
- security. *See also* WDFAS (Windows Defender Firewall with Advanced Security)
  - account management, 44–45
  - Controlled Folder Access, 4–5
  - DCs (domain controllers), 38–40
    - configure authentication policy silos, 41–42
    - KRBTGT account password, 40–41
    - restrict scheduled tasks, 40
    - restricting access, 43–44
  - exploit protection, 2
    - application-level settings, 3–4
    - system-level settings, 3
  - hardening, 1
  - Protected Users group, 34–35
  - RODCs (read-only domain controllers), 36
    - decommissioning, 38
    - local administrators, 38
    - partial attribute set, 37–38
    - password replication, 36–37
  - WDAC (Windows Defender Application Control), 5–6
    - Audit mode, 6–7
    - features, 6
    - Turn on Virtualization Based Security policy, 7
    - Virtual Secure Mode, enabling, 6
  - Windows Defender Credential Guard, 9
    - Enabled With UEFI Lock option, 10
    - features and solutions, 9
    - LSA (Local Security Authority) protection, 10
    - Protected Users group, 9
    - requirements, 9–10
    - restrictions, 9
    - Turn on Virtualization Based Security policy, 10
- Server Migration tool, 168
- service accounts, 18–20
- Set-ADServiceAccount cmdlet, 20
- Set-ClusterOwnerNode cmdlet, 96
- Set-ExecutionPolicy cmdlet, 243
- Set-MpPreference cmdlet, 4–5, 8
- Set-ProcessMitigation cmdlet, 4
- shadow copies, 135–136
- SIEM (security information and event management), 56
- silos claim, 42
- site recovery
  - for Azure VMs, 144–145
  - networking, 145–146
  - policy, 146–147
  - for on-premises VMs, 142–144
- site-to-site VPN, troubleshooting, 213
- smart paging, 150
- SmartScreen, 11
- snapshots, 236–237

## SOAR (security orchestration, automation, and response)

- SOAR (security orchestration, automation, and response), 56
- split-brain scenarios, 92
- Start-MpScan cmdlet, 8
- Start-MpWDOScan cmdlet, 8
- start-service msdepsvc cmdlet, 175
- storage. *See also* backup and recovery
  - Azure Data Box, 166
  - configuring for failover clustering, 93
  - IaaS virtual machines
    - troubleshooting, 226
  - RSVs (Azure Recovery Services vaults), 121–122
    - alerts, 123
    - deleting, 123
    - GRS (geo-redundant storage), 122
    - LRS (locally redundant storage), 122
    - ZRS (zone-redundant storage), 122
  - tiering, 151–152
- Storage Migration Service, 158–159
  - cutover phase, 162–164
  - migrate to Azure virtual machines, 164
- Storage Migration Service Orchestrator, 159
- Storage QoS, 151
- store and forward replication, 239
- stretch clusters, 92
- System Insights, 202

## T

- Test-Cluster cmdlet, 90, 112
- threat and vulnerability management, 58
- tombstone lifetime, 232–234
- tool(s). *See also* WSMT (Windows Server Migration Tools)
  - Active Directory Migration, 193–194
  - Azure Migrate
    - Discovery and Assessment, 167–168
    - Server Migration, 168
  - Credential Guard Readiness, 7
  - DCDiag, 246–247
  - LAPS (Local Administrator Password Solution), 32–34
  - Local Group Policy Object, 17
  - Microsoft Defender for Cloud
    - adaptive application controls, 59
    - adaptive network hardening, 60
    - Azure security baselines, 58
    - file integrity monitoring, 59
    - JIT (just-in-time) access, 58
    - threat and vulnerability management, 58

- Network Watcher, 210–211
- Policy Analyzer, 15–17
- repadmin, 241
- vssadmin, 135–136
- Web Deploy, 175
- troubleshooting
  - AD DS (Active Directory Domain Services)
    - DNS (Domain Naming System), 246–247
    - KCC (Knowledge Consistency Checker), 239
    - replication, 238
    - SYSVOL, recovering, 242
  - authentication
    - identity synchronization, 242–243
    - pass-through, 244–245
    - password hash synchronization, 243–244
    - Seamless SSO, 245–246
  - Azure VPN, 208–212
    - point-to-site, 212–213
    - site-to-site, 213
  - connectivity
    - hybrid network, 210–211
    - on-premises, 213–214
  - DHCP, 219
  - DNS, 214–215
    - cache locking, 217
    - event logs, 216–217
    - netmask ordering, 218
    - recursion, 217
    - server tests, 216
    - socket pool, 217
    - zone level statistics, 218
  - IaaS virtual machines
    - boot failures, 221–223
    - connection issues, 227–229
    - deployment failures, 220
    - extension issues, 225
    - performance issues, 224
    - storage issues, 226
- trustlets, 9
- tunnel rules, 73–74
- Turn on Virtualization Based Security policy, 7

## U

- Uninstall-ADDSDomainController cmdlet, 248
- uninstalling, Windows Defender, 8
- Uninstall-WindowsFeature cmdlet, 8
- Update-ClusterFunctionalLevel cmdlet, 108
- Update-MpSignature cmdlet, 8

updates. *See also* WSUS (Windows Server Update Services)  
 deploying, 63–64  
 upgrading, S2D cluster, 114–115  
 user rights assignment policy, 12–15  
 USNs (update sequence numbers), 239–240

## V

validating, failover cluster, 89–90  
 virtual account, 20  
 Virtual Secure Mode, WDAC (Windows Defender Application Control), 6  
 VM Insights, 209–210  
 VMs (virtual machines). *See also* Hyper-V VMs  
   CPU groups, 151  
   encrypted, restoring, 140  
   extension issues, troubleshooting, 225  
   failover clustering, 98  
   laaS  
     administrative connections, 227–229  
     backup pre-check, 137  
     boot failures, troubleshooting, 221–223  
     deployment failures, troubleshooting, 220  
     disk encryption, 83–84  
     enabling backup, 136–137  
     migrating physical workloads, 171–172  
     migrating VM workloads, 169–171  
     repairing, 223  
   importing, 178  
   JIT (just-in-time) access, 58, 229  
   load balancing, 101–102  
   migration, 164  
   monitoring, 209–210  
   performance, troubleshooting, 224  
   on-premises, site recovery, 142–144  
   recover to new Azure VM, 138–139  
   restore and overwrite, 138  
   restore disks, 139–140  
   restoring individual files, 140–141  
   site recovery, 144–145  
     for Azure VMs, 144–145  
     network, 145–146  
     for on-premises VMs, 142–144  
 VPN, troubleshooting, 208–212  
   point-to-site, 212–213  
   site-to-site, 213  
 VSS (Volume Shadow Copy Services), 135–136  
 vssadmin, 135–136

## W

WDAC (Windows Defender Application Control), 5–6  
   Audit mode, 6–7  
   features, 6  
   Turn on Virtualization Based Security policy, 7  
   Virtual Secure Mode, enabling, 6  
 WDFAS (Windows Defender Firewall with Advanced Security), 68  
   connection security rules, 72–73  
     authentication exemptions, 74  
     tunnel rules, 73–74  
   domain isolation rules, 75–76  
   firewall profiles, 68–69  
   inbound rules, 69–70  
   IPsec, 71–72  
   outbound rules, 70–71  
 Web Deploy, 175, 187  
 Windows Admin Center  
   managing failover clusters, 109–110  
   monitoring servers, 201–202  
   troubleshooting remote VM connectivity, 229  
 Windows Defender. *See also* Defender for Identity  
 Credential Guard, 9  
   Enabled With UEFI Lock option, 10  
   features and solutions, 9  
   LSA (Local Security Authority) protection, 10  
   Protected Users group, 9  
   requirements, 9–10  
   restrictions, 9  
   Turn on Virtualization Based Security policy, 10  
 gMSA (group-Managed Service Account), 19–20  
   managing, 8  
   uninstalling, 8  
 Windows Server. *See also* migration  
   audit policies, 21  
     auditpol.exe, 22–23  
     expression-based, 21–22  
     file and folder, 22  
     settings, 21  
   Azure Update Management, 65–66  
   backup and recovery, 119  
   BitLocker, 78  
     Group Policy, 79  
     manage and recover encrypted volumes, 82  
     Network Unlock, 80–81  
     recovery, 79–80  
     requirements, 78–79  
   Controlled Folder Access, 4–5

- exploit protection, 2
  - application-level settings, 3–4
  - system-level settings, 3
- failover clustering, 87–88. *See also* S2D (Storage Spaces Direct)
  - CAU (Cluster-Aware Updating), 106–107
  - cluster sets, 99
  - configure workload options, 96
  - configuring network adapters, 95
  - create a Cloud witness, 100
  - CSV (Cluster Shared Volumes), 93
  - dynamic quorum, 91–92
  - File Server cluster role, 96–97
  - floating IP address, 100
  - installing, 89
  - load balancing, 100
  - manage using Windows Admin Center, 109–110
  - modify quorum options, 93–94
  - NLB (network load balancing), 102–105
  - prerequisites, 89
  - prestige cluster computer objects, 90
  - recover a failed cluster node, 107
  - S2D (Storage Spaces Direct), 112–113, 115–116
  - site spanning, 91–92
  - split-brain scenarios, 92
  - storage configuration, 93
  - stretch clusters, 92
  - upgrade a node to Windows Server 2022, 108–109
  - validating, 89–90
  - virtual machine, 98
  - VM load balancing, 101–102
  - VM Monitoring, 98
  - workgroup clusters, 90–91
  - workloads between nodes, 109
- Group Policy, user rights assignment policy, 12–15
- NLTM, disabling, 17–18
- S2D (Storage Spaces Direct), 110–111
  - cache drives, 113
  - DAX (direct access), 114
  - deployment options, 111–112
  - mirror-accelerated parity, 112
  - nested resiliency, 112
  - properties, 111
  - resiliency types, 113
- SCT (Security Compliance Toolkit), 15
  - Local Group Policy Object tool, 17
  - Policy Analyzer tool, 15–17
- service accounts, 18–20
- SmartScreen, 11
- Storage Migration Service, 158–159
- Storage Migration Service Orchestrator, 159
- System Insights, 202
- Update Services, 60
  - automatic approval rules, 64–65
  - autonomous and replica modes, 61
  - deploying updates, 63–64
  - groups, 62–63
  - policies, 63
  - products, security classifications, and languages, 60–61
  - security roles, 62
  - update files, 61–62
- WDAC (Windows Defender Application Control), 5–6
  - Audit mode, 6–7
  - features, 6
  - Turn on Virtualization Based Security policy, 7
  - Virtual Secure Mode, enabling, 6
- Windows Defender
  - Credential Guard, 9–11
  - uninstalling, 8
- Windows Server Backup, 128–129
  - backup locations, 129–130
- MARS agent, 131–132
  - deploying, 132–134
  - minimum retention limits, 132
  - restore from backups, 130–131
  - restore to an alternative location, 131
  - role and application backup, 130
- workgroup clusters, 90–91
- workload. *See also* VMs (virtual machines)
  - cluster, 96
  - migration. *See* migration
- WSMT (Windows Server Migration Tools), 172–174
  - installing, 174
  - requirements, 174
- WSUS (Windows Server Update Services), 60
  - automatic approval rules, 64–65
  - autonomous and replica modes, 61
  - deploying updates, 63–64
  - groups, 62–63
  - policies, 63
  - products, security classifications, and languages, 60–61
  - security roles, 62
  - update files, 61–62

## X-Y-Z

ZRS (zone-redundant storage), 122