



Microsoft Azure Fundamentals

Exam Ref AZ-900

Jim Cheshire

Contents

1. Cover Page
2. Title Page
3. Copyright Page
4. Dedication Page
5. Contents at a glance
6. Contents
7. Acknowledgments
8. About the Author
9. Introduction
 1. Organization of this book
 2. Microsoft certifications
 3. Quick access to online references
 4. Errata, updates, & book support
 5. Stay in touch
10. Preparing for the exam
11. Chapter 1. Understand cloud concepts
 1. Skill 1.1: Describe the benefits and considerations of using cloud services
 2. Skill 1.2: Describe the differences between Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), and Software-as-a-Service (SaaS)
 3. Skill 1.3: Describe the differences between public, private, and

hybrid cloud models

4. Thought experiment
5. Thought experiment answers
6. Chapter summary

12. Chapter 2. Understand core Azure services

1. Skill 2.1: Understand the core Azure architectural components
2. Skill 2.2: Describe some of the core products available in Azure
3. Skill 2.3: Describe some of the solutions available on Azure
4. Skill 2.4: Understand Azure management tools
5. Thought experiment
6. Thought experiment answers
7. Chapter summary

13. Chapter 3. Understand security, privacy, compliance, and trust

1. Skill 3.1: Understand securing network connectivity in Azure
2. Skill 3.2: Describe core Azure Identity services
3. Skill 3.3: Describe security tools and features of Azure
4. Skill 3.4: Describe Azure governance methodologies
5. Skill 3.5: Understand monitoring and reporting options in Azure
6. Skill 3.6: Understand privacy, compliance, and data protection standards in Azure
7. Thought experiment
8. Thought experiment answers
9. Chapter summary

14. Chapter 4. Understand Azure pricing and support

1. Skill 4.1: Understand Azure subscriptions
2. Skill 4.2: Understand planning and management of costs

3. Skill 4.3: Understand the support options available in Azure
4. Skill 4.4: Describe Azure service level agreements
5. Skill 4.5: Understand service lifecycle in Azure
6. Thought experiment
7. Thought experiment answers
8. Chapter summary

15. Index

16. Code Snippets

1. i
2. ii
3. iii
4. iv
5. v
6. vi
7. vii
8. viii
9. ix
10. x
11. xi
12. xii
13. xiii
14. xiv
15. xv
16. xvi
17. xvii
18. xviii
19. xix

20. xx

21. 1

22. 2

23. 3

24. 4

25. 5

26. 6

27. 7

28. 8

29. 9

30. 10

31. 11

32. 12

33. 13

34. 14

35. 15

36. 16

37. 17

38. 18

39. 19

40. 20

41. 21

42. 22

43. 23

44. 24

45. 25

46. 26

- 47. 27
- 48. 28
- 49. 29
- 50. 30
- 51. 31
- 52. 32
- 53. 33
- 54. 34
- 55. 35
- 56. 36
- 57. 37
- 58. 38
- 59. 39
- 60. 40
- 61. 41
- 62. 42
- 63. 43
- 64. 44
- 65. 45
- 66. 46
- 67. 47
- 68. 48
- 69. 49
- 70. 50
- 71. 51
- 72. 52
- 73. 53

- 74. 54
- 75. 55
- 76. 56
- 77. 57
- 78. 58
- 79. 59
- 80. 60
- 81. 61
- 82. 62
- 83. 63
- 84. 64
- 85. 65
- 86. 66
- 87. 67
- 88. 68
- 89. 69
- 90. 70
- 91. 71
- 92. 72
- 93. 73
- 94. 74
- 95. 75
- 96. 76
- 97. 77
- 98. 78
- 99. 79
- 100. 80

- 101. 81
- 102. 82
- 103. 83
- 104. 84
- 105. 85
- 106. 86
- 107. 87
- 108. 88
- 109. 89
- 110. 90
- 111. 91
- 112. 92
- 113. 93
- 114. 94
- 115. 95
- 116. 96
- 117. 97
- 118. 98
- 119. 99
- 120. 100
- 121. 101
- 122. 102
- 123. 103
- 124. 104
- 125. 105
- 126. 106
- 127. 107

128. 108

129. 109

130. 110

131. 111

132. 112

133. 113

134. 114

135. 115

136. 116

137. 117

138. 118

139. 119

140. 120

141. 121

142. 122

143. 123

144. 124

145. 125

146. 126

147. 127

148. 128

149. 129

150. 130

151. 131

152. 132

153. 133

154. 134

155. 135

156. 136

157. 137

158. 138

159. 139

160. 140

161. 141

162. 142

163. 143

164. 144

165. 145

166. 146

167. 147

168. 148

169. 149

170. 150

171. 151

172. 152

173. 153

174. 154

175. 155

176. 156

177. 157

178. 158

179. 159

180. 160

181. 161

- 182. [162](#)
- 183. [163](#)
- 184. [164](#)
- 185. [165](#)
- 186. [166](#)
- 187. [167](#)
- 188. [168](#)
- 189. [169](#)
- 190. [170](#)
- 191. [171](#)
- 192. [172](#)
- 193. [173](#)
- 194. [174](#)
- 195. [175](#)
- 196. [176](#)
- 197. [177](#)
- 198. [178](#)
- 199. [179](#)
- 200. [180](#)
- 201. [181](#)
- 202. [182](#)
- 203. [183](#)
- 204. [184](#)
- 205. [185](#)
- 206. [186](#)
- 207. [187](#)
- 208. [188](#)

- 209. [189](#)
- 210. [190](#)
- 211. [191](#)
- 212. [192](#)
- 213. [193](#)
- 214. [194](#)
- 215. [195](#)
- 216. [196](#)
- 217. [197](#)
- 218. [198](#)
- 219. [199](#)
- 220. [200](#)
- 221. [201](#)
- 222. [202](#)
- 223. [203](#)
- 224. [204](#)
- 225. [205](#)
- 226. [206](#)
- 227. [207](#)
- 228. [208](#)
- 229. [209](#)
- 230. [210](#)
- 231. [211](#)
- 232. [212](#)
- 233. [213](#)
- 234. [214](#)
- 235. [215](#)

236. [216](#)

237. [217](#)

238. [218](#)

239. [219](#)

240. [220](#)

241. [221](#)

242. [222](#)

243. [223](#)

244. [224](#)

245. [225](#)

246. [226](#)

247. [227](#)

248. [228](#)

249. [229](#)

250. [230](#)

251. [231](#)

252. [232](#)

253. [233](#)

254. [234](#)

255. [235](#)

256. [236](#)

257. [237](#)

258. [238](#)

259. [239](#)

260. [240](#)

261. [241](#)

262. [242](#)

263. 243

264. 244

265. 245

266. 246

267. 247

268. 248

269. 249

270. 250

271. 251

272. 252

273. 253

274. 254

275. 255

276. 256

277. 257

278. 258

279. 259

280. 260

281. 261

282. 262

283. 263

284. 264

285. 265

286. 266

287. 267

288. 268

289. 269

290. 270

291. 271

292. 272

293. 273

294. 274

295. 275

296. 276

297. 277

298. 278

299. 279

300. 280

301. 281

302. 282

303. 283

304. 284

Exam Ref AZ-900 Microsoft Azure Fundamentals

Jim Cheshire



Exam Ref AZ-900 Microsoft Azure Fundamentals

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

Copyright © 2019 by Pearson Education

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.pearsoned.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-978-0-1357-3218-2

ISBN-0-1357-3218-2

Trademarks

Microsoft and the trademarks listed at <https://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 authors, 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 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 or (800) 382-3419.

For government sales inquiries, please contact
governmentsales@pearsoned.com.

For questions about sales outside the U.S., please
contact intlcs@pearson.com.

Editor-in-Chief

Brett Bartow

Executive Editor

Loretta Yates

Sponsoring Editor

Charvi Arora

Development Editor

Troy Mott

Managing Editor

Sandra Schroeder

Senior Project Editor

Tracey Croom

Editorial Production

Backstop Media

Copy Editor

Liv Bainbridge

Indexer

MAP Systems

Proofreader

Jana Gardner

Technical Editor

Timothy Warner

Cover Designer

Twist Creative, Seattle

*I dedicate this book to my wife, Becky, my
daughter, Hope, and my son, James.*

—Jim Cheshire

Contents at a glance

Introduction

Preparing for the exam

CHAPTER 1 Understand cloud concepts

CHAPTER 2 Understand core Azure services

**CHAPTER 3 Understand security, privacy,
compliance, and trust**

**CHAPTER 4 Understand Azure pricing and
support**

Index

Contents

Introduction

Organization of this book

Microsoft certifications

Quick access to online references

Errata, updates, & book support

Stay in touch

Preparing for the exam

Understand cloud concepts

Skill 1.1: Describe the benefits and
considerations of using cloud services

High Availability

Scalability, elasticity, and agility

Fault tolerance and disaster
recovery

Economic benefits of the cloud

Skill 1.2: Describe the differences between
Infrastructure-as-a-Service (IaaS),
Platform-as-a-Service (PaaS), and

Software-as-a-Service (SaaS)

Infrastructure-as-a-Service (IaaS)

Platform-as-a-Service (PaaS)

Software-as-a-Service (SaaS)

Comparing service types

Skill 1.3: Describe the differences between
public, private, and hybrid cloud models

The public cloud

The private cloud

The hybrid cloud

Thought experiment

Thought experiment answers

Chapter summary

Understand core Azure services

Skill 2.1: Understand the core Azure
architectural components

Azure regions

Availability zones

Azure Resource Manager (ARM)

Resource groups

Skill 2.2: Describe some of the core
products available in Azure

Azure compute products

Azure networking products

Azure storage products

Azure database products

The Azure Marketplace and its
usage scenarios

Skill 2.3: Describe some of the solutions
available on Azure

Internet of Things (IoT)

Big Data and analytics

Artificial Intelligence

Serverless computing

Skill 2.4: Understand Azure management
tools

The Azure portal

Azure and PowerShell

Azure CLI

Azure Advisor

Thought experiment

Thought experiment answers

Chapter summary

Understand security, privacy, compliance, and

trust

Skill 3.1: Understand securing network connectivity in Azure

Azure Firewall

DDoS Protection

Network Security Groups

Choosing an appropriate Azure security solution

Skill 3.2: Describe core Azure Identity services

Azure Active Directory

Multi-factor authentication

Skill 3.3: Describe security tools and features of Azure

Azure Security Center

Azure Key Vault

Azure Information Protection

Azure Advanced Threat Protection

Skill 3.4: Describe Azure governance methodologies

Azure Policy

Role-based access control

Locks

Azure Advisor

Skill 3.5: Understand monitoring and reporting options in Azure

Azure Monitor

Azure Service Health

Skill 3.6: Understand privacy, compliance, and data protection standards in Azure

Microsoft Privacy Statement

Trust Center

Service Trust Portal

Compliance Manager

Azure Government

Azure Germany

Thought experiment

Thought experiment answers

Chapter summary

Understand Azure pricing and support

Skill 4.1: Understand Azure subscriptions

Azure subscription

Uses and options with Azure subscriptions

Skill 4.2: Understand planning and

management of costs

Options for purchasing Azure
products and services

Options around Azure free account

Factors affecting costs

Zones

The pricing calculator

The total cost of ownership (TCO)
calculator

Best practices for minimizing Azure
costs

Azure Cost Management

Skill 4.3: Understand the support options
available in Azure

Support plans

How to open a support case

Available support channels outside
of support plans

Knowledge Center

Skill 4.4: Describe Azure service level
agreements

Service level agreement (SLA)

Determine the SLA for a particular

Azure product or service

Skill 4.5: Understand service lifecycle in
Azure

Public and private preview features

How to access preview features

General availability

Monitoring feature updates

Thought experiment

Thought experiment answers

Chapter summary

Index

Acknowledgments

I'd like to express my deep gratitude to the following people, without whom this book would not have been possible.

Thank you to Loretta for bringing me into this project. After two decades of working together on numerous projects, you still seem to find a way to bring freshness and excitement to each one. To Troy for your always-present ear when I needed to bounce an idea off someone, and for your experienced counsel during the editing process. Thank you to Liv for your unwavering work during copy editing and helping to tighten things up. Thanks to Tim for all the times you made me take a second look at my approach, and for adding real value with your ideas. Finally, thank you to all the people at Microsoft Press who worked so hard to create this book from the digital manuscript.

About the Author

JIM CHESHIRE is a technology enthusiast with over 25 years of experience in various roles within IT. Jim has authored more than 15 books on technology, and he's held numerous training sessions on Microsoft Azure, both in private enterprises and through Safari's Live Training program. Jim is heavily involved in Azure and is in his 21st year at Microsoft. He's currently working as an engineer in Azure App Service.

Introduction

Both businesses and individuals are adopting cloud technologies at a breakneck pace, and Microsoft Azure is often the choice for cloud-based applications and services. The purpose of the AZ-900 exam is to test your understanding of the fundamentals of Azure. The exam includes high-level concepts that apply across all of Azure to important concepts that are specific to a particular Azure service. Like the exam, this book is geared towards giving you a broad understanding of Azure itself and of many of the common services and components in Azure.

While we've made every effort possible to make the information in this book accurate, Azure is rapidly evolving, and there's a chance that some of the screens in the Azure portal are slightly different now than they were when this book was written. It's also possible that other minor changes have taken place such as minor name changes in features and so on.

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 MSDN, TechNet, 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: <http://aka.ms/examlist>. 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 <http://www.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

<https://MicrosoftPressStore.com/ExamRefAZ900/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:

<https://MicrosoftPressStore.com/ExamRefAZ900/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 <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 <https://support.microsoft.com>.

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. Although there is no substitute for on-the-job experience, preparation through study and hands-on practice can help you prepare for the exam. We recommend that you augment 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 publicly 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.

Chapter 1. Understand cloud concepts

Cloud computing has been part of information technology (IT) for over 20 years. During that time, it has evolved into a complex collection of cloud services and cloud models. Before you begin the process of moving to the cloud, it's important that you understand key concepts and services related to the cloud.

Important Have you read page xix?

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

There are many reasons for moving to the cloud, but one of the primary benefits is removing some of the IT burden from your own company. The cloud allows you to take advantage of a cloud provider's infrastructure and investments, and it makes it easier to maintain consistent access to your applications and data. You'll also gain the benefit of turn-key solutions for backing up data and ensuring your applications can survive disasters and other availability problems. Hosting your data and applications in the cloud is often more cost-

effective than investing in infrastructure and on-premises IT resources.

Once you decide to take advantage of the cloud, you need to understand the different cloud offerings available to you. Some cloud services provide an almost hands-off experience, while others require you to manage some of the systems yourself. Finding the right balance for your needs requires that you fully understand each type of service.

This chapter covers the benefits of using the cloud, the different cloud services that are available, and cloud models that enable a variety of cloud configurations.

Skills covered in this chapter:

- Describe the benefits and considerations of using cloud services
- Describe the differences between Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), and Software-as-a-Service (SaaS)
- Describe the differences between Public, Private, and Hybrid cloud models

SKILL 1.1: DESCRIBE THE BENEFITS AND CONSIDERATIONS OF USING CLOUD SERVICES

Today's companies rely heavily on software solutions and access to data. In fact, in many cases, a company's most valuable assets are directly tied to data and

applications. Because of that, investment in IT has grown tremendously over the past couple of decades. Reliance in on-premises IT departments worked well in the early days of IT, but access to data and applications has become such a critical part of day-to-day operations that localized IT systems have become inefficient on many levels.

When making decisions about what to move to the cloud and the benefit associated with cloud solutions, evaluate these decisions against the benefits that cloud computing can provide.

This section covers:

- High availability
- Scalability and elasticity
- Agility
- Fault tolerance and disaster recovery
- Principles of economies of scale
- Differences between capital expenditures and operations expenditures
- Consumption-based model

High Availability

The availability of data and applications is a core requirement for any application, whether it is on-premises or in the cloud. If your data or application isn't

available to you, nothing else matters. There are many reasons why you may lose availability, but the most common issues are:

- A network outage
- An application failure
- A system, such as a virtual machine, outage
- A power outage
- A problem with a reliant system such as an external database

In a perfect world, you experience 100% availability, but if any of the above problems occur, that percentage will begin to decrease. Therefore, it's critical that your infrastructure minimize the risk of problems that impact availability of your application.

Cloud providers offer a *service-level agreement* (SLA) that guarantees a certain level of availability as a percentage. An SLA will usually guarantee an uptime of close to 100%, but it only covers systems that are controlled by the cloud provider.

An application hosted in the cloud might be one that is developed by your company, but it can also be one provided to you by the cloud provider.

Network outage

All applications require some level of network connectivity. Users of an application require network

connectivity to the computers that run the application. The application requires network connectivity to required back-end systems such as database servers. Applications may also call into other applications using a network. If any of these network connections fail, they can cause a lack of availability.

More Info Planning for Network Outages

A network failure doesn't have to mean that your application or data is unavailable. If you plan carefully, you can often avoid an application problem when a network problem occurs. We'll cover that in more detail when we discuss fault tolerance later in this chapter.

Cloud providers invest a lot of money in network infrastructure, and by moving to the cloud you gain the benefit of that infrastructure and the additional reliability that comes with it. If something within that infrastructure fails, the cloud provider diagnoses and fixes it, often before you even realize there's a problem.

Application failure

An application failure is often the result of a software bug, but it can also be caused by application design.

More Info Application Design and the Cloud

You don't need to understand application design concepts for the AZ-900 exam, but if you're interested in learning more about application design and the cloud, Microsoft has a good reference at: <https://docs.microsoft.com/en-us/azure/architecture/patterns/>.

In some cloud scenarios, you are still responsible for application failures, but your cloud provider likely provides you with tools that you can use to diagnose these failures more easily. For example, Azure offers a service called Application Insights that integrates with your application to give you detailed information about the performance and reliability of your application. Application developers can often use this information to get right to the code where a problem is happening, dramatically reducing the time needed for troubleshooting.

Cloud providers offer other features that can reduce availability impacts caused by application failure. You can often test new versions of an application in a protected environment without impact to real users. When you're ready to move actual users to a new version, you can often move a small number of users first to ensure things are working correctly. If you discover problems, the cloud often makes it easy to roll things back to the prior version.

System outage

A system outage occurs when the computer running a particular system becomes unavailable. In the on-premises world, that computer might be a server running a database or another part of the application. In the cloud, these systems run inside of *virtual machines*,

or VMs.

VMs are software-based computers that run on a physical computer. A single computer can run multiple VMs, and each VM has its own isolated operating system and applications. All VMs running on a computer share the CPU, memory, and storage of the host computer they run on.

Note VMS Aren't Just for the Cloud

VMs make it easy to add additional computers when necessary, and they allow you to better manage computer resources such as CPU, disk space, and memory. For that reason, VMs are commonplace in most businesses.

Depending on the cloud service you choose, you may or may not be responsible for maintaining VMs. However, whether you or your cloud provider maintain them, the cloud provider will constantly monitor the health of VMs and will have systems in place to recover an unhealthy VM.

Power Outage

Reliable electricity is critical to availability. Even a quick power flicker can cause computers to reboot and systems to restart. When that happens, your application is unavailable until all systems are restored.

Cloud providers invest heavily in battery-operated

power backup and other redundant systems in order to prevent availability problems caused by power outages. In a situation where a large geographic area is impacted by a power outage, cloud providers offer you the ability to run your application from another region that isn't impacted.

Problems with a reliant system

Your application may use systems that aren't in the cloud or that are hosted by a different cloud provider. If those systems fail, you may lose availability. By hosting your application in the cloud, you gain the benefit of troubleshooting, alerting, and diagnosis tools that the cloud provider offers.

Now that you have an understanding of some of the things that can impact availability, and some general advantages of the cloud in helping to alleviate those problems, let's review some of the specific ways that the cloud can help you ensure high availability.

Scalability, elasticity, and agility

Computing resources aren't free. Even if you're using virtual machines, the underlying resources such as disk space, CPU, and memory cost money. The best way to minimize cost is to use only the resources necessary for your purposes. The challenge is that resource needs can change often and quickly.

Consider a situation where you are hosting an application in the cloud that tracks sales data for your company. If your sales staff regularly enter information on daily sales calls at the end of the day, you might need additional computing resources to handle that load. Those same resources aren't needed during the day when the sales staff is making sales calls and not using the application.

You might also host a web application in the cloud that is used by external customers. Depending on the usage pattern, you might want to add additional computing resources on certain days or during certain times. You might also need to quickly adapt to more users if your company receives unexpected publicity from the media or some other means.

Scaling and *elasticity* allow you to easily deal with these kinds of scenarios. Scaling is the process of adding additional resources or additional power for your application. There are two variations of scaling: horizontal scaling (often referred to as *scaling out*) and vertical scaling (often referred to as *scaling up*).

When you scale out, you add additional VMs for your application. Each VM you add is identical to other VMs servicing your application. Scaling out provides additional resources to handle additional load.

When you scale up, you move to a new VM with

additional resources. For example, you may determine that you need a more powerful CPU and more memory for your application. In that case, scaling up will allow you to move your application to a more powerful VM.

Note Scaling up Often Adds Features

When you scale up, you often not only add more CPU power and memory, but you also often gain additional features because of the added power. For example, scaling up might give you solid-state disk drives or other features not available at lower tiers.

Figure 1-1 shows an example of scaling up a web application hosted in Azure.




Figure 1-1 Scaling up a web application in Azure

Real World Scaling Goes Both Ways

In addition to scaling out and scaling up, you can also scale in and scale down to decrease resource usage. In a real-world situation, you would want to increase computing resources when needed, reducing them when demand goes down.

Cloud providers make it easy to scale your application, and they offer the ability to scale automatically based on the usage pattern for your

application. You can scale automatically based on things like CPU usage and memory usage, and you can also scale based on other metrics that are specific to the type of application. The concept of automatically scaling is referred to as *elasticity*.



Exam Tip

In Azure, you can scale automatically by configuring Auto-Scale. Auto-Scale is an Azure service that can automatically scale applications running in many Azure services based on usage patterns, resource utilization, time of day, and much more.

One of major benefits of the cloud is that it allows you to quickly scale. For example, if you are running a web application in Azure and you determine that you need two more VMs for your application, you can scale out to three VMs in seconds. Azure takes care of allocating the resources for you. All you have to do is tell Azure how many VMs you want and you're up and running. This kind of speed and flexibility in the cloud is often called cloud *agility*.

[More Info](#) [More Information on Scaling Best Practices](#)

For more information on scaling in Azure, see the documentation at: <https://docs.microsoft.com/azure/architecture/best-practices/auto-scaling>.

Fault tolerance and disaster recovery

In a complex cloud environment, things are bound to go wrong from time to time. In order to maintain a high level of availability, cloud providers implement systems that monitor the health of cloud resources and take action when a resource is determined to be unhealthy, thereby ensuring that the cloud is *fault tolerant*.



Exam Tip

Don't confuse fault tolerance with scaling. Scaling allows you to react to additional load or resource needs, but it's always assumed that all of the VMs you are using are healthy. Fault tolerance happens without any interaction from you, and it's designed to automatically move you from an unhealthy system onto a healthy system in the event that things go wrong.

In addition to monitoring the health of VMs and

other resources, cloud providers design their infrastructure in such a way as to ensure fault tolerance. For example, if you have an application running on two VMs in Azure, Microsoft ensures that those two VMs are allocated within the infrastructure so that they are unlikely to be impacted by system failures.

More Info Fault Tolerance in Azure

You don't have to understand the technical details of how Azure implements fault tolerance for the AZ-900 exam, but if you're interested in learning more, check out:
<https://msdn.microsoft.com/magazine/mt422582.aspx>.

Fault tolerance is designed to deal with failure at a small scale; moving you, for example, from an unhealthy VM to a healthy VM. However, there are times when much larger failures can occur. For example, natural disasters in a region can impact all resources in that particular region. Not only can something like that impact availability, but without a plan in place, disasters can also mean the loss of valuable data.

Real World Disaster Recovery and Governments

Depending on what kind of data you store, you may be required to have a disaster recovery plan in place. Cloud providers typically comply with standards imposed by laws such as HIPAA, and they often provide compliance tools you can use to ensure compliance. You'll learn more about compliance and Azure in Chapter 3, "Understand security, privacy, compliance, and trust."

Disaster recovery not only means having reliable backups of important data, but it also means that the cloud infrastructure can replicate your application's resources in an unaffected region so that your data is safe and your application availability isn't impacted. Disaster recovery plans are commonly referred to as *Business Continuity and Disaster Recovery* (BCDR) plans, and most cloud providers have services that can help you develop and implement a plan that works for your particular needs.

Economic benefits of the cloud

So far we've talked only about the availability benefit of moving to the cloud, but there are also economic benefits. Let's consider both the on-premises model and the cloud model.

On-Premises Model

In the on-premises model, a business purchases physical computer hardware to be used for its IT needs. Because these computers are physical assets that are intended to be used for more than one year, they are usually purchased as *capital expenses*.

There are several drawbacks to this model. When a business purchases computer hardware, it will typically keep that hardware in service until the return on that

investment is realized. In the fast-evolving environment of computers, that can mean that hardware is outdated long before it makes financial sense to replace it. Another major drawback to this method is that it is not an agile approach. It may take months to requisition and configure new hardware, and in the era of modern IT, that approach often makes no sense.

More Info Tying Up Money

Businesses need money for day-to-day operations, and when you have large amounts of money tied up in capital expenses, it can dramatically reduce the amount of money you can put toward your daily operations.

Cloud model

When you move to the cloud, you no longer rely on your on-premises computing hardware. Instead, you essentially rent hardware from the cloud provider. Because you aren't purchasing physical assets, you move your IT costs from capital expenses to *operating expenses*, or day-to-day expenses for your business. Unlike capital expenses, operating expenses are tracked on a month-by-month basis, so it's much easier to adjust them based on need.

Another major benefit of the cloud model is reduced costs. When you use cloud resources, you are using resources made available from a large pool of resources

owned by the cloud provider. The cloud provider pays for these resources up-front, but because of the large scale of resources they purchase, the cost to the cloud provider is greatly reduced. The reduction in cost that is realized when purchasing large numbers of a resource is referred to as the *principle of economies of scale*, and those savings are passed on to consumers of the cloud.

Cloud providers take these savings a step further by offering the ability to use only those computing resources you require at any particular time. This is typically referred to as a *consumption-based model*, and it's often applied at many levels in cloud computing. As we've already discussed, you can scale your application to use only the number of VMs you need, and you can choose how powerful those VMs are. You can adjust their number and power as your needs require. However, many cloud providers also offer services that allow you to pay only for time that you consume computer resources. For example, you can have application code hosted in a cloud provider and pay only for time that the code is actually executing on a VM. When no one is using the application, you don't pay for any resources.

More Info Consumption-Based Computing

For an example of a consumption-based model, see Serverless computing in [Chapter 2](#), “[Understand core Azure services](#).”

As you can see, the cloud model offers many economic benefits over the on-premises model, and that's just one reason why businesses are rapidly moving to the cloud.

SKILL 1.2: DESCRIBE THE DIFFERENCES BETWEEN INFRASTRUCTURE-AS-A-SERVICE (IAAS), PLATFORM-AS-A-SERVICE (PAAS), AND SOFTWARE-AS-A-SERVICE (SAAS)


As you've learned, one of the benefits of moving to the cloud is that you offload some of the responsibility of your infrastructure to the cloud provider. Moving to the cloud, however, is not an all-or-nothing kind of thing. When you're evaluating your use of the cloud, you need to balance your need for controlling resources against the convenience of allowing the cloud provider to handle things for you.

Offerings in the cloud are typically referred to as *services*, and in this skill section, we're going to discuss the three primary types of cloud services:

Infrastructure-as-a-Service (IaaS), *Platform-as-a-Service (PaaS)*, and *Software-as-a-Service (SaaS)*. Each type of service comes with advantages and disadvantages, and the easiest way to visualize them is by using the cloud pyramid as shown in [Figure 1-2](#). The

bottom of the cloud pyramid represents the greatest amount of control over your resources, but it also represents the greatest amount of responsibility on your part. The top of the pyramid represents the least amount of control, but also the least amount of responsibility.

Less Control



More Control




Figure 1-2 The cloud pyramid

This section covers:

- Infrastructure-as-a-Service (IaaS)
- Platform-as-a-Service (PaaS)
- Software-as-a-Service (SaaS)
- Comparing service types

Infrastructure-as-a-Service (IaaS)

Infrastructure refers to the hardware that your application uses, and IaaS refers to the virtualized infrastructure offered by a cloud provider. When you create an IaaS resource, the cloud provider allocates a VM for your use. In some cases, the cloud provider might do the basic operating system install for you. In other situations, you may need to install the operating system yourself. In either case, you are responsible for installing other necessary services and your application.

Because you control the operating system install and installation of other services, IaaS gives you plenty of control over your cloud resources. However, it also means that you are responsible for making sure your operating system is patched with security updates, and if something goes wrong in the operating system, you're responsible for troubleshooting it. The cloud provider is only responsible for providing the VM. You do, however,

benefit from the underlying infrastructure in the area of fault tolerance and disaster recovery that we discussed earlier.

More Info Remote Access to IaaS VMs

You will have remote access to your IaaS VMs so that you can interact with them just as if you were using them in your on-premises environment. When you move to PaaS and SaaS services, you typically lose that capability because the infrastructure is managed by the cloud provider.

In [Figure 1-3](#), you see an IaaS VM in the Azure portal. The Ubuntu Server, a Linux operating system, has been chosen for the VM. Once the VM is up and running, it will be using Ubuntu Server 18.04. Unless an updated is installed, it will always be running that version. Microsoft will never install patches or version updates for me.

Create a virtual machine

[Basics](#) [Disks](#) [Networking](#) [Management](#) [Guest config](#) [Tags](#) [Review + create](#)

Create a virtual machine that runs Linux or Windows. Select an image from Azure marketplace or use your own customized image. Complete the Basics tab then Review + create to provision a virtual machine with default parameters or review each tab for full customization.
Looking for classic VMs? [Create VM from Azure Marketplace](#)

PROJECT DETAILS

Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

* Subscription ⓘ

Jim's Personal Azure Account ✓

* Resource group ⓘ

(New) AZ900 ✓

[Create new](#)

INSTANCE DETAILS

* Virtual machine name ⓘ

LinuxDocker ✓

* Region ⓘ

South Central US ✓

Availability options ⓘ

No infrastructure redundancy required ✓

* Image ⓘ

Ubuntu Server 18.04 LTS ✓

[Browse all images and disks](#)

* Size ⓘ

Standard D2s v3
2 vcpus, 8 GB memory
[Change size](#)

Review + create

Previous

Next : Disks >

Figure 1-3 Creating an IaaS VM in Azure

Once you have an IaaS VM running in the cloud, you

gain access to many services the cloud provider offers. For example, Microsoft offers Azure Security Center to ensure the security of your IaaS VMs, Azure Backup to make backing up data easy, Azure Log Analytics to help with troubleshooting any problems you might have, and much more.

More Info More Information On IaaS and Azure

For more information on IaaS and Azure, see the documentation at:
<https://azure.microsoft.com/overview/what-is-iaas/>.

IaaS services allow you to control costs effectively, because you only pay for them when you are using them. If you stop your IaaS VM, your billing stops for the resource. This makes IaaS an ideal choice if you need developers to have a platform for testing an application during release. Developers can start an IaaS VM, test the application as a team, and then stop the IaaS VM when testing is complete.

Another popular use of IaaS is when you need one or more powerful VMs for a temporary period. For example, you might need to analyze a large amount of data for a project. By utilizing IaaS VMs for your project, you can keep costs to a minimum, create resources quickly as you need them, and gain all the processing power you need.

IaaS services benefit from scaling and elasticity that

we discussed earlier. If you need more VMs, you can scale out to accommodate that and then scale in when those resources are no longer needed. If you need more CPU power, more memory, or more disk space, you can quickly scale up to gain those benefits and then scale down when they're no longer needed.

In a nutshell, IaaS services are a great choice if you want to let someone else manage the hardware infrastructure (which can include both the computers and the network) related to your application, but you want to maintain control of what's installed in the operating system. In an IaaS environment, the cloud provider isn't going to install something on the operating system for you, so the current state of what's installed on your VMs is always known to you. If this is important for your particular needs, IaaS may be the right choice for you. IaaS is also a great choice if you occasionally need high-end VMs for specific needs.

IaaS is also a great choice if you want your application and configuration in the cloud, but you want the option of not paying for it when you aren't using it. By stopping your VM, you can avoid the costs associated with it, and when you need to use your application again, you can simply start your VM and pick up right where you left off.

Platform-as-a-Service (PaaS)

In a PaaS environment, a cloud provider still provides the infrastructure for you, but they also provide the operating system, software installed in the operating system to help you connect to databases and network systems (often referred to as *middleware*), and many features that enable you to build and manage complex cloud applications.

PaaS sits right in the middle of the cloud pyramid. PaaS services offer you the flexibility of controlling the application, but they offload management and control of the underlying systems to the cloud provider. If you are deploying your own application to the cloud and you want to minimize your management investment, a PaaS service is often the best choice.

Suppose you need to run a web application that uses the PHP framework to connect to a back-end database system. If you were to choose IaaS for your application, you'd need to ensure that you install and configure PHP on your VM. You'd then need to install and configure the software necessary to connect to your back-end database. In a PaaS scenario, you simply deploy your web application to the cloud provider, and everything else is taken care of for you.

In [Figure 1-4](#), we have a web application in Azure App Service, one of the PaaS offerings in Azure. It has been created on a VM that's maintained by Microsoft. Notice

the option of choosing either Linux or Windows, but the operating system is still managed by Microsoft. We also have the option of enabling Application Insights, a service in Azure that provides deep insight into how an application is performing, making it easier to troubleshoot problems if they occur.

Dashboard > New > Web App > App Service plan

Web App

Create

* App name

jimsphapp

azurewebsites.net

* Subscription

Jim's Personal Azure Account

* Resource Group

Create new

Use existing

ExamRefRG

* OS

Windows

Linux

* Publish

Code

Docker Image

* App Service plan/Location

ServicePlande593f0f-a72b(Centr...

Application Insights

Disabled

Create

Automation options

App Service plan

Select a plan for the web app

An App Service plan is the container for your app. The App Service plan settings will determine the location, features, cost and compute resources associated with your app.

Create new

ServicePlande593f0f-a72b(S1)

Central US

1 instance(s)

Figure 1-4 Creating a Web App in Azure App Service

One more interesting thing in [Figure 1-4](#) is the option to publish either your code or a Docker image. Docker is a technology that makes it easy to package your application and the components that it requires into a *container* that you can then deploy and run on another computer in another environment, as long as that computer has Docker installed on it. In Azure App Service, I don't have to worry about Docker installation or configuration. It's automatically included on all App Service VMs as part of Microsoft's PaaS offering, and it's completely managed and maintained by Microsoft.

In a PaaS offering, cloud providers offer numerous application frameworks such as PHP, Node.js, ASP.NET, .NET Core, Java, Python, and more. The cloud provider usually provides multiple versions of each framework so you can choose a version that you know is compatible with your application. The cloud provider will also ensure that common components necessary for data connectivity from your application to other systems is installed and configured. That usually means that your application code works without you having to do any kind of complex configuration. In fact, this is one of the main benefits of using a PaaS service; you can often move your application from on-premises to a cloud environment by simply deploying it to the cloud.

This concept is often referred to as *lift-and-shift*.

Because the cloud provider controls the operating system and what's installed on the VM, they can provide additional capabilities to you by adding their own features. For example, suppose you want to add a log-in feature to your web application, and you want to allow users to log in with a Microsoft account, a Facebook account, or a Google account. If you wanted to add this capability on-premises, or in an IaaS environment, you need some developers to build it for you, a task that isn't easy and one that requires specialized knowledge. You'd have to either have developers in your company who already have those skills, or you'd have to hire them. However, cloud providers often offer features like this in their PaaS services, and enabling them is as easy as flipping a switch and doing some minor configuration specific to your app.

A PaaS service also benefits from all of the other enhancements offered by the cloud; you get fault tolerance, elasticity, easy and quick scaling, backup and disaster recovery features, and more. In fact, features such as backing up and restoring data are oftentimes more user-friendly and feature-rich in a PaaS environment because the cloud provider installs customized software on the PaaS VMs to add functionality.

As you can see, there are real benefits to allowing the cloud provider to control what's installed on the VMs running your application, but there can also be drawbacks. For example, the cloud provider controls when patches and updates are applied to both the operating system and to other components installed on the VMs. You'll usually be given advance notice of major changes so that you can test your application on-premises first and avoid any downtime, but you do lose the flexibility and control of deciding when to update the VM.

More Info More Information On Paas and Azure

For more information on PaaS offerings in Azure, see:
<https://azure.microsoft.com/overview/what-is-paas/>.

Software-as-a-Service (SaaS)

As you've learned, IaaS requires you to control both the operating system and middleware components along with your application. When you move to PaaS, you offload the control of the operating system and middleware components to the cloud provider, and you're responsible only for your application code. As you move to the top of the cloud pyramid and into the SaaS realm, the cloud provider controls everything. In other words, a SaaS service is software provided by a cloud provider that's installed on infrastructure completely

controlled by the hosting provider.

SaaS services offer you the flexibility of a pay-as-you-go model. Essentially, you rent your software from a service provider. Users of the software usually access the software from a web browser, but they may also install applications that will only work as long as you are paying for the SaaS service. One huge benefit of web-based software is that it works from just about any device, including smart phones. Because of that, SaaS services enable connectivity and productivity for field staff using devices they already own.

When using a SaaS service, not only do you benefit from using software written and maintained by someone else, but you can also benefit from allowing the cloud provider to maintain and configure the application. For example, if your company offers corporate email, you can choose to use Microsoft's Office 365 SaaS service. By using the Exchange Online service in Office 365, you can take advantage of enterprise-ready email solutions without having to hire IT staff and build infrastructure to support it. Instead, Microsoft maintains the system for you. Not only do you benefit from the flexibility and reliability of the cloud, but you can also rest easy knowing that Microsoft is ensuring your Exchange services are always available to your users.

SaaS services aren't just for the enterprise. In fact, most people use SaaS services all the time without even realizing it. If you use Hotmail or Gmail or another online email service, you're using a SaaS service. The cloud provider hosts the email software in the cloud, and you log in and use that software using your web browser. You don't have to know anything about the software. The cloud provider can offer new features with software updates, and those new features are available to you automatically without any action on your part. If the cloud provider finds a problem with the software, they can resolve it with a patch without you even realizing anything happened.

More Info More Information On SaaS and Azure

For more information on SaaS services and Azure, see:
<https://azure.microsoft.com/overview/what-is-saas/>.

Comparing service types

We've already discussed some of the advantages and disadvantages of each type of cloud service, and the cloud pyramid provides a visual representation of how types of cloud services differ related to your responsibility and what you can control. In order to solidify these concepts, let's look at a comparison of each service type.

As you've learned, IaaS provides you with the greatest

flexibility. You can install your own software and your own components, and you control when the software and operating system are updated. An additional benefit is that you pay for your resources only when they're being used, so IaaS has the ability to reduce your operational expenses. Even though you can save costs by turning off VMs you aren't using, the higher costs associated with installing and maintaining your VMs might offset that benefit.

PaaS services offer you some of the same flexibility of IaaS services without the need to manage the infrastructure. In a PaaS service, you are responsible only for the application that's installed in the cloud. This can be your own application, or an application developed by someone else (for example, a WordPress system or an e-commerce solution), but in either case, you are responsible for the application. PaaS services are popular for developer teams who are looking to move on-premises applications to the cloud easily and quickly, and they typically offer many different deployment options to make that as easy as possible. PaaS services also offer more features than IaaS services, because the cloud provider installs their own software and features on the platform. Any application running in a PaaS service, however, can be impacted by updates and version changes in the underlying software, and that can mean increased costs associated with

testing an application before the cloud provider rolls out changes.

SaaS services are quite a bit different than IaaS or PaaS services because they are completely managed and maintained by the cloud provider. You don't have the option of installing any of your own software with a SaaS service, so the deciding factor is related entirely to whether or not the provided-software meets your needs. The benefit of a SaaS service is that it largely removes the IT burden from your company, and it enables everyone in your company to access the software on multiple devices from just about anywhere Internet access is available. You also benefit from data backup that the cloud provider includes in their infrastructure. If you have a need to customize the application or have any control over its configuration, however, SaaS may not be a good choice for you.

Real World Dealing With the Complexities of Modern It

Deciding on a particular cloud service type can be straightforward in some cases, but it can also be complicated depending on your needs. For example, you might be in an industry that requires some of your information to be stored only on-premises. You might also have some older systems that aren't ready to move to the cloud, but you need your cloud applications to use those older systems. In the next skill section, you'll learn more about how to deal with such complexities.

SKILL 1.3: DESCRIBE THE DIFFERENCES BETWEEN PUBLIC,

PRIVATE, AND HYBRID CLOUD MODELS

In the simplest sense, the cloud represents infrastructure and applications that are accessible over the Internet. The examples covered so far are the more traditional cloud experience where anyone on the Internet can access your application. While you might have some means of authenticating people using your application so that the wrong people don't get access, your application is still running on VMs that are connected to the Internet and are accessible over public networks.

The traditional cloud model is referred to as the *public cloud*. In addition to a public cloud model, businesses can also use a *private cloud* where the infrastructure is dedicated to them. Finally, a *hybrid cloud* model represents a mixture of public and private cloud models.

More Info Community Clouds

You might see references to a fourth cloud model called the community cloud. A community cloud is similar to a private cloud, but instead of resources being dedicated to a single company, they are dedicated to a community of companies or individuals who manage it together. For example, hospitals might use a community cloud that's explicitly designed to handle the Health Insurance Portability and Accountability Act of 1996 (HIPAA) and other health care regulations. Financial institutions might also share a community cloud that enforces regulations and policy related to banks and financial trading.

Community clouds aren't part of the AZ-900 exam, but it's still important to understand what the term means in case you come across it while preparing for the exam.

This section covers:

- The public cloud
- The private cloud
- The hybrid cloud

The public cloud

The most common cloud model is the public cloud. In a public cloud model, you use shared infrastructure that is accessible on a public network. The network, storage, and VMs that your application uses are provided by a cloud provider and shared between all consumers of the public cloud. Microsoft Azure is an example of a public cloud.

The public cloud model is beneficial in that it makes it easy and fast to move to the cloud. Because the cloud provider already has the infrastructure in place and configured for you, all you have to do is decide on the type of cloud service you want and you're off and running. You also benefit from the ability to scale quickly and efficiently because the cloud provider has resources already provisioned and ready for your use when needed.

As we discussed earlier, another advantage to the public cloud model is that you can control costs more

efficiently because you only pay for the resources you are using. If you need to scale out to more VMs, the cloud provider has them available and waiting for you. You don't have to maintain a pool of resources yourself. Instead, you take advantage of the resources the cloud provider has invested in.

Important Multi-Tenant Environment

Because you are sharing resources in a public cloud with other people who are using that public cloud, you'll often see public clouds referred to as a multi-tenant environment.

While the flexibility and convenience of the public cloud is attractive, it comes with some disadvantages. First of all, you do give up some control of the infrastructure when using the public cloud. How much control depends on where you land on the cloud pyramid, but no matter what, the cloud provider is going to control some portion of your infrastructure.

There may also be security concerns with operating in the public cloud. The network involved in the public cloud is the public Internet, and it's available to anyone with an Internet connection. That means you will need to have security measures in place to avoid unauthorized access to your application and data. Cloud providers realize this, and they provide security measure to help protect you, but those measures may

not meet your security requirements.

Another disadvantage of the public cloud is that it locks you into the specific configuration defined by the cloud provider. For example, suppose you have an application that needs a large amount of disk storage, but you only need a single-CPU system to run it. In order to meet your disk space requirements, the cloud provider might require you to scale up to a high-powered, multi-CPU VM, thereby increasing your costs unnecessarily.

More Info More Information On Public Clouds

For more information on public clouds and Azure, see:
<https://azure.microsoft.com/overview/what-is-a-public-cloud/>.

The private cloud

The private cloud model provides many of the attractive benefits of the cloud (things like easy scaling, and elasticity) in a private environment that is dedicated to a single company. A private cloud can be hosted in an on-premises environment, but it can also be hosted on a third-party hosting provider.

Important Single-Tenant Environment

Because the resources in a private cloud are dedicated to a single organization, you will often see the private cloud referred to as a single-tenant environment.

Two of the main reasons why companies choose a private cloud are: privacy and regulatory concerns. Unlike the public cloud, private clouds operate on a private network that is only accessible by a single organization. Businesses like banks and medical providers may have regulations in place that require certain data be inaccessible from the Internet, and in those situations, a private cloud might be a good choice. Another common consumer of private clouds is the cruise ship industry. Cruise ships operate in remote areas where Internet access isn't available, but they still want to take advantage of the benefits of the cloud for day-to-day operations of complex ship systems.



Exam Tip

You'll often hear that a private cloud consists of infrastructure that is owned by an individual company, but that's not actually always true. If a company runs a private cloud on-premises, they will usually own the hardware and infrastructure used for the private cloud, but it's also possible to host a private cloud in a third-party data center. In that situation, the infrastructure is owned by the hosting

provider, but it's still completely dedicated to the single company paying for the private cloud.

The bottom line is that the difference between a public and a private cloud is the privacy of infrastructure and data. It doesn't really matter who owns the infrastructure

There are some disadvantages to a private cloud. If you are hosting your private cloud on-premises, you will likely spend as much on IT as you would in a non-cloud environment. You will have to pay for hardware and virtualized systems for your cloud, and you'll need IT staff who are capable of managing the software and infrastructure for your cloud.

Avoiding IT costs is one of the primary reasons that companies choose to use a third-party hosting provider for private clouds, but that choice also has some drawbacks. For example, once you offload management of your private cloud to a third-party, you lose control of important considerations, such as the security of your data. It's often impossible to achieve full transparency when dealing with third-party providers, and you can't always guarantee that data on your private cloud network will remain secured in a way that you require.

More Info More Information On Private Clouds

For more information on private clouds, see:

<https://azure.microsoft.com/overview/what-is-a-private-cloud/>.

The hybrid cloud

As you might expect, hybrid clouds are a mixture of public and private clouds. In a hybrid cloud environment, you may have an application that is running within the public cloud, yet it accesses data that is securely stored on-premises. You might also have a scenario where your application and most of its resources are located on a private cloud, but you want to use services or infrastructure that are located in a public cloud. Indeed, the various scenarios that are suitable for a hybrid model are almost endless.

Hybrid cloud models are often a company's first foray into the cloud. Many companies have legacy on-premises systems that are expensive to move to the cloud, yet you may want to take advantage of some of the benefits of the cloud. In such a scenario, a company might move only part of a particular system to the cloud, leaving the legacy system on-premises until a later time.

Not all companies adopting a hybrid cloud model are doing so because of legacy systems. In some situations, a company may want to maintain complete control over

part of their infrastructure or data. They may decide to build out on-premises infrastructure in tandem with building their public cloud presence.

Important Hybrid Doesn't Always Include On-Premises

Remember, a private cloud is a cloud dedicated to a single organization. It doesn't have to be located on-premises. It can also be hosted at a third-party data center, so a hybrid cloud model might be the combination of a third-party data center and a public cloud.

When companies adopt a hybrid model, they often require the capability of connecting the private, on-premise network with the public cloud network. Cloud providers offer many technologies to make that possible. In Microsoft Azure, Virtual Networks, Hybrid Connections, and Service Bus are just some examples of such technologies.

More Info More Information On Azure Network Offerings

We'll cover some of the Azure networking offerings in [Chapter 2, Skill 2.2](#).

While it might not be immediately obvious, a hybrid cloud model comes with several challenges. First of all, application development teams will need to ensure that data shared between the public and private cloud is compatible. This might require some specialized development skills and complex troubleshooting. The

networking complexities in a hybrid environment can also be quite challenging, especially because network infrastructure at third-party providers may introduce problems that are difficult to troubleshoot. Finally, spreading application resources between a public and a private cloud may cause application slowdowns due to the geographical distance between systems running the application and the data the application uses. All of these situations have to be carefully evaluated when deciding to use a hybrid cloud model.

In order to make hybrid cloud easier for its customers, Microsoft provides Azure Stack. Azure Stack is sold as a package, including software and validated hardware to run it. Azure Stack allows you to run Azure services on-premises, making it easy to then transfer applications to the cloud with a minimal amount of work. Because the hardware is part of Azure Stack and has been validated by Microsoft, you don't have the burden of attempting to determine hardware needs in order to deploy Azure Stack, but you do have to manage the on-premises hardware.

THOUGHT EXPERIMENT

Let's apply what you've learned in this chapter. You can find the answers in the section that follows.

You work for Contoso Medical Group (CMG), and your manager is frustrated with one of your commonly-

used applications. The CMG IT department is resource-constrained, and they are having difficulty ensuring the application is always available.

The development team has been updating the application frequently, but due to a lack of knowledge in deployment methods, they only have the option of directly copying files, and this is causing problems with tracking changes that are being made. At the same time, the development team has no data to show whether the application is running correctly.

The problem became critical two days ago when a deadline was approaching for updating medical records. The application experienced way more usage than normal, and the system was quickly overloaded and became unresponsive. The IT team determined the problem was the server running low on resources, but it took them two hours to build a second server to handle the load.

Your manager has come to you asking for a solution that addresses all of these issues. Whatever solution you offer must take into account that the medical data in this application is covered under HIPAA, and your manager wants CMG to retain all control of the data. Your manager also wants to carefully control costs.

You've decided that CMG should move the application to the cloud, but you need to sell the idea to

your manager.

Answer the following questions:

1. What type of cloud service would you recommend?
2. How would you justify your choice related to the problems being encountered by the IT team?
3. How would you justify your choice related to the problems being encountered by the development team?
4. What other benefits will please your manager if your advice is followed?
5. How can you meet the requirements related to the medical records and the need to control them?

THOUGHT EXPERIMENT ANSWERS

In this section, we'll discuss the answers from the previous section.

1. A PaaS service makes the most sense in this situation. An IaaS environment would require your IT department to manage the VMs, and that would not meet your requirements. A SaaS service provides the software to you, and in this case, you need to run your company's custom application in the cloud.

2. The IT department is short on resources and is challenged in keeping the application available. In a PaaS service, the management of the VMs running the application is offloaded to the cloud provider. The cloud provider also offers an SLA so that your application is always available. The IT team will also benefit from easy scaling offered in a cloud environment, and instead of two hours, they can add more servers almost instantly.
3. In a PaaS service, the cloud provider offers flexible deployment options that make it easy to deploy an application using the method you prefer. They also provide logging so that the development team can track changes made to the application. Diagnostic features in a PaaS service (such as Azure's Application Insights) provide detailed data on how an application is performing and can alert you to code problems in an application.
4. Your manager wants to lower costs, and moving to the cloud should meet that need. Your IT department has already built a second server, so that when additional need is required, you can meet it. However, the increased usage was temporary. Even so, it was related to a deadline

for filing records, and the next time that deadline occurs, you'll need that second server. By moving to the cloud, you benefit from easy scaling and elasticity so that you can scale out when you need the second server to handle load, and then you can easily scale back in to reduce your costs.

5. By adopting a hybrid cloud model, you can keep your sensitive medical data on-premises, while benefiting from the application itself running in the cloud.

CHAPTER SUMMARY

In this chapter, you learned some of the general concepts related to the cloud. You learned about the advantages of moving to the cloud, you learned about the different cloud service types, and you learned about the different cloud models available to you. Here are the key concepts from this chapter.

- Cloud providers offer service-level agreements (SLAs) that guarantee a certain level of availability, but only for those systems that are controlled by them.
- Moving to the cloud can help avoid downtime caused by network outages, system outages, and power outages. It can also help you if you need to diagnose problems with an application or problems with an external system that your application uses.
- You can scale up (or vertically) when you want to add additional CPUs or more memory using a more powerful VM.

- You can scale out (or horizontally) if you want to add more VMs to handle additional load.
- Cloud providers give you ways to automatically scale based on usage patterns, resource utilization, and times of day. This is referred to as *elasticity*.
- Cloud providers monitor the health of the infrastructure. When a VM becomes unhealthy, the cloud provider can automatically move you to a healthy VM without you having to do anything. This is called *fault tolerance*.
- Cloud providers also operate across multiple data centers that are in different regions of the world. If a natural disaster (or any other disaster) happens in one region, you can switch over to another region, assuming you have replicated your environment in multiple regions. This kind of planning is called Business Continuity and Disaster Recovery planning, and cloud providers often have features in place to make implementing a plan easy. This is often referred to as disaster recovery.
- Because you are using infrastructure owned by the cloud provider, moving to the cloud reduces your *capital expenses*, the major expenses that are incurred for infrastructure and other major purchases. Cloud providers take advantage of the *principle of economies of scale* by purchasing large amounts of infrastructure to be used by cloud consumers.
- Day-to-day expenses (*operational expenses*) can also be reduced in the cloud because you pay only for those resources you are using at any particular time. This *consumption-based model* is a key benefit of the cloud.
- Infrastructure-as-a-Service (IaaS) offers infrastructure running in the cloud, but you have to maintain the operating system and what's installed on that infrastructure. IaaS services offer you the most control in the cloud, but they also carry the largest management burden.
- Platform-as-a-Service (PaaS) offloads the management of the

infrastructure, and it also offloads the operating system and components installed on the VMs to the cloud provider. You are responsible for your application. PaaS services also offer many additional features that make it easy to add functionality to an application without having to write complex code. Development teams also have a wide variety of deployment methods available, and the cloud provider often automates much of that process.

- Software-as-a-Service (SaaS) provides a hosted application in the cloud that is most commonly accessed using a web browser. In a SaaS service, the cloud provider manages everything for you. You are essentially renting the use of the software from the cloud provider. A big benefit of SaaS is that it makes applications easily-accessible by employees in the field on any device.
- The public cloud model is sometimes referred to as a multi-tenant environment. Multiple companies and users share the same infrastructure. VMs and other infrastructure are allocated to users as they need them, and when they no longer need them, they are returned to the pool to be used by other users. The network is available publicly over the Internet, but you do have the ability to put security methods in place to control access to your resources.
- The private cloud model is sometimes referred to as a single-tenant environment. All infrastructure is private to an individual or a company, and the network is only available within the private cloud itself. It is not exposed to the Internet. In many cases, the infrastructure used in a private cloud is owned by the company, but not always. It's possible to host a private cloud in a third-party data center.
- A hybrid cloud model is a mixture of the public and private cloud models. Hybrid clouds are often used when a company needs to use on-premises resources in a cloud application.

Chapter 2. Understand core Azure services

In Chapter 1, “Understand cloud concepts,” you learned about the cloud and how you can benefit from using cloud services. Microsoft Azure was mentioned, but not in a lot of detail.

In this chapter, we dive into the many services and solutions that Azure offers. You’ll gain an understanding of the key concepts in Azure’s architecture, which apply to all Azure services. We cover Azure datacenters and ways that Microsoft implements fault tolerance and disaster recovery by spreading Azure infrastructure across the globe. You’ll also learn about availability zones, which are Microsoft’s solution for ensuring your services aren’t impacted when a particular Azure datacenter experiences a problem.

You’ll also discover how to manage and track your Azure resources, and how you can work with resources as a group using Azure resource groups. You’ll learn how to use resource groups to not only plan and manage Azure resources, but also how resource groups can help you categorize your operational expenses in

Azure.

In order to really understand resource groups and how Azure works under the hood, it's important to understand Azure Resource Manager (ARM), the underlying system that Azure uses to manage your resources. You'll learn about the benefits that ARM provides, and you'll see how ARM opens up some powerful possibilities for quickly and easily deploying real-world solutions to Azure.

Once you have the foundational understanding of Azure, you'll dig into some of the core products that Microsoft provides, such as Azure Compute, networking, storage, and database offerings, which are covered from an Azure perspective. You'll learn about some of the products available in each of these areas, and you'll get a feel for how Azure products work together. Along the way, you'll learn about the Azure Marketplace and how it enables the creation and deployment of complex solutions with minimal work on your part, and because of the “under the hood” knowledge you'll have from earlier in the chapter, the Azure Marketplace won't seem like black magic.

You'll even learn about some of the hottest technology areas today and what Azure has to offer in those areas. This includes the Internet of Things (IoT) and how you use Azure to connect and manage devices

of all kinds. Azure can help you analyze huge amounts of data using big data and analytics products, and you'll learn how these offerings can help you control costs.

One of the hot technologies right now is artificial intelligence, or AI. Azure offers a comprehensive AI platform that includes some powerful machine learning components, and we'll talk about what Azure offers in this area and how you can use AI and machine learning to create powerful and insightful solutions. We'll wrap up with coverage of serverless computing in Azure and how you can create powerful and flexible services in Azure without spending a lot of money, and often without spending anything at all!

In addition, you'll learn about the tools that Microsoft offers for creating and managing your Azure services, including the Azure portal, which is a web browser-based management tool that offers great tools for digging into your Azure resources and easily managing them. We also cover how to use command-line tools with PowerShell and the Azure command-line interface. And, we'll wrap everything up with a look at Azure Advisor, Microsoft's service that gives you best-practices advice for your Azure services.

If you think that's a lot to cover, you're right! It's important for you to have an understanding of all of these topics in order to pass the AZ-900 exam. With the

foundational knowledge of the cloud from Chapter 1, you'll find that understanding Azure-specific concepts will be easier than you think.

Skills covered in this chapter:

- Understand the core Azure architectural components
- Describe some of the core products available in Azure
- Describe some of the solutions available on Azure
- Understand Azure management tools

SKILL 2.1: UNDERSTAND THE CORE AZURE ARCHITECTURAL COMPONENTS

If you were to ask any CEO to list the five most important assets of their company, it is likely that the company's data would be near the top of the list. The world we live in revolves around data. Just look at companies like Facebook and Google. These companies offer services to us that we like. Everyone likes looking at pictures from friends and family on Facebook (mixed in with things that we don't like so much), and who doesn't use Google to look for things on the Internet? Facebook and Google don't offer those services because they want to be nice to us. They offer those services because it's a way for them to collect a large amount of data on their customers, and that data is their most valuable asset.

Facebook and Google aren't alone. Most companies

have vast amounts of data that is key to their business, and keeping that data safe is at the cornerstone of business decisions. That's why many companies are hesitant to move to the cloud. They're afraid of losing control of their data. Not only are they afraid that someone else might gain access to sensitive data, but they're also concerned about losing data that would be difficult (or even impossible) to recreate.

Microsoft is keenly aware of those fears, and Azure has been designed from the ground up to instill confidence in this area. Let's look at some core architectural components that help Microsoft deliver on the cloud promise.

This section covers:

- Azure regions
- Availability zones
- Azure Resource Manager (ARM)
- Resource groups

Azure regions

The term “cloud” has a tendency to make people think of Azure as a nebulous entity that you can't clearly see, but that would be a mistake. While there certainly are logical constructs to Azure, there are also physical components to it. After all, at the end of the day, we're

talking about computers!

In order to provide Azure services to people around the world, Microsoft has created boundaries called geographies. A geography boundary is oftentimes the border of a country, and there's good reason for that. There are often regulations for data handling that apply to an entire country, and having a geography defined for a country allows Microsoft to ensure that data-handling regulations are in place. Many companies (especially ones that deal with sensitive data) are also much more comfortable if their data is contained within the confines of the country in which they operate.

There are numerous geographies in Azure. For example, there's a United States geography, a Canada geography, a UK geography, and so on. Each geography is broken out into two or more regions, each of which is typically hundreds of miles apart. As an example, within the United States geography, there are many regions, including the Central US region in Iowa, the East US region in Virginia, the West US region in California, and the South Central US region in Texas. Microsoft also operates isolated regions that are completely dedicated to government data due to the additional regulations that governmental data requires.



Exam Tip

The fact that each geography contains at least two regions separated by a large physical distance is important. That's how Azure maintains disaster recovery, and it's likely this concept will be included on the exam. We'll cover more about this later in this chapter.

At each region, Microsoft has built datacenters (physical buildings) that contain the physical hardware that Azure uses. These datacenters contain climate-controlled buildings that house the server racks containing physical computer hardware. They also have complex and reliable network infrastructure to provide the networking power.

More Info Customers Only See Regions

When a customer is creating Azure resources, only the region is visible. The concept of geographies is an internal implementation of Azure that customers don't really have visibility of when using Azure.

Each datacenter has an isolated power supply and power generators in case of a power outage. All of the network traffic entering and exiting the datacenter goes over Microsoft's own fiber-optic network, on fiber owned or leased by Microsoft. Even data that flows

between regions across oceans travels over Microsoft's fiber-optic cables that traverse the oceans.

More Info Datacenter Power

As of 2018, all of Microsoft's datacenters were using at least 50% natural power consisting of solar power, wind power, etc. By 2020, the goal is 60%, and the long-term goal is to use 100% sustainable power.

In order to remove reliance on third-party power providers, Microsoft is also investing in the development of natural gas-powered, fully-integrated fuel cells for power. Not only do fuel cells provide clean power, but they also remove the power fluctuations and other disadvantage of relying on the power grid.

To ensure that data in Azure is safe from disasters and failures due to possible problems in a particular region, customers are encouraged to replicate data in multiple regions. If, for example, the South Central US region is hit by a devastating tornado (not out of the question in Texas), data that is also replicated to the North Central US region in Illinois is still safe and available. In order to ensure that applications are still performing as quickly as possible, Microsoft guarantees round-trip network performance of 2-milliseconds or less between regions.

Availability zones

The fact that regions are physically separated by hundreds of miles protects Azure users from data-loss and application outages due to disasters at a particular region. However, it's also important that data and

applications maintain availability when a problem occurs at a particular datacenter within a region. For that reason, Microsoft developed availability zones.

Note Availability Zone Availability

Availability zones aren't available in all Azure regions. For the most up-to-date list of availability zone-enabled regions, see: <https://docs.microsoft.com/azure/availability-zones/az-overview>.

There are at least three availability zones within each enabled region, and because each availability zone exists within its own datacenter in that region, each has a water supply, cooling system, network, and power supply that is isolated from other zones. By deploying an Azure service in two or more availability zones, you can achieve high-availability in a situation where there is a problem in one zone.



Exam Tip

Availability zones provide high-availability and fault tolerance, but they may not help you with disaster recovery. If there is a localized disaster, such as a fire in a datacenter housing one zone, you will benefit from availability zones. Because availability

zones are located in the same Azure region, if there is a large-scale natural disaster such as a tornado, you may not be protected. In other words, availability zones are just one facet to an overall disaster recovery and fault tolerant design.

Because Availability zones are designed to offer enhanced availability for infrastructure, not all services support availability zones. For example, Azure has a service called App Service Certificate that allows you to purchase and manage an SSL certificate through Azure. It wouldn't make any sense to host an App Service Certificate within an availability zone because it's not an infrastructure component.

As of right now, availability zones are supported with the following Azure services.

- Windows Virtual Machines
- Linux Virtual Machine
- Virtual Machine Scale Sets
- Managed Disks
- Load Balancer
- Public IP address
- Zone-redundant storage
- SQL Database

- Event Hubs
- Service Bus (Premium tier only)
- VPN Gateway
- ExpressRoute
- Application Gateway (currently in preview)
- App Service Environments (currently in preview in limited regions)

Note Keep Up With Changes in Azure

You can keep up with all the news related to Azure updates by watching the Azure blog at <https://azure.com/blog>.

By deploying your service to two or more availability zones, you ensure the maximum availability for that resource. In fact, Microsoft guarantees a service level agreement (SLA) of 99.99% uptime for Azure Virtual Machines only if two or more VMs are deployed into two or more zones. Figure 2-1 illustrates the benefit of running in multiple zones. As you can see, even though availability zone 3 has gone offline for some reason, zones 1 and 2 are still operational.




Figure 2-1 Azure Virtual Machine inside of three availability zones



Exam Tip

Don't confuse availability zones with availability sets. Availability sets allow you to create two or more virtual machines in different physical server

racks in an Azure datacenter.

Microsoft guarantees a 99.95% SLA with an availability set.

An availability zone allows you to deploy two or more Azure services into two distinct datacenters within a region. Microsoft guarantees a 99.99% SLA with availability zones.

There are two categories of services that support availability zones: *zonal* services and *zone-redundant* services. Zonal services are services such as virtual machines, managed disks used in a virtual machine, and public IP addresses used in virtual machines. In order to achieve high-availability, you must explicitly deploy zonal services into two or more zones.

Note Managed Disks and Public Ip Addresses

When you create a virtual machine in Azure and you deploy it to an availability zone, Azure will automatically deploy the managed disk(s) and public IP address (if one is configured) to the same availability zone automatically.

Zone-redundant services are services such as zone-redundant storage and SQL Databases. To use availability zones with these services, you specify the option to make them zone-redundant when you create

them. (For storage, the feature is called ZRS or zone-redundant storage. For SQL Database, there is an option to make the database zone-redundant.) Azure takes care of the rest for you by replicating data to automatically multiple availability zones.

Azure Resource Manager (ARM)

Almost all systems that are moved to the cloud consist of more than one Azure service. For example, you might have an Azure virtual machine for one part of your app, your data might be in an Azure SQL Database, you might have some sensitive data stored in Azure Key Vault, and you might have a web-based portion of your app hosted in Azure App Service.

If you have to manage all of these different Azure services separately, it can be quite a headache, and if you have multiple applications in the cloud, it can be even worse. Not only would it be confusing to keep track of which services are related to which applications, but when you add in the complexity of deploying updates to your application, things can really become disorganized.

In order to make it easier to deploy and manage Azure services, Microsoft developed Azure Resource Manager, or ARM. ARM is a service that runs in Azure, and it's responsible for all interaction with Azure services. When you create a new Azure service, ARM

authenticates you to make sure you have the right access to create that resource, and then it talks to a *resource provider* for the service you're creating. For example, if you're creating a new web app in Azure App Service, ARM will pass your request on to the Microsoft.Web resource provider, because it knows all about web apps and how to create them.



Exam Tip

There are resource providers for every Azure service, but the names might not always make sense. For example, the Microsoft.Compute resource provider is responsible for creating virtual machine resources.

You don't have to know details on resource providers for the AZ-100 exam, but you should understand the general concept, because you are expected to know about Azure Resource Manager.

Later in this chapter, you'll learn about using the Azure portal to create and manage Azure services. You'll also learn about how you can use command-line tools to

do the same thing. Both the portal and the command-line tools work by using ARM, and they interact with ARM using the ARM application programming interface, or API. The ARM API is the same whether you're using the portal or command-line tools, and that means you get a consistent result. It also means that you can create an Azure resource with the portal and then make changes to it using command-line tools, allowing you the flexibility that cloud consumers need.

More Info Visual Studio and ARM

Visual Studio, Microsoft's development environment for writing applications, also has the ability to create Azure resource and deploy code to them. It does this using the same ARM API that tools we've mentioned use. In fact, you can think of the ARM API as your interface into the world of Azure. You really can't create or manage any Azure services without going through the ARM API.

The flow of a typical ARM request to create or manage a resource is straightforward. A tool such as the Azure portal, command-line tools, or Visual Studio makes a request to the ARM API. The API passes that request to ARM where the user is authenticated and authorized to perform the action. ARM then passes the request to a resource provider, and the resource provider creates the new resource or modifies an existing resource. Figure 2-2 illustrates this flow and features a small sampling of the many Azure services that are available.

TOOLS



Azure Portal



Command Line



Visual Studio



Azure Resource Manager API



Azure Resource Manager (ARM)



Resource Provider



AZURE RESOURCES

Figure 2-2 Azure Resource Manager

The request that is made to ARM isn't a complicated, code-based request. Instead, ARM uses *declarative syntax*. That means that, as a consumer of Azure, you tell ARM what you want to do and ARM does it for you. You don't have to tell ARM *how* to do what you want. You simply have to tell it what you want. To do that, ARM uses files that are encoded in JavaScript Object Notation (or JSON) called *ARM templates*.

Note ARM Templates

You don't need to know how to use ARM templates for the AZ-900 exam, but in order to grasp how ARM works, you really need to at least know a little about them.

In the most basic sense, an ARM template contains a list of resources that you want to either create or modify. Each resource is accompanied by properties such as the name of the resource and properties that are specific to that resource. For example, if you were using an ARM template to deploy a Web App in App Service, your ARM template would specify the region you want your app to be created in, the name of the app, the pricing plan for your app, any domain names you want your app to use, and so forth. You don't have to know how to set all those properties. You simply tell ARM to do it (you declare your intent to ARM), and ARM takes

care of it for you.

More Info More On Arm Templates

ARM templates are incredibly powerful, but they're also pretty simple. If you want to read more about how to use ARM templates, check out the documentation at: <https://docs.microsoft.com/azure/azure-resource-manager/resource-group-authoring-templates>.

There's one more important aspect to ARM template deployment. When you're deploying multiple resources (which, as pointed out, is a typical real-world scenario), you often have service dependencies. In other words, you are deploying one or more services that rely on another services already being created.

Think, for example, of a situation where you're deploying a certificate to be used with a web app. One of the properties you need to set on the web app is the certificate that you want to use, but if that certificate hasn't been deployed yet, your deployment will fail. ARM allows you to specify dependencies so you can avoid issues like this. You simply tell ARM that the web app depends on the certificate and ARM will ensure the certificate's deployment is completed before it deploys the web app.

As you can see, ARM has many benefits, and you should be aware of these for your exam:

- ARM allows you to easily deploy multiple Azure resources at

once.

- ARM makes it possible to reproduce any deployment with consistent results at any point in the future.
- ARM allows you to create declarative templates for deployment instead of requiring you to write and maintain complex deployment scripts.
- ARM makes it possible to set up dependencies so that your resources are deployed in the right order every time.

Now let's talk about another aspect of ARM that helps you to manage Azure resources, and that's resource groups.

Resource groups

You should now be realizing that moving to the cloud may not be as simple as it first seemed. Creating a single resource in Azure is pretty simple, but when you're dealing with enterprise-level applications, you're usually dealing with a complex array of services. Not only that, but you might be dealing with multiple applications that use multiple services, and they might be spread across multiple Azure regions. Things can certainly get chaotic quickly.

Fortunately, Azure provides a feature in ARM that helps you deal with this kind of problem :the resource group. A resource group is a logical container for Azure services. By creating all Azure services associated with a particular application in a single resource group, you

can then deploy and manage all of those services as a single entity.

Organizing Azure resources in a resource group has many advantages. First of all, you can easily set up deployments using an ARM template. ARM template deployments are typically for a single resource group. You can deploy to multiple resource groups, but doing so requires you to set up a complicated chain of ARM templates.

Another advantage to resource groups is that you can name a resource group with an easily-recognizable name so that you can see all Azure resources used in a particular application at a glance. This might not seem so important until you actually start deploying Azure resources and realize that you have many more resources than you first thought. For example, when you create an Azure Virtual Machine, Azure creates not only a virtual machine, but also a disk resource, a network interface, a public IP resource, and a network security group. If you're looking at all your Azure resources, it can be hard to differentiate which resources go with which app. Resource groups solve that problem.

In [Figure 2-3](#), you can see a lot of Azure services. Some of these were automatically created by Azure in order to support other services, and in many cases,

Azure gives the resource an unrecognizable name.

☐	NAME	TYPE	RESOURC...	LOCATION	SUBSCRI...
☐	 900rgdiag	Storage acc...	900RG	South Centr...	Jim's Perso...
☐	 900RG-vnet	Virtual netw...	900RG	South Centr...	Jim's Perso...
☐	 EComVM	Virtual mac...	WebStorefr...	South Centr...	Jim's Perso...
☐	 EComVM_OsDisk_1_1d...	Disk	WEBSTORE...	South Centr...	Jim's Perso...
☐	 ecomvm34	Network int...	WebStorefr...	South Centr...	Jim's Perso...
☐	 EComVM-ip	Public IP ad...	WebStorefr...	South Centr...	Jim's Perso...
☐	 EComVM-nsg	Network sec...	WebStorefr...	South Centr...	Jim's Perso...
☐	 greatappalready	App Service	Test	Central US	Jim's Perso...
☐	 jwc900	SQL server	WebStorefr...	Central US	Jim's Perso...
☐	 900StoreDB (jwc900/...	SQL database	WebStorefr...	Central US	Jim's Perso...
☐	 ServicePlan9dbd216e-...	App Service ...	WebStorefr...	Central US	Jim's Perso...
☐	 UbuVM	Virtual mac...	900RG	South Centr...	Jim's Perso...
☐	 UbuVM_OsDisk_1_973...	Disk	900RG	South Centr...	Jim's Perso...
☐	 ubuvm97	Network int...	900RG	South Centr...	Jim's Perso...

Figure 2-3 All my Azure resources

In Figure 2-4, you can see resources that are in the WebStorefront resource group. These are the Azure resources used in the e-commerce storefront.

Dashboard > Resource groups > WebStorefront

WebStorefront
Resource group

+ Add
Edit columns
Delete resource group
Refresh
Move
Assign tags
Delete

Overview
Activity log
Access control (IAM)
Tags
Events

Settings
Quickstart
Resource costs
Deployments
Policies
Properties
Locks
Automation script

Monitoring
Insights (preview)
Alerts
Metrics
Diagnostic settings

Subscription (change)
Jim's Personal Azure Account

Deployments
3 Succeeded

Subscription ID
2ed0e6a0-aaf7-4f5f-aff6-bf25ffdef188

Tags (change)
[Click here to add tags](#)

Filter by name...
All types
All locations
No grouping

11 items
☐ Show hidden types

NAME	TYPE	LOCATION
EComVM	Virtual machine	South Central US
EComVM_OsDisk_1_1daee8c7f45b4205b14c5c93a5546...	Disk	South Central US
econvm34	Network interface	South Central US
EComVM-ip	Public IP address	South Central US
EComVM-nsg	Network security group	South Central US
jwc900	SQL server	Central US
900StoreDB (jwc900/900StoreDB)	SQL database	Central US
ServicePlan9dbd216e-8674	App Service plan	Central US
webstore900	App Service	Central US
webstorefrontdiag	Storage account	South Central US

Figure 2-4 An Azure resource group

It's convenient to see all of the resources associated with a particular app, but you aren't locked into that paradigm. This is a useful example, because it's a common use of resource groups, but you can organize your resource groups any way you choose. Notice in [Figure 2-4](#) that you see resources in several different Azure regions (Regions are in the Location column). If you have access to multiple Azure subscriptions, "you can also" have resources from multiple subscriptions in a single resource group.

If you look at the left side of [Figure 2-4](#), you'll see a menu of operations that you can perform on your resource group. We won't go into all of these because it's out of scope for the AZ-900 exam, but there are a few that are helpful in understanding the benefit of resource groups.

If you click on **Resource Costs**, you can see the cost of all of the resources in this resource group. Having that information at your fingertips is especially helpful in situations where you want to make sure certain departments in your company are charged correctly for the used resources. In fact, some companies will create resource groups for each department rather than creating them scoped to applications. Having a Sales and Marketing resource group or an IT Support resource group, for instance, can help you immensely in reporting and controlling costs.



Exam Tip

An Azure resource can only exist in one resource group. In other words, you can't have a virtual machine in a resource group called WebStorefront and also in a resource group called SalesMarketing, because it must be in one group or the other. You can move Azure resources from one resource group to another.

You can also click on Automation Script and Azure will generate an ARM template that you can use to deploy all of these Azure resources. This is useful in a situation where you want to deploy these resources at a later time, or when you want to deploy them to another Azure subscription.

If you click on Tags, you can apply one or more tags that you choose to your resource group. A tag consists of a name and a value. For example, suppose a company is participating in two trade events: one in Texas and one in New York. You have also created a lot of Azure resources to support those events. You want to view all of the Azure resources for a specific event, but they're spread out across multiple resource groups. By adding a

tag to each resource group that identifies the event it's associated with, you can solve this problem.

In [Figure 2-5](#), you can see the tags associated with a WebStorefront resource group. This resource group has been assigned a tag named EventName, and the value of that tag is ContosoTexas. By clicking on the cube icon to the right of the tag, you can view all resources that have that tag.




Figure 2-5 Tagging a resource group

To view all of your tags, choose **All Services** from the main menu in the portal, and then click on Tags as shown in [Figure 2-6](#).




Figure 2-6 Viewing all tags

You can apply a tag to most Azure resources, not just resource groups. It's also important to understand that by adding a tag to a resource group, you are not adding that tag to the resources within the resource group. If you have a web app in the WebStorefront resource group, that web app does not inherit the tag that is applied to the resource group. Because of that, tags add an additional layer of flexibility and powerful when viewing your Azure resources.



Exam Tip

Tags can also help you organize your Azure billing expenses. When you download your Azure invoice, resource tags will appear in one of the columns, and because Azure invoices can be downloaded as comma-separated values, you can use tools like Microsoft Excel to filter based on tags.

When you delete a resource group, all of the resources in that resource group are automatically deleted. This makes it easy to delete multiple Azure Resources in one easy step. Suppose you are testing a scenario and you need to create a couple of virtual machines, a database, a Web App, and more. By placing

all these resources in one resource group, you can easily delete that resource group after your testing and Azure will automatically delete all of the resources in it for you. This is a great way to avoid unexpected costs associated with resources you are no longer using.

Throughout this skill section, you've learned about some of the benefits of using Azure. Because Azure regions are spread out across the world in different geographies, you can be assured that your data and apps are hosted where you need them to be and that any regulations or data requirements are complied with. You learned that there are multiple datacenters in each region, and by deploying your applications in availability zones, you can avoid impact from a failure in a particular datacenter.

You also learned about Azure Resource Manager (ARM) and how it can help you achieve consistent deployments to Azure and to manage your Azure resources easily. Finally, you learned about using resource groups to organize your Azure resources and how to categorize billing using tags. In the next skill section, you'll learn details about some of the specific products that are core to Azure.

SKILL 2.2: DESCRIBE SOME OF THE CORE PRODUCTS AVAILABLE IN AZURE

As we went over the core Azure architectural

components, you noticed some references to some of the products available in Azure. There were also some details about the Azure portal, but we'll cover that in detail in Skill 2.4. In this skill section, we'll talk about some of the core Azure products in four different categories:

- **Azure compute** This refers to the resources that provide computing power to run your applications. Azure offers both IaaS and PaaS compute products.
- **Azure networking** These products provide connectivity between Azure resources, and to and from the Internet or your on-premises resources.
- **Azure storage** These products give you secure and reliable cloud storage for your data.
- **Azure database** These products provide highly-scalable solutions for hosting databases of many varieties.

Note Using Azure

In this skill section, you'll create a couple of Azure resources, so you'll need an Azure subscription. If you don't have one, you can get a free trial by going to: <https://azure.microsoft.com/free/>.

This section covers:

- Azure compute products
- Azure networking products
- Azure storage products
- Azure database products
- The Azure Marketplace and its usage scenarios

Azure compute products

Azure compute products allow you to easily and dynamically allocate resources that are needed for any computing task. You can create compute resources quickly when you need them, and when your needs grow, you can scale those resources to handle additional requirements. By using Azure compute resources for your computing needs, you can more easily control costs because you don't pay for resources unless you need them. You can also allocate infrastructure much more quickly than you can in the on-premises world, and you can benefit from the economies of scale that Azure affords and use extremely powerful computers that you might not otherwise be able to afford.

Some examples of compute products in Azure are Azure Virtual Machines, Azure App Service, container offerings in Azure, and serverless computing. (Serverless computing is covered in Skill 2.3).

Azure virtual machines

A virtual machine (VM) is a software-based computer that runs on a physical computer. The physical computer is considered the *host*, and it provides the underlying physical components such as disk space, memory, CPU power, and so on. The host computer

runs software called a hypervisor that can create and manage one or more VMs, and those VMs are commonly referred to as *guests*.

The operating system on a guest doesn't have to be the same operating system that the host is running. If your host is running Windows 10, you can run a guest that uses Windows Server 2016, Linux, or many other operating systems. This flexibility makes VMs extremely popular. However, because the VMs running on a host use the physical systems on that host, if you have a need for a powerful VM, you'll need a powerful physical computer to host it.

By using Azure Virtual Machines, you can take advantage of powerful host computers that Microsoft makes available when you need computing power, and when you no longer need that power, you no longer have to pay for it.

To create an Azure Virtual Machine, log into the Azure portal using your Azure account and then follow these steps as shown in [Figures 2-7 through 2-9](#).

1. Click **Create A Resource**.
2. Click **Compute**.
3. Click **Ubuntu Server**.




Figure 2-7 Creating a virtual machine

4. Next to Resource Group, click **Create New** to create a new resource group.
5. Enter **TestRG** as the resource group name and click **OK**.
6. Enter **TestVM** as your VM name.
7. Scroll down and select **Password** for the authentication type.
8. Enter a username for your administrator account.
9. Enter a password you'd like to use for your administrator account.
10. Confirm the password.

11. Leave all the other settings as they are and click **Review + Create** to validate your settings.

More Info Virtual Machine Settings and Options

There are many more options you can choose for your VM. We could have clicked **Next : Disks**, as shown in [Figure 2-9](#), to move to additional pages that contain many more options. You can also click one of the tabs (**Disks**, **Networking**, **Management**, and so on as shown in [Figure 2-8](#)) to change specific settings. However, if you choose, you can use the default settings like we've done by clicking **Review + Create** as soon as you've entered the information Azure requires for a VM.

Dashboard > New > Create a virtual machine

Create a virtual machine

Basics Disks Networking Management Guest config Tags Review + create

Create a virtual machine that runs Linux or Windows. Select an image from Azure marketplace or use your own customized image. Complete the Basics tab then Review + create to provision a virtual machine with default parameters or review each tab for full customization.
Looking for classic VMs? [Create VM from Azure Marketplace](#)

PROJECT DETAILS

Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

* Subscription ⓘ

Jim's Personal Azure Account

* Resource group ⓘ

5 (New) TestRG

Create new 4

INSTANCE DETAILS

* Virtual machine name ⓘ

6 TestVM

* Region ⓘ

South Central US

Availability options ⓘ

No infrastructure redundancy required

* Image ⓘ

Ubuntu Server 18.04 LTS

[Browse all images and disks](#)

Figure 2-8 Virtual machine settings

Dashboard > New > Create a virtual machine

Create a virtual machine

* Size ⓘ **7** Standard D2s v3
2 vcpus, 8 GB memory
[Change size](#)

ADMINISTRATOR ACCOUNT

Authentication type ⓘ **7** ☒ Password ☐ SSH public key

* Username ⓘ **8** jmc ✓

* Password ⓘ **9** ✓

* Confirm password ⓘ **10** ✓

Login with Azure Active Directory (Preview) ⓘ ☐ On ☒ Off

INBOUND PORT RULES

Select which virtual machine network ports are accessible from the public internet. You can specify more limited or granular network access on the Networking tab.

* Public inbound ports ⓘ ☒ None ☐ Allow selected ports

Select inbound ports

11 [Review + create](#) [Previous](#) [Next: Disks >](#)

All traffic from the internet will be blocked by default. You will be able to change inbound port rules in the VM > Networking page.

Figure 2-9 Virtual machine settings

After you click **Review + Create**, Azure will validate your settings to make sure you haven't left anything

out. Once your validation has passed, you will see a Create button. Click the **Create** button to start the deployment of your new VM.

More Info How Azure Deploys your VM

When you click Create to create your VM, the Azure portal is actually using an ARM template to deploy your VM. That ARM template contains parameters that are replaced with the information you entered for your VM. Every VM that is created in Azure is created using an ARM template. This ensures that the deployments are consistent.

As your VM is being deployed, you'll see the status displayed in the Azure portal as shown in Figure 10-10. You can see the Azure resources that are created to support your VM. You can see the resource name, the resource type (which starts with the resource provider), and the status of each resource.

... Your deployment is underway

Check the status of your deployment, manage resources, or troubleshoot deployment issues. Pin this page to your dashboard to easily find it next time.



Deployment name: CreateVm-Canonical.UbuntuServer-18.04-LTS-20190203095907

Subscription: [Jim's Personal Azure Account](#)

Resource group: [TestRG](#)

DEPLOYMENT DETAILS [\(Download\)](#)

Start time: 2/3/2019, 10:17:36 AM
Duration: 2 minutes 1 second
Correlation ID: 11fe3143-98dd-490e-9498-b9cfa760e55e

RESOURCE	TYPE	STATUS	OPERATION DETA...
 TestVM-nsg	Microsoft.Network...	OK	Operation details
 TestRG-vnet	Microsoft.Network...	Created	Operation details
 TestVM-ip	Microsoft.Network...	OK	Operation details
 testrgdiag898	Microsoft.Storage...	Accepted	Operation details

Figure 2-10 Virtual machine settings

Once all the resources required for your VM are created, your VM will be considered fully deployed. You'll then be able to click the **Go To Resource** button to see the management interface for your VM in the Azure portal as shown in [Figure 2-11](#).




Figure 2-11 Viewing a virtual machine

Our new VM is a guest on a physical computer with

an Azure datacenter. In that datacenter is a physical rack of computer servers, and our VM is hosted on one of those servers. The host computer is managed by Microsoft, but the VM is managed by you, because this is an IaaS offering in Azure.

Note VMs and Billing

You are charged for Azure VMs as long as they are running. To stop billing for this VM, click the Stop button at the top of the screen shown in [Figure 2-11](#). Azure will save the current state of the VM and billing will stop. You won't be able to use the VM while it's in a stopped state, but you will also avoid the billing of that VM. Keep in mind that unless you have configured a static IP address for your VM, your IP address will likely change the next time you start it.

You can also stop a VM from within the guest operating system on the VM, but when you do that, you will still be charged for the resources the VM uses because it's still allocated to you. That means you'll still incur charges for managed disks and other resources.

As of right now, this VM is susceptible to downtime due to three types of events: *planned maintenance*, *unplanned maintenance*, and *unexpected downtime*.

Planned maintenance refers to planned updates that Microsoft makes to the host computer. This includes things like operating system updates, driver updates, and so on. In many cases, updates won't impact your VM, but if Microsoft installs an update that requires a reboot of the host computer, your VM will be down during that reboot.

Azure has underlying systems that constantly

monitor the health of computer components. If one of these underlying systems detects that a component within the host computer might fail soon, Azure will flag the computer for unplanned maintenance. In an unplanned maintenance event, Azure will attempt to move your VM to a healthy host computer. When it does this, it preserves the state of the VM, including what's in memory and any files that are open. It only takes Azure a short time to move the VM, during which time it's in a paused state. In a case where the move operation fails, the VM will experience unexpected downtime.

In order to ensure reliability when a failure occurs in a rack within the Azure datacenter, you can (and you should) take advantage of a feature called *availability sets*. Availability sets protect you from maintenance events and downtime caused by hardware failures. To do that, Azure creates some underlying entities in an availability set called *update domains* and *fault domains*. (In order to protect yourself in the event of maintenance events or downtime, you must deploy at least two VMs into your availability set transpose).

Fault domains are a logical representation of the physical rack in which a host computer is installed. By default, Azure assigns two fault domains to an availability set. If a problem occurs in one fault domain (one computer rack), the VMs in that fault domain will

be impacted, but VMs in the second fault domain will not be. This protects you from unplanned maintenance events and unexpected downtime.

Update domains are designed to protect you from a situation where the host computer is being rebooted. When you create an availability set, Azure creates five update domains by default. These update domains are spread across the fault domains in the availability set. If a reboot is required on computers in the availability set (whether host computers or VMs within the availability set), Azure will only reboot computers in one update domain at a time and it will wait 30 minutes for computers to recover from the reboot before it moves on to the next update domain. Update domains protect you from planned maintenance events.

Figure 2-12 shows the diagram that Microsoft uses to represent an availability set. In this diagram, the fault domains FDo, FD1, and FD2 encompass three physical racks of computers. UDo, UD1, and UD2 are update domains within the fault domains. You will see this same representation of an availability set within other Azure training as well, but it's a bit misleading because update domains are not tied to a particular fault domain.




Figure 2-12 Microsoft documentation representation of an availability set

Figure 2-13 shows a better representation of an availability set, with five VMs in the availability set. There are two fault domains and three update domains. When VMs were created in this availability set, they were assigned as follows:

- The first VM is assigned Fault Domain 0 and Update Domain 0.
- The second VM is assigned Fault Domain 1 and Update Domain 1.
- The third VM is assigned Fault Domain 0 and Update Domain 2.
- The fourth VM is assigned Fault Domain 1 and Update Domain 0.

- The fifth VM is assigned Fault Domain 0 and Update Domain 1.




Figure 2-13 A better representation of an availability set

You can verify the placement of fault domains and update domains by creating five VMs in an availability set with two fault domains and three update domains. If you then look at the availability set created in the Azure

portal as shown in [Figure 2-14](#), you can see the same configuration depicted in [Figure 2-13](#).

The screenshot shows the Azure portal interface for an availability set named 'WebAvailabilitySet'. The left sidebar contains navigation links: Overview, Activity log, Access control (IAM), Tags, Settings, Virtual machines, Properties, Locks, Automation script, Support + troubleshooting, and New support request. The main content area displays the following details:

- Resource group (change): ATest
- Location: South Central US
- Subscription (change): Jim's Personal Azure Account
- Subscription ID: 2ed0e6a0-aaf7-4f5f-aff6-bf25ffdef188
- Managed: Yes

Summary statistics (highlighted with a red box):

- Fault domains: 2
- Update domains: 3
- Virtual machines: 5

Below these details is a table of virtual machines (highlighted with a red box):

NAME	STATUS	FAULT DOMAIN	UPDATE DOMAIN
VM1	Running	0	0
VM2	Running	1	1
VM3	Running	0	2
VM4	Running	1	0
VM5	Running	0	1

Figure 2-14 An availability set in the Azure portal showing fault domains and update domains

Notice in [Figure 2-14](#) that the availability set is named WebAvailabilitySet. In this availability set, we run five VMs that are all running a web server and host the

website for an application. Suppose you need a database for this application, and you want to host that database on VMs as well. In that situation, you would want to separate the database VMs into their own availability set. As a best-practice, you should always separate your workloads into separate availability sets.

Availability sets certainly provide a benefit in protecting from downtime in certain situations, but they also have some disadvantages. First of all, every machine in an availability set has to be explicitly created. While you can use an ARM template to deploy multiple virtual machines in one deployment, you still have to configure those machines with the software and configuration necessary to support your application.

An availability set also requires that you configure something in front of your VMs that will handle the distribution of traffic to those VMs. For example, if your availability set is servicing a website hosted on the VMs, you'll need to configure a load balancer that will handle the job of routing users of your website to the VMs that are running it.

Another disadvantage to availability sets relates to cost. In a situation where your VM needs changed often based on things like load on the application, you might find yourself paying for many more VMs than you need.

Azure offers another feature for VMs called *scale sets*

that solves these problems nicely. When you create a scale set, you tell Azure what operating system you want to run and then you tell Azure how many VMs you want in your scale set. You have many other options such as creating a load balancer or gateway and so forth. Azure will create as many VMs as you specified (up to 1,000) in one easy step.

More Info Using a Custom Image

The default set of templates for VMs are basic and include only the operating system. However, you can create a VM, install all of the necessary components you need (including your own applications), and then create an image that can be used when creating scale sets.

For more information on using custom images, see:

<https://docs.microsoft.com/azure/virtual-machine-scale-sets/virtual-machine-scale-sets-deploy-app#build-a-custom-vm-image>.

Scale sets are deployed in availability sets automatically, so you automatically benefit from multiple fault domains and update domains. Unlike VMs in an availability set, however, VMs in a scale set are also compatible with availability zones, so you are protected from problems in an Azure datacenter.

As you might imagine, you can also scale a scale set in a situation where you need more or fewer VMs. You might start with only one VM in a scale set, but as load on that VM increases, you might want to automatically add additional VMs. Scale sets provide that functionality by using Azure's auto-scale feature. You define scaling

rules that use metrics like CPU, disk usage, network usage, and so forth. You can configure when Azure should add additional instances and when it should scale back and deallocate instances. This is a great way to ensure availability while reducing costs by taking advantage of the elasticity that auto-scale provides.

More Info Scaling and Availability Sets

Before the introduction of scale sets, you had the ability to configure auto-scale rules for an availability set. You'll probably still see third-party documentation and training that talks about scaling availability sets, but that functionality has been replaced with scale sets.

Microsoft guarantees an SLA of 99.95% when you use a multi-VM deployment scenario, and for most production scenarios, a multi-VM deployment is preferred. However, if you use a single-instance VM, and you use premium storage, Microsoft guarantees a 99.9% SLA. Premium storage uses solid-state drives (SSDs) that are located on the same physical server that is hosting the VM for enhanced performance and uptime.

Containers in Azure

It's becoming pretty commonplace for companies to move applications between “environments,” and this type of thing is even more prevalent when it comes to the cloud. In fact, one of the most complex aspects of

moving to the cloud is dealing with the complexities of moving to a new environment. To help with this problem and to make it easier to shift applications into new environments, the concept of *containers* was invented.

A container is created using a zipped version of an application called an *image*, and it includes everything the application needs to run. That might include a database engine, a web server, and so on. The image can be deployed to any environment that supports the use of containers. Once there, the image is used to start a container the application runs in.

In order to run an application in a container, a computer needs to have a container runtime installed on it. The most popular container runtime is Docker, a runtime developed and maintained by a company called Docker Inc. Docker not only knows how to run applications in containers, but it also enforces certain conditions to ensure a secure environment.

More Info Docker Images

You aren't limited to your own images. In fact, Docker runs a repository of images that you are free to use in your own applications. You can find it at: <https://hub.docker.com>.

Each container operates within an isolated environment. It has its own network, its own storage,

and so on. Other containers running on the same machine cannot access the data and systems used by another container. This makes containerized applications an ideal solution when security is a concern.

Azure offers numerous technologies for hosting containers. Azure Container Instances (ACI) is a PaaS service that makes it easy to start a container with minimal configuration. You simply tell ACI where to find the image (using either a Docker tag or a URL to the image) and some basic configuration for the VM you want the container to run on.

Azure creates server resources as needed to run your container, but you're not paying for an underlying VM. Instead, you pay for the memory and CPU that your container uses. That translates into extremely low costs in most cases. For example, if your ACI app is running on a machine with 1 CPU and 1 GB of memory and you use the app for 5 minutes a day, at the end of the month, your cost would be less than 5 cents!

Note Containers Use their Own Operating System

The operating system for a container is actually part of the image. The VM that you are configuring when you create an ACI app is the VM that runs the container runtime. Even so, it's important that you choose an operating system that's compatible with your container. A Docker image that was built for Linux will not run on a Windows host and vice versa.

ACI is designed to work with simple applications. You can define a container group and run multiple containers within an ACI instance, but if you have an application that is used heavily by many people and that might need to take advantage of scaling, ACI isn't a good choice for you. Instead, Azure's Kubernetes Service (AKS) would be a better choice.

Kubernetes is a container orchestration service. This means that it's responsible for monitoring containers and ensuring that they're always running. It can also scale to add additional containers when the needs require it to, and it can then scale back when the needs are reduced.

Kubernetes creates containers in a *pod*. A pod is a group of related containers, and containers within a pod are able to share resources. This is one of the advantages to using Kubernetes, because it releases you from the resource-sharing restriction typically imposed in a multi-container environment. However, a container in one pod is not able to share resources with a container in another pod.

The computer that Kubernetes pods are running on is called a *node* or a *worker*. This computer must have a container runtime such as Docker running on it. In addition to pods, the node also runs several services that are required for Kubernetes to manage the pods,

and so on. There will typically be multiple nodes within a Kubernetes instance, and they are all controlled by a master node called the Kubernetes *master*. The entire environment of the master and all of its nodes is called a Kubernetes *cluster*.

A Kubernetes master contains all of the configuration and services necessary to manage the orchestration of pods and other Kubernetes entities. Configuring a master can be complex, and it is by far the most laborious task of using Kubernetes. For that reason, services such as Azure Kubernetes Service (AKS) are becoming more popular.

AKS offloads the burden of dealing with the Kubernetes master to Microsoft. When you create a Kubernetes cluster in AKS, Azure creates the master and the nodes for you. All you have to do is deploy your containers, and you're up and running with a managed Kubernetes cluster.

AKS simplifies the creation of a Kubernetes cluster, but it also makes it extremely easy to manage a cluster (see [Figure 2-15](#)). Operations, such as upgrading a cluster or scaling a cluster, are simple using the Azure portal menu options. You can also get detailed information on your cluster, including each node that's running in the cluster.




Figure 2-15 An AKS cluster in the Azure portal

While AKS makes adopting and managing Kubernetes easier, it doesn't completely obfuscate Kubernetes. In order to deploy your applications, you still need to understand how to use Kubernetes, and in some cases you'll need to use the Kubernetes command line. Azure,

however, makes it far easier than doing all of the legwork and maintenance yourself. Even better, AKS in Azure is free. You only pay for the Azure computer for resources that you use within your cluster.

For a true PaaS experience in container hosting, Microsoft offers Web App for Containers, a feature of Azure App Service. When you create a Web App for Containers app, you specify the OS you want (either Windows or Linux) and you specify the location of the Docker image (see [Figure 2-16](#)). The image can be in Docker Hub, a private registry, or in Azure Container Services.

Dashboard > New > Web App for Containers >

Web App for Containers

Create

App name: contosojc ✓

Subscription: Jim's Personal Azure Account

Resource Group: WebTest ✓

OS: Linux

App Service plan/Location: ServicePlane0845c1e-9f6e/Centr...

Configure container >

Single Container | Docker Compose (Preview)

Image source: Quickstart, Azure Container Registry, Docker Hub, Private Registry

Repository Access: Public, Private

Image and optional tag (eg 'image:tag'): jamesche/dockertest

Startup File:

Figure 2-16 Creating a Web App for Containers app

Containers that are running in Web App for Containers enjoy the benefits of all of the PaaS features of Azure App Service. Microsoft manages the infrastructure that's involved, so you only have to worry about the application contained in the image.

Unlike ACI, you pay for Azure App Service whether you're using the application or not, because your

application is running on a dedicated VM in App Service. That VM is associated with an App Service plan, and each App Service plan is associated with a specific pricing tier. You can change the pricing tier of your App Service Plan at any time. For example, if you decide that your application needs more memory than you first thought, you can scale up to a higher tier and get more memory. App Service takes care of moving your app to the new VM.

App Service also makes it easy to scale out by using Azure auto-scale. Just like scaling a VM scale set, you can specify metrics that are used to determine when to scale your app. Keep in mind, however, that you pay for each VM that you use, so if you scale out to a large number of VMs, you're going to see an equally large bill at the end of the month.

Another benefit of using Web App for Containers is that, because it's a true PaaS service, it offers many turnkey features that you can use in your application without having to deal with complicated development or configuration issues. For example, if you want to enforce authentication in your application, and you want users to be able to use their Microsoft Account, Facebook, Twitter, or Google login credentials, you can configure that easily with App Service Authentication as shown in [Figure 2-17](#).




Figure 2-17 Configuring authentication for Web App for Containers

Azure networking products

Applications in Azure are almost always composed of

multiple Azure services working together. Even though these multiple services rely on each other for the application to function, they should not be tightly integrated. Instead, applications should be designed using a *loosely-coupled architecture*.

In a loosely-coupled architecture, each component of an application can be replaced or updated without breaking the application. In order to design applications in this way, you have to separate out the various components, and they need to operate in their own tier of the application. It's this separation of components that allows you to be more flexible in the implementation details of your application, and it's a critical component to an application designed for the cloud. Applications designed in this way are referred to as *N-tier* applications.



Exam Tip

The AZ-900 exam isn't an exam for developers, so we won't go into any level of detail about application design. It is important for you to understand the concept of multi-tier applications, however, so that you understand why Azure's networking features work the

way that they do.

Suppose you have an application that records sales data for your company. Users enter their sales records, and the application performs some analysis on them, and then stores the information in a database. The application uses three tiers: a web tier, a middle tier, and a data tier.

The web tier is a website running in Azure App Service. It's there only to give the user a way to interact with the application. It doesn't handle any logic. It simply takes what the user inputs and passes it on to the middle tier where the work actually happens.

The middle tier (or application tier) is where all of the application logic exists. This is where the application analyzes the sales data for trends, and applies business rules to it as it's running in an Azure Virtual Machine. The data tier is where you store sales data, but the middle tier can also retrieve sales data from it when you need to display reports. The data tier consists of an Azure SQL Database. [Figure 2-18](#) shows a diagram of the application.




Figure 2-18 An example of an n-tier application.

Here are a few requirements for this application.

- Only the web tier can talk to and from the Internet.
- The web tier can talk to the middle tier, but it cannot talk to the data tier.
- The middle tier can talk to and from the web tier and the data tier.
- The data tier can talk to and from the middle tier, but it cannot talk to the web tier.

These requirements are typical for an N-tier design, and they help to keep data secure and prevent security issues with the application. Since each of these tiers is running in a separate Azure service, they can't talk to

each other by default. In order to communicate between the tiers of your application, you need a computer network, and that's where Azure's networking products come into play.

Azure virtual network

An Azure virtual network (often called a VNET) allows Azure services to communicate with each other and with the Internet. You can even use a VNET to communicate between your on-premises resources and your Azure resources. When you created the virtual machine earlier in this chapter, Azure created a VNET for you. Without that VNET, you wouldn't be able to remote into the VM, or use the VM for any of your applications. You can also create your own VNET and configure it any way you choose.

An Azure VNET is just like any other computer network. It's comprised of a network interface card (a NIC), IP addresses, and so on. You can break up your VNET into multiple subnets and set up a portion of your network's IP address space for those subnets. You can then configure rules that control the connectivity between those subnets.

Figure 2-19 illustrates an Azure VNET that we might use for the sales application. The VNET uses IP addresses in the 10.0.0.0 address range and each subnet has its own range of addresses. IP address ranges in

VNETs are specified using classless inter-domain routing (CIDR) notation, and a discussion of that is far outside of the scope of this exam. However, with the configuration shown in [Figure 2-19](#), we have 65,536 IP addresses available in our VNET, and each subnet has 256 IP addresses allocated to it. (The first four IP addresses and the last IP address in the range are reserved for Azure's use, so you really only have 251 addresses to use in each subnet.) This is a typical design because you still have a large number of addresses available in your network for later expansion into additional subnets.




Figure 2-19 Your application in an Azure Virtual Network

In most cases, you create VNETs before you create

the resources that use them. If you go back and look at [Figure 2-10](#), you'll see that Azure has automatically created a VNET for the VM. It does that because you can't use a VM unless there's a network associated with it. While you can connect a VNET to an existing VM, you can't move a VM into another network. For that reason, you create your VNET before you create your VM.

Our web tier, on the other hand, is running in Azure App Service, a PaaS offering. This is running on a VM that Microsoft manages, so Microsoft has created and manages the VM and its network. In order to use that tier with the VNET, App Service offers a feature called VNET Integration that allows you to integrate a web app in App Service with an existing VNET.

The IP addresses within the VNET at this point are all private IP addresses. They allow resources within the VNET to talk to each other, but you can't use a private IP address on the Internet. You need a public IP address in order to give the Internet access to your web tier.

More Info Outbound Internet Connectivity

A public IP address doesn't have to be assigned to a resource in order for that resource to connect outbound to the Internet. Azure maintains a pool of public IP addresses that can be dynamically assigned to a resource if it needs to connect outbound. That IP address is not exclusively assigned to the resource, so it cannot be used for inbound communication from the Internet to the Azure resource.

Since the web tier is running on Azure App Service (a PaaS service), Microsoft manages the public-facing network for us. You get Internet access on that tier without having to do anything. If you want to run the web tier on an IaaS VM instead, configure the public IP address for the web tier. In those situations, Azure allows you to create a Public IP Address resource and assign it to a virtual network.

More Info Network Security Groups

Azure offers a feature called Network Security Groups that allow you to enforce rules about what kind of traffic is allowed on the VNET. We'll cover Network Security Groups in [Chapter 3, "Understand security, privacy, compliance, and trust."](#)

Azure load balancer

It is easy to scale out the web tier in our sales application when needed. App Service takes care of ensuring that load is distributed across all of the VMs we're using. App Service uses a load balancer to do this, and one of the advantages of choosing a PaaS offering for the web tier is that you don't have to worry about Managing it. If you use an IaaS VM running a web server for the web tier, you may want to have more than one VM in order to handle additional load if needed. [Figure 2-20](#) represents what the web tier might look like using an IaaS model.




Figure 2-20 The web tier using an IaaS model with Azure VMs

This kind of configuration is typical in order to maintain high-availability in your application, but it does add an additional layer of complexity. Since each of these VMs has its own public IP address, a user is going to use only one VM. Ideally you have a system in place that ensures if one of these VMs experiences a problem, any traffic is sent to the other VMs. In addition to that, when there's high load, you want to spread the load across all three of these VMs. The solution to this problem is to use Azure LoadBalancer.

Azure Load Balancer is inside of the VNET, but it sits between the user and the subnet. When a user connects to the web tier, she connects to the load balancer's IP

address, not the IP address of one of my VMs. The load balancer routes requests into the web tier to the VMs, and it can use rules to ensure that traffic is equally distributed between them. If one of the VMs goes down and doesn't respond, the load balancer can send that traffic to another VM without the user even realizing there's a problem.

Figure 2-21 shows the web tier from Figure 2-20 with Azure Load Balancer added to the mix. Notice that the public IP is now on the load balancer, and the VMs are using the private IP addresses from the subnet.




Figure 2-21 The web tier with Azure Load Balancer

Azure Load Balancer isn't just for distributing traffic from the Internet. In order for our application to maintain high-availability, we should ensure the same scalability of other tiers, and Azure Load Balancer can sit within other tiers as well to ensure that load is distributed, and to ensure the application maintains the high-availability necessary for the business.

Azure Application Gateway

Using Azure Load Balancer for the web tier is a perfectly suitable solution, but since the web tier only uses HTTP traffic for the web site, we can gain additional features specific to HTTP traffic by using Azure Application Gateway.

Azure Application Gateway is a load balancer that's specifically designed to deal with HTTP traffic. Because Application Gateway understands HTTP traffic, it can make decisions based on that HTTP traffic. For example, Application Gateway can:

- Route traffic to a specific VM or pool of VMs based on the URL.
- Use a cookie to ensure that a user is always routed to the same VM in a situation where that VM contains state information on that user that must be maintained.
- Display a customized error page, complete with your company branding, when a page isn't found or when an error occurs.

- Handle the SSL traffic for your site so that your application tiers don't have the overhead of dealing with decrypting traffic.

You can also add Web Application Firewall (WAF) to Application Gateway. WAF is designed to stop known vulnerabilities from making it into your VNET, allowing you to operate in a more secure environment. If a request attempts to enter your network and it's determined to be a threat, it's rejected at the gateway and never makes it to your application.

VPN Gateway

In some cases, you may need your application hosted in Azure to talk to an on-premises resource. We talked about these types of scenarios in [Chapter 1](#) when we covered hybrid cloud scenarios. To implement such a system, you can use Azure's VPN Gateway.

VPN Gateway connects your on-premises resources to your Azure VNET using a virtual private network, or VPN. Traffic that flows over this VPN is encrypted. There are multiple configurations for VPN Gateway connections as shown in [Table 2-1](#).

Table 2-1 Type of VPN Gateway Connections

Connection Type	Description
Site-to-Site	Connects your VNET to a single on-premises location.

VPN (S2S)	Requires a VPN with public-facing IP address on-premises. A multi-site variant allows you to connect to multiple on-premises locations.
Point-to-Site VPN (P2S)	Connects one specific on-premises client PC to your VNET. Multiple clients can connect, but each one connects over its own VPN client.
VNET-to-VNET	Connects two Azure VNETs to each other. Useful in situations where you have two VNETs in different Azure regions and you want to securely connect them.

More Info VNET PEERING

As an alternative to using VNET-to-VNET connections, you can use VNET peering to establish communication between two Azure VNETs in the same region, and you can use global VNET peering to connect VNETs in different Azure regions. Peering is typically used in a scenario where you don't require a gateway for connectivity to on-premises resources.

For more information on VNET peering, see <https://docs.microsoft.com/azure/virtual-network/virtual-network-peering-overview>.

Azure Content Delivery Network

Azure Content Delivery Network (CDN) is an effective way of delivering large files or streaming content over the Internet. It makes the downloading of large files much faster by caching the files in multiple geographical locations so that users can get the files from a server as close to them as possible. CDNs are typically used with images, videos, and other similarly large files.

A CDN works by storing a cached version of files on a *point-of-presence* (POP) server that is located on the outside edge of a network. These servers (called *edge* servers) are able to serve content without having to go through the entire network, a process which adds time to a request.

Microsoft has CDN edge servers located across the globe, so when a user requests large files from any geographical location, it can serve a cached copy that is as close to the user's location as possible. The content on an edge server has a *time-to-live* (TTL) property associated with it that tells the edge server how long it should keep the cached copy. If a TTL isn't specified, the default TTL time is seven days. Once a cached copy is removed, the next time that resource is requested, the edge server will make a request to the server where the original copy of the resource is located. It will then cache it again for future users until the TTL expires.

Azure Traffic Manager

Azure Traffic Manager is a domain name system (DNS)-based system that's designed to enhance the speed and reliability of your application. To use Traffic Manager, you configure *endpoints* within Traffic Manager. An endpoint is simply a resource that you want users to connect to. Traffic Manager supports public IP addresses connected to Azure VMs, web apps running in

App Service, and cloud services hosted in Azure. An endpoint can also be a resource located on-premises or even at another hosting provider.

Once you've configured your endpoints, you specify routing rules that you want Traffic Manager to apply to them. There are many routing rules available in Traffic Manager.

- **Priority** All traffic is sent to a primary endpoint, but backup endpoints are available in case the primary endpoint experiences an outage.
- **Weighted** Traffic is distributed across endpoints. By default, all traffic is distributed evenly, but you can specify a weight for each endpoint and traffic will be distributed as you specify.
- **Performance** Traffic Manager determines the endpoint with the lowest network latency from the user's location and uses that endpoint.
- **Geographic** Traffic is routed based on the geographic location of the DNS server that queries Traffic Manager.
- **Multivalue** Returns all valid endpoints that use the specified Internet protocol version, either IPv4 or IPv6.
- **Subnet** Traffic is routed based on the end-user IP address range.

One important thing to remember is that Traffic Manager is DNS-based. That means that a user never directly talks to Traffic Manager. Traffic Manager is only used for the DNS lookup. Once an IP address is known for the desired endpoint, all subsequent requests bypass Traffic Manager entirely. Also, because Traffic Manager

is DNS-based, the actual traffic between the user and the resource is never sent through Traffic Manager.

Azure storage products

Azure offers many options for storing data. Whether you need to store data temporarily on a disk mounted to a VM, or you need to be able to store data long-term, Azure has an option to fit your needs.

Azure Blob Storage

Azure Blob Storage is designed for storing unstructured data, which has no defined structure. That includes text files, images, videos, documents, and much more. An entity stored in Blob Storage is referred to as a *blob*. There are three types of blobs in Azure Storage.

- **Block blobs** Used to store files used by an application.
- **Append blobs** They are like block blobs, but append blobs are specialized for append operations. For that reason, they are often used to store constantly updated data like diagnostic logs.
- **Page blobs** They are used to store virtual hard drive (.vhd) files that are used in Azure virtual machines. We'll cover these in Azure Disk Storage later in this chapter.

Blobs are stored in storage containers. A container is used as a means of organizing blobs, so you might have a container for video files, another container for image files, and so on. The choice, however, is entirely up to you.

Microsoft offers numerous storage tiers that are priced according to how often the data is accessed, how long you intend to store the data, and so on. The Hot storage tier is for data you need to access often. It has the highest cost of storage, but the cost for accessing the data is low. The Cool storage tier is for data that you intend to store for a longer period and not access quite as often. It has a lower storage cost than the Hot tier, but access costs are higher. You're also required to keep data in storage for at least 30 days.

Microsoft also offers an Archive storage tier for long-term data storage. Data stored in the Archive tier enjoys the lowest storage costs available, but the access costs are the highest. You must keep data in storage for a minimum of 180 days in the Archive tier. Because data in the Archive tier isn't designed for quick and frequent access, it can take a very long time to retrieve it. In fact, while the Hot and Cold access tiers guarantee access to the first byte of data within milliseconds, the Archive tier only guarantees access to the first byte within 15 hours.

If you're planning on moving data from on-premises into Azure Storage, there are many options available to you. You can use Azure Storage Explorer, a free tool available from Microsoft, to upload data. You can also use command line tools that Microsoft provides for uploading to Azure Storage.

If you want to move a large amount of data, Microsoft offers a service called Data Box. Data Box has an online service called Data Box Edge that makes copying data to Azure Storage as easy as copying it to a hard drive on your system. For even larger amounts of data, Microsoft offers a Data Box offline service where they will ship you hard drives. You simply copy your data to the hard drives, encrypt the drives with BitLocker, and then ship them back to Microsoft. They even offer Data Box Heavy, a service where they'll ship you a rugged device on wheels that can hold up to 1 petabyte of data!

Azure Queue Storage

A message queue is a component in an application that can store messages that an application uses to know what tasks to take. For example, you may have an application that performs image manipulation on pictures, and some of those manipulations might take much longer than others. If you have thousands of people using the application, a message queue can help to ensure a responsive and reliable application by allowing one component to put messages in the queue and your image manipulation component can then retrieve those messages, perform the manipulation, and put a message back on the queue.

Azure Queue Storage provides a cloud-based message queue that can be accessed securely from application

components located anywhere. They can be located in the cloud or on-premises. Queue Storage can asynchronously process millions of messages up to 64KB in size. The sender of the message expects the receiver to take action on it only when it's ready. You can think of this in the same way that email works. You send an email to a receiver and the receiver deals with it when they have time. You don't expect an immediate response.

More Info Authorization To Queue Storage

Access to Queue Storage is protected and authorized using either Azure Active Directory or a shared key.

To access Queue Storage, your application uses the APIs available for the language the application was written in. Microsoft provides APIs for use with .NET, Java, Node.js, C++, PHP, Python, and Ruby.

Azure Disk Storage

Disk storage in Azure refers to disks that are used in virtual machines. Azure creates a disk for you when you create a VM, which is automatically designated for temporary storage. This means that data on that disk will be lost if there's a maintenance event on the VM. If you need to store data for a longer period of time that will persist between VM deployments and maintenance

events, you can create a disk using an image stored in Azure Storage.

Azure disks are available as both traditional hard disks (HDD) and solid-state drives (SSD). Azure Standard HDD Disk are cheaper and designed for non-critical data. SSD disks are available in a Standard tier for light use and as Azure Premium Disk for heavy use.

Azure disks are available as either Managed Disks or unmanaged disks. All Azure disks are backed by page blobs in Azure Storage. When you use unmanaged disks, they use an Azure Storage account in your Azure subscription, and you have to manage that account. This is particularly troublesome because there are limitations in Azure Storage, and if you have heavy disk usage, you may end up experiencing downtime due to throttling.

When you move to Managed Disks, Microsoft handles the storage account, and all storage limitations are removed. All you have to worry about is your disk. You can leave the Storage account in Microsoft's hands.

More Info Managed Disks

Microsoft recommends Managed Disks for all new VMs. They also recommend that all VMs currently using unmanaged disks move to Managed Disks.

Perhaps an even more important reason to use

Managed Disks is that by doing so, you avoid a possible single point of failure in your VM. When you use unmanaged disks, there is a possibility that the Azure Storage accounts backing up your disks might exist within the same storage scale unit. If a failure occurs in that scale unit, you will lose all of your disks. By ensuring that each Managed Disk is in a separate scale unit, you avoid the situation of a single point of failure.

Azure Files

Azure disks are a good option for adding a disk to a virtual machine, but if you just need disk space in the cloud, it doesn't make sense to take on the burden of managing a virtual machine and its operating system. In those situations, Azure Files is the perfect solution.

Note Azure Files And Azure Storage

Azure Files shares are backed by Azure Storage, you so will need a storage account to create an Azure Files share.

Azure Files is a completely managed file share that you can mount just like any SMB file share. That means existing applications that use network attached storage (NAS) devices or SMB file shares can use Azure Files without any special tooling, and if you have multiple applications that need to access the same share, that will work with Azure Files, too.



Exam Tip

You can mount Azure Files shares on Azure VMs and on-premises on Windows, Linux, and MacOS. You can't, however, use Windows 7 or Windows Server 2008 to mount an Azure Files share on-premises because those operating systems only support SMB 2.1.

Also, because Azure Files shares use SMB, you'll need to make sure that TCP port 445 is open on your network. On Windows, you can use the Test-NetConnection PowerShell cmdlet to test connectivity over port 445. For more information, see:

<https://docs.microsoft.com/azure/storage/files/storage-how-to-use-files-windows>.

One possible problem with using Azure Files is the remote location of files. If your users or applications are using a file share mapped to Azure Files, they might experience longer than usual file transfer times because the files are in Azure. To solve that problem, Microsoft introduced Azure File Sync.

Install Azure File Sync on one or more servers in your local network and it will keep your files in Azure Files synchronized with your on-premises server. When users or applications need to access those files, they can access the local copy quickly. Any changes you make to the centralized Azure Files share are synchronized to any servers running Azure File Sync.

Azure database products

Most applications use some kind of database to store data that can be retrieved through queries and used in the application. Azure provides numerous database solutions, and if you're going to move to the cloud, it's important for you to understand the differences between them.

Azure SQL Database

Azure SQL Database is a PaaS offering for SQL Server database hosting. Microsoft manages the platform, so all you have to worry about is your database and the data in it.

SQL Server databases are *relational databases* made up of tables of data, and each table has a schema that defines what the data should look like. For example, the schema may define that your data contains an ID number, a first name, a last name, and a date. Any data that you add to the table must follow the schema, so it

must have all of the fields defined in the schema.

A database will contain many tables of data that are related to each other, and by using specialized queries, developers can return data that is a result of joining related data from multiple tables. For example, you might have a Customers table and an Orders table, each with a “CustomerID” field that identifies a customer. By querying and joining the data from both of these tables, you can provide a user with an invoice showing all of their orders. This relationship between the tables is how relational databases got their name, as shown in Figure 2-22.



Table: Customers

ID	LAST NAME	FIRST NAME	ADDRESS	CITY	STATE	ZIP
1001	Jennifer	Jones	123 Main Street	Dallas	TX	77778
1003	Robert	Earl	645 Plain Road	New York	NY	11234
1211	Raj	Pratap	7786 Chain Street	Wander	NM	84420
1220	Allison	Quandry	998 Trouble Road	Rainy	ID	45596

Table: Orders

OrderID	ITEM NUMBER	ORDER DATE	CUSTOMER
190243	10003448	02/04/2019	1211
190124	11338277	01/23/2019	1001
181233	10011287	12/23/2018	1211
190219	11121998	02/01/2019	1003

Figure 2-22 Two tables in a relational database

Note Relational Databases

SQL Server isn't the only relational database system. There are many relational database systems, including Oracle, PostgreSQL, and MySQL.

Azure offers three different deployment options for Azure SQL Database: single database, elastic pool, and managed instance.

A single database is simply a database running in a hosted SQL Server instance running in Azure. Microsoft manages the database server, so all you have to worry about is the database itself. Microsoft provides two different purchase models for single databases. [Table 2-2](#) shows these models and how they differ.

Table 2-2 Single database purchase models

Database Transaction Unit (DTU) Model	VCore Model
Good choice for users who don't need a high degree of flexibility with configuration and who want fixed pricing.	Good choice if you need a high level of visibility and control of individual resources (such as memory, storage, and CPU power) your database uses.
Pre-configured limits for transactions against the database, and pre-configured storage, CPU, and memory configurations.	Flexibility in CPU power, memory, and storage with storage charged on a usage basis.
Basic and Standard offerings, along with a Premium tier for production databases with a large number of transactions.	General Purpose and Business Critical offerings to provide lower costs when desired and high-performance and availability when required.
Ability to scale to a higher tier when needed.	Ability and flexibility to scale CPU, memory, and storage as needed.

Backup storage and long-term retention of data provided for an additional charge.	Backup storage and long-term retention of data provided for an additional charge.

An elastic pool consists of more than one database (and often many databases) all managed by the same SQL Database server. This solution is geared towards SaaS offerings where you may want to have multiple users (or maybe even each user) to be assigned their own database. You can easily move databases into and out of an elastic pool, making it ideal for SaaS.

In some cases, being able to scale a single database to add additional power is sufficient. If your application has wide variations in usage and you find it hard to predict usage (such as with a SaaS service), however, being able to add more databases to a pool is much more desirable. In an elastic pool, you are charged for the resource usage of the pool versus individual databases, and you have full control over how individual databases use those resources. This makes it possible to not only control costs, but also to ensure that each database has the resources it needs while still being able to maintain predictability in expenses. What's more, you can easily transition a single database into an elastic pool by simply moving the database into a pool.

Note Pricing Models Of Elastic Pools

The pricing model information in [Table 2-2](#) also applies to elastic pools.

Your resources aren't applied to an individual database, however, they are applied to the pool.



Exam Tip

While you can scale up and down easily with Azure SQL Database by moving to a higher tier or adding compute, memory, and storage resources, relational databases don't scale horizontally. There are some options available for scaling out a read-only copy of your database, but in general, relational databases don't offer the capability of scaling out to provide additional copies of your data in multiple regions.

A managed instance is explicitly designed for customers who want an easy migration path from on-premises or another non-Azure environment to Azure. Managed instances are fully compatible with SQL Server on-premises, and because your database server is integrated with an isolated VNET and has a private IP address, your database server can sit within your private Azure VNET. The features are designed for users who want to lift and shift an on-premises database to Azure

without a lot of configuration changes or hassle. Both the General Purpose and Business Critical service tiers are available.

Microsoft developed the Azure Database Migration Service (DMS) to make it easier for customers to easily move on-premises databases or databases hosted elsewhere in the cloud to a managed instance. The DMS works by walking you through a wizard experience to tell Azure which database(s) and table(s) you want to migrate from your source database to Azure SQL Database. It will then use the Azure VNET that comes with the managed instance to migrate the data. Once the data has been migrated, DMS sets up synchronization between the source database and Azure SQL Database. This means that as long as the source database remains online, any changes made to it will be synchronized with the managed instance in Azure SQL Database.

More Info Dms And On-Premises Databases

In order to migrate an on-premises database, you must have connectivity between Azure and your on-premises network over VPN or using a service such as ExpressRoute.

For more information on ExpressRoute, see:
<https://docs.microsoft.com/azure/expressroute/expressroute-introduction>.

Azure Cosmos DB

As you’ve seen in our discussion about SQL Server databases, relational databases lock you into a specific structure for your data. While there’s certainly a place for relational databases, as companies began to collect more and more data, they began to seek out a more flexible way to store that data. This eventually led to what are called NoSQL database systems.

In a NoSQL database system, you’re not locked into a schema for your data. If you’re storing information like that shown in the Customers table in [Figure 2-22](#), and you want to start storing customer birthdays as well, you simply add the birthday to your data and add it to the database. The database doesn’t care what kind of data there is and what fields there are.

There are four types of NoSQL database systems: key-value, column, document, and graph. [Table 2-3](#) lists each of these types and some information about them.

Table 2-3 NoSQL database systems

System	Description	Common Use
Key-value	Stores data that is tied to a unique key. Pass in the key and the database returns the data.	Since the value can be just about anything, key-value databases have many uses.
Column	NoSQL databases are called <i>keyspaces</i> , and a keyspace contains column	Storing user-profile data for a website. Also, because column databases scale well

	families. A column contains rows and columns like a relational table, but each row can have its own set of columns. You aren't locked into a schema.	and are extremely fast, they are well-suited to storing large amounts of data.
Document	Data is stored as a structured string of text called a document. This can be HTML, JSON, and so forth. This is similar to a key-value database except that the document is a structured value.	Same as key-value, but document databases have advantages. They scale well horizontally, and they allow you to query against the value and return portions of the value. A key-value database query returns the entire value associated with the key.
Graph	Stores data and the relationships between each piece of data. Data is stored in nodes, and relationships are drawn between nodes.	Many systems use graph databases because they are extremely fast. A social network might use a graph database because it would be easy to store relationships between people and also things those people like, and so forth.

There are many different NoSQL database systems, and most of them are geared toward a particular database model. Microsoft offers a hosted NoSQL database system in Azure called Cosmos DB, and Cosmos DB supports all of the NoSQL database types.

Microsoft has built some custom code around Cosmos DB so that developers can use their existing skills with other database systems with a Cosmos DB database. This makes it easy for existing applications to begin taking advantage of Cosmos DB without engineers having to write new code.

When you create a Cosmos DB database, you choose the API you want to use, and this determines the database type for your database. The database API types are:

- **Core (SQL)** Creates a document database that you can query using SQL syntax that you might be familiar with from using relational databases.
- **Azure Cosmos DB for MongoDB API** Used for migrating a MongoDB database to Cosmos DB. MongoDB databases are document databases.
- **Cassandra** Used for migrating a Cassandra database to Cosmos DB. Cassandra databases are column databases.
- **Azure Table** Used for migrating data stored in Azure Table Storage to Cosmos DB. This creates a key-value database.
- **Gremlin** Used for migrating Gremlin databases to Cosmos DB. Gremlin databases are graph databases.

The reason Microsoft calls these APIs is because they are just that. They are application programming interfaces that allow developers who are already using an existing NoSQL database technology to migrate to Cosmos DB without having to change their code.

Another huge advantage to Cosmos DB is a feature Microsoft calls turnkey global distribution. This feature takes advantage of the horizontal scalability of NoSQL databases and allows you to replicate your data globally with a few clicks. In the Azure portal, you can simply click on the region(s) where you want data replicated, as shown in [Figure 2-23](#). Once you click Save, Cosmos DB will begin to replicate data, which will be available in the selected regions. This makes it easy to ensure that users have the fastest experience possible with an application.



Figure 2-23 Easy replication across the globe with Cosmos DB

The Azure Marketplace and its usage scenarios

You’ve learned about many of the products and services available in Azure, but there are many available products outside of what we’ve discussed. Not only does Microsoft offer many additional services, but third-party vendors also provide a wide array of resources you can use in Azure. All of these resources are available in a

single repository called the Azure Marketplace.

To access the Azure Marketplace, click on **Create A Resource** in the Azure portal as shown in [Figure 2-24](#). This will display a list of categories you can choose from. It will also show a list of popular offerings from all categories. You can click on a category to see all templates in that category, and you can click a template in the list of popular templates, enter a search term, or even click **See All** to see all templates that are available.

Microsoft Azure

Search resources, services, and docs

<<

Dashboard > New

Create a resource

Home

Dashboard

All services

FAVORITES

All resources

Resource groups

App Services

App Service plans

SQL databases

Virtual machines

Load balancers

Storage accounts

Virtual networks

Security Center

Cost Management + Billing

Help + support

Azure Cosmos DB

Container services (deprec...

Azure Active Directory

Monitor

Advisor

New

Search the Marketplace

Azure Marketplace

See all

Popular

Get started

Recently created

Compute

Networking

Storage

Web

Mobile

Containers

Databases

Analytics

AI + Machine Learning

Internet of Things

Integration

Security

Identity

Developer Tools

Management Tools

Software as a Service (SaaS)

Blockchain

Windows Server 2016 VM

Quickstart tutorial

Ubuntu Server 18.04 VM

Learn more

Web App

Quickstart tutorial

SQL Database

Quickstart tutorial

Serverless Function App

Quickstart tutorial

Cosmos DB

Quickstart tutorial

Kubernetes Service

Quickstart tutorial

DevOps Project

Quickstart tutorial

Storage Account

Quickstart tutorial

Figure 2-24 The Azure Marketplace

If you click **See All**, you'll be taken to the full Marketplace experience where you can filter based on pricing, operating systems, and publisher as shown in Figure 2-25.



Figure 2-25 Filtering the Azure Marketplace



Exam Tip

All of the templates in the Azure Marketplace are ARM templates that deploy one or more Azure services. Remember from our earlier discussion of Azure Resource Manager that all ARM deployments are deployed using ARM templates. The Marketplace is no different.

Some of the templates in the Marketplace deploy a single resource. For example, if you click on the Web App template, it will create a Web App running in Azure App Service. Other templates create many resources that combine to make an entire solution. For example, you can create a DataStax Enterprise database cluster and the template will create between 1 and 40 DataStax Enterprise nodes.

You are billed for each Marketplace offering on your Azure invoice, so if you create a DataStax Enterprise cluster with 40 nodes, you won't see separate billing for 40 VMs, VNets, and so on. Instead, you'll see a bill for a DataStax Enterprise cluster. This makes billing much easier to understand.

As shown in [Figure 2-26](#), many Marketplace templates provide links to documentation and other information to help you get the most out of the

template. If you decide that you don't want to immediately create the resources, you can click **Save For Later** and the template will be added to your saved list that you can access by clicking **My Saved List** as shown in the upper-left corner in [Figure 2-25](#).




Figure 2-26 Marketplace links and Saved List

SKILL 2.3: DESCRIBE SOME OF THE SOLUTIONS AVAILABLE ON AZURE

In the Skill 2.2 section, you learned about some of the core products in Azure. In this section, you will learn about some of the most cutting-edge technologies that

are available in Azure today. This includes the Internet of Things (IoT), Big Data and analytics, artificial intelligence (AI), and serverless computing in Azure.

This section covers:

- Internet of Things (IoT)
- Big Data and analytics
- Artificial Intelligence
- Serverless computing

Internet of Things (IoT)

Many of us don't live in high-tech smart homes, so we might not realize just how big IoT is becoming. To put it into context, the popular statistics portal Statista reports that there are over 25 billion IoT connected devices today, and that number is expected to grow to a staggering 75 billion by the year 2025. There are approximately 3.2 billion people on the Internet today, and the entire world's population is only around 8 billion. These IoT devices eclipse the human race in number, and the amount of information they collect and share is mind-boggling.

To help companies manage devices and handle the data they are collecting, Azure has several services that are targeted at IoT, including IoT Hub and IoT Central.

Azure IoT Hub

In order to make more sense out of Azure's IoT services, let's consider a theoretical company named ContosoPharm, which in this example is a pharmaceutical company with a large, multi-story building where they store drugs under development, along with sensitive components used in research. These items must be under strict climate control. If the temperature or humidity moves outside of a very tight range, it results in the loss of priceless materials.

In order to protect their investment, ContosoPharm uses IoT connected climate-control systems, along with IoT connected generators and lighting systems. These systems constantly monitor the environment and send alerts if something goes wrong. There are approximately 5,000 IoT devices in the building, and ContosoPharm must meet the following requirements for all those devices.

- They must update firmware on the IoT devices easily, and in a staged way, so that all of them aren't updated at the same time.
- They must alter the settings on the devices, such as changing alert levels, but these settings are specific to the physical location of the devices in the building.
- They must ensure that any connectivity to the devices is completely secure.

IoT Hub can easily solve all of these problems. IoT

devices are added to IoT Hub, and you can then manage them, monitor them, and send messages to them, either individually or to groups that you create. You can add up to 1,000,000 IoT devices to a single IoT Hub.

Figure 2-27 shows an IoT device added to the IoT Hub for ContosoPharm.




Figure 2-27 An IoT device in IoT Hub

From IoT Hub, you can send messages to devices (called cloud-to-device, or C2D messaging) or from your

device to IoT Hub (called device-to-cloud, or D2C messaging). You can also intelligently route messages to Event Hub, Azure Storage, and Service Bus based on the content in the message.

When you add a new IoT device, IoT Hub creates a connection string that uses a shared access key for authentication. This key prevents unauthorized access to your IoT Hub. Once connected, messages between your device and IoT Hub are encrypted for additional security.

In addition to messages, you can also use IoT Hub to send files to your devices. This allows you to easily update the firmware on your devices in a secure way. To update the firmware on an IoT device, you simply copy the firmware to the device. The device will detect the firmware and will reboot and flash the new firmware to the device.

One important concept in IoT Hub is the concept of what's called a *device twin*. Every IoT device in IoT Hub has a logical equivalent that's stored in IoT Hub in JSON format. This JSON representation of the device is called a device twin, and it provides important capabilities.

Each device twin can contain metadata that adds additional categorization for the device. This metadata is stored as tags in the JSON for the device twin, and it's

not known to the actual device. Only IoT Hub can see this metadata. One of ContosoPharm's requirements was to update firmware in a staged way instead of updating all devices at the same time. They can achieve that by adding tags for the device twins from their devices that might look like the following:

[Click here to view code image](#)

```
"tags": {  
  "deploymentLocation": {  
    "department": "researchInjectibles",  
    "floor": "14"  
  }  
}
```

They can then choose to send firmware files only to devices on the 14th floor, for instance, or say to devices in the researchInjectibles department. [Figure 2-28](#) shows the device twin configuration in IoT Hub with tags set for the location of the device. Notice the “building” tag with a value of null. This is a tag that was previously set on the device twin, and by setting it to null, the tag will be removed.




Figure 2-28 Device twin showing tags set in the JSON

The device twin also contains the properties for the IoT device. There are two copies of every property. One is the “reported” property, and the other is the “desired”

property. You can change a device property in IoT Hub by changing the “desired” property to a new value. The next time the device connects to IoT Hub, that property will be set on the device. Until that happens, the “reported” property will contain the last value the device reported to IoT Hub. Once the property is updated, the “reported” and “desired” property will be equal.

The reason IoT Hub uses this method for setting properties is that it may not always have a connection to every device. For example, if a device puts itself to sleep to save power, IoT Hub can’t write property changes to that device. By keeping a “desired” and “reported” version of every property, IoT Hub always knows if a property needs to be written to a device the next time the device connects to IoT Hub.

To help with users who want to add a large number of IoT devices to IoT Hub, Microsoft offers the IoT Hub Device Provisioning Service, or DPS. The DPS uses enrollment groups to add devices to your IoT Hub. The concept is that once the device wakes up (oftentimes for the first time if it’s a new device), it needs to know that it should connect to your IoT Hub. In order to do that, the DPS needs to uniquely identify the device, and it does that with either a certificate or via a trusted platform module chip.

Once DPS confirms the identity of the device, it can

use the enrollment group details to determine which IoT Hub it should be added to. It will then provide the device with the connection information to connect to that IoT Hub. In addition to that, the enrollment group can also provide the initial configuration for the device twin. This allows you to specify properties such as a firmware version that the device needs to have in when it starts.

As your devices send messages to IoT Hub, you can route those messages to Azure Storage, Event Hub, and various other endpoints. You can choose the type of messages you want to route, and you can also write a query to filter which messages are routed. In [Figure 2-29](#), we have configured a route that sends messages to Azure Blob Storage. You can see in the query that we are only going to route those messages that come from a device with a device twin containing the tag for our researchInjectibles department.

Dashboard > ContosoPharmHub - Message routing > Add a route

Add a route

* Name ⓘ
 ✓

* Endpoint ⓘ
 ✓ [+ Add](#)

* Data source ⓘ
 ✓

* Enable route ⓘ

Create a query to filter messages before data is routed to an endpoint. [Learn more](#)

Routing query ⓘ

```
1 $twin.tags.deploymentLocation.department = "researchInjectibles"
```

Figure 2-29 Adding a message route in IoT Hub

There are two pricing tiers for IoT Hub: Basic and Standard. Each tier offers multiple editions that offer pricing based on the number of messages per day for each IoT Hub unit. When you scale an IoT Hub, you add additional units. This adds the ability to handle more messages at an increased price. [Table 2-4](#) shows the editions and pricing for the Basic tier. [Table 2-5](#) shows editions and pricing for the Standard tier.

Table 2-4 IoT Hub Basic tier pricing

Edition	Monthly Price per IoT Hub Unit	Messages per day per IoT Hub Unit
B1	\$10	400,000
B2	\$50	6,000,000
B3	\$500	300,000,000

Table 2-5 IoT Hub Standard tier pricing

Edition	Monthly Price per IoT Hub Unit	Messages per day per IoT Hub Unit
Free	Free	8,000
S1	\$25	400,000
S2	\$250	6,000,000
S3	\$2,500	300,000,000

***Exam Tip***

Pricing for scale in IoT Hub is pretty clear. Most enterprises will choose the Standard tier because of the additional functionality available in that tier. They will then choose an edition that meets their minimal needs for messages. When they need additional

messages during spikes, they'll scale to more IoT Hub units.

For example, assume that ContosoPharm message needs are approximately 5,000,000 per day. They would choose the S2 pricing tier and pay \$250 per month if they are running 1 IoT Hub unit. If the number of messages increase to 8,000,000 (either due to configuration changes or the addition of additional IoT devices), they would likely choose to scale to 2 IoT Hub units. Doing so would give them 12,000,000 messages per day at a cost of \$500 per month.

Note Changing Pricing Tier

You cannot change to a lower pricing tier after you create your IoT Hub. If you create your IoT Hub in the Standard tier, you cannot change it to the Basic tier. If you create an IoT Hub in the Standard tier using the S1, S2, or S3 edition, you cannot change it to the Free edition.

It's also important to note that the following features are only available in the Standard tier.

- Device Streams for streaming messages in near real-time
- Cloud-to-device messaging

- Device management, device twin, and module twin
- IoT Edge for handling IoT Devices at the edge of the network where they reside

If you use the Device Provisioning Service, there's a charge of \$0.10 for every 1,000 operations.

Azure IoT Central

IoT Hub is a great way to manage and provision devices, and it provides a robust means of dealing with messages. You can even use Azure Stream Analytics to route messages to Power BI for a near real-time dashboard of device messages, but doing that requires a bit of complex configuration. If you're looking for a first-class experience in monitoring IoT devices without having to do complex configuration, IoT Central is a good choice.

IoT Central is a SaaS offering for IoT devices. Unlike IoT Hub, you don't have to create any Azure resources to use IoT Central. Instead, you browse to <https://apps.azureiotcentral.com> and create your app within the web browser interface as shown in [Figure 2-30](#).




Figure 2-30 The homepage of Azure IoT Central

To create an IoT app, click on **New Application**. This opens the Create Application screen shown in [Figure 2-31](#), where you can choose either the Trial plan (which does not require an Azure subscription) or Pay-As-You-Go using an Azure subscription. If you choose a Trial plan, you have 7 days to test IoT Central with any number of devices at no charge, and you can upgrade the app to Pay-As-You-Go at any time within those 7 days. If you choose Pay-As-You-Go, you pay based on the number of devices you have, but the first five devices are always free.

Create Application

We just need a few things from you, so we can create your application

Choose payment plan

☒ Trial	☐ Pay-As-You-Go
Free trial for 7 days. No subscription required.	Price is based on the number of devices you use. Free for the first 5 devices. Subscription required. Learn more

Select an application template

☒ Sample Contoso	☐ Sample Devkits	☐ Custom Application
Get started with a predefined application for a connected device.	Want to connect a Raspberry PI or MXChip IoT DevKit? Start with this predefined app and get them connected in minutes.	Start with a blank template and define your application from scratch.

Figure 2-31 Creating a new IoT Central app

You also have the choice of choosing a template, or creating a blank template. The Sample Contoso

template creates a sample app with a simulated refrigerated vending machine device. If you have a Raspberry PI or an MXChip IoT DevKit from the Azure IoT Starter Kit, you can use the Sample Devkits template. It has device templates so that you can add these devices to your IoT Central app. Finally, the Custom Application template allows you to start from scratch and add any IoT devices you may have.

After you select your template, scroll down to specify the name for your app and the URL. You can use the default names or specify your own, but it's recommended to use your own so you can easily identify your app. Also, once your app has been created, you access it directly by using the URL you specify, so you may want that to be descriptive as well.

If you're using Pay-As-You-Go, you'll need to specify an Azure Active Directory associated with your subscription, your Azure subscription, and the region where you want to create your app. (It's best to choose a region that's geographically close to your IoT devices if possible.) Click **Create** to finish the creation of your app as shown in [Figure 2-32](#).

Application Name * ⓘ

Sample Contoso 1ku8dr89jfh

URL * ⓘ

sample-contoso-1ku8dr89jfh .azureiotcentral.com

Directory * ⓘ

Earth-Shaking Sites (jameschelive.onmicrosoft.com) ▼

Azure Subscription * ⓘ Don't have a subscription? [Create subscription](#) ⓘ

Jim's Personal Azure Account ▼

Region * ⓘ

East US ▼

By clicking "Create" you agree to the [Subscription Agreement](#) ⓘ and [Privacy Statement](#) ⓘ. Provisions in the agreement with respect to pricing, cancellation fees, payment, and data retention do not apply to "Trial". "Pay-As-You-Go" requires an Azure subscription, and you acknowledge that this service is licensed to you under the terms applicable to your [Azure Subscription](#) ⓘ.

Create

Figure 2-32 Specifying an app name, URL, and Azure subscription information

In [Figure 2-30](#), you can see that we’ve already created an app called ContosoPharm. When you click on that app, you see a menu on the left side of the page, and if you click on Device Explorer, you can see any devices added as shown in [Figure 2-33](#).




Figure 2-33 The IoT device in IoT Central

Add a new device by clicking the plus sign as shown in [Figure 2-34](#). You have the option of adding a real device if you have one, but you can also add a simulated device.

Adding simulated devices is a good way to get everything set up the way you want them in IoT Central and then you can add real devices at a later time.



Figure 2-34 Adding a device in IoT Central

Note Simulated Devices Is An IoT Central-Only Feature

The ability to create a simulated device is specific to IoT Central. IoT Hub doesn't offer this capability.

Every page within your app can be edited directly in the browser. [Figure 2-35](#) shows the home page for the IoT Central app. If you click on the Edit button, you can remove tiles, add tiles, and edit information in tiles in a point and click interface right within my browser.




Figure 2-35 Editing a page in IoT Central

The reason we see an Edit button is because this user is set as the administrator of this application. IoT Central gives you control over who can do what using roles. There are three built-in roles you can assign a user to.

- **Application Administrator** Users in this role have full access to the application and can edit page and add new users.
- **Application Builder** Users in this role can edit pages, but they can't perform any administrative tasks such as adding users, changing user roles, changing application settings, and so on.
- **Application Operator** Users in this role can use the application, but they can't edit any pages and they can't perform administrative tasks.

In some situations, these built-in roles may not offer the flexibility you need, so Microsoft is working on allowing you to define your own roles with custom permissions.

To administer your application, click on **Administration** on the menu on the left as shown in [Figure 2-36](#). You can then add and remove users, adjust user roles, change the application name or URL, add a custom image for your application, and so on. You can also copy or delete your application from this screen.

ContosoPharm

Search



<

Home

Device Explorer

Device Sets

Analytics

Jobs

Application Builder

Continuous Data Export

Administration

Administration

Application Settings

Users

Roles


Billing

Device Connection

Access Tokens

Save

Application Image ⓘ



Application Name * ⓘ

ContosoPharm

Application URL * ⓘ

contosopharmazureiotcentral.com

Copy Application ⓘ

Create a copy of this application, minus any device instances, device data history, and user data. The copy will be a Pay-As-You-Go application that you'll be charged for.

Copy

Figure 2-36 Administering an application in IoT Central

If you click on a device, you can look at information coming from the device's sensors. In [Figure 2-37](#), you can see the humidity and temperature sensors on a F14_TempMonitor device. Humidity is the top line and temperature is the bottom line. As you can see, we're experiencing a small rise in temperature and a pretty strong rise in humidity.




Figure 2-37 Administering an application in IoT Central

If you want a better view of data from your device, you can click on Dashboard as shown in the top of the screen in [Figure 2-37](#). The dashboard, like other pages in your application, is customizable so that you can see exactly the data you want. [Figure 2-38](#) shows a dashboard created for a device.




Figure 2-38 Creating a dashboard for your device

Note Dashboards

Dashboards are for a single device. If you want to see customized information for more than one device, you can add tiles for the devices to your home page located at https://<your_app_name>.azureiotcentral.com.

IoT Central also allows you to easily configure rules that will monitor your devices and perform an action you choose when your rule is activated. In [Figure 2-39](#), we are configuring a rule that will activate when humidity reaches 60 or above. Notice that we also have a live historical view of the metric in a graph on the right so that you can make more intelligent decisions about the thresholds.




Figure 2-39 Creating a rule

When a rule is triggered, IoT Central can send an email to someone with details of what happened. You can also choose to trigger a webhook, make a call to an Azure Function, run a workflow in an Azure Logic App, or run a workflow in Microsoft Flow. These options provide the flexibility to perform almost any task when a rule is triggered.

When you have a large number of devices, it's convenient to group devices into a device set so that you can take action on many devices at a time. To create a device set, specify a condition that should be met for a device to be added to the set. In [Figure 2-40](#), we're creating a device set for all devices that have F14 in the name. If the name contains "F14," the device is automatically added to the device set. Even when adding a new device at a later time, it will be part of this device set if the name contains "F14."

the main menu of your application to configure your job. A job can modify properties, change settings, or send commands to devices. In Figure 2-41, we create a job that will turn the IR sensor on for all devices in our device set.

Jobs

▷ Run ✕ Cancel

Create job

Name * ⓘ
Set IR Sensor

Description ⓘ
Turn on the IR sensor on devices.

Device set * ⓘ
Fourteenth Floor Devices

Job type * ⓘ
Settings

Settings +

Set IR to ON ✕
IR

 ON

 MXChip (1.0.0)

☒ 1 / 1 device selected

☒	Name	Manufactured In	Die Number	Device Location
☒	F14_TempMonitor		4	47°36'23"N, 122°19'56"W

Figure 2-41 Creating a job

IoT Central also allows you to perform analytics on metrics from devices in a device set. For example, you

can look at all devices that registered temperatures above a certain level. For even richer analytics of data, you can configure IoT Central to continuously export data from your devices to Azure Blob Storage, Azure Event Hubs, or Azure Service Bus.

Big Data and analytics

Businesses collect tremendous amounts of data from many different sources. As you've already learned, Microsoft offers an SLA on Azure services that are in the area of 99.9%+ for availability. Microsoft doesn't put that number out there and then just cross their fingers that nothing goes wrong. They maintain enormous amounts of data on how the Azure infrastructure is operating, and they use that data to predict problems and react to them before they impact customers.

Because of the sheer enormity of the Azure infrastructure, you can just imagine how much data is being generated for every single system in that infrastructure, and in order to meet SLAs, they have to be able to reliably analyze that data in real time. How exactly do they do that? You can't really throw that amount of data at a VM or a pool of VMs without overloading the system to the point of failure.

The problem of actually doing anything with the vast data we collect is common across all businesses, and this is what we mean by *big data*. Big data means more

data than you can analyze through conventional means within the desired time-frame.

By placing big data in a data warehouse, you can then use massive amounts of computing power to analyze multiple pieces of data in parallel, and you can perform analysis on the data much more quickly than you otherwise could.

We'll get into the analysis of big data later in this chapter. First, let's talk about where to store big data. Microsoft has two Azure offerings for storing big data for analysis: Azure SQL Data Warehouse and Azure Data Lake Storage. They are similar in purpose but quite different in design.



Exam Tip

Azure Blob Storage can also be used as a data store for big data. However, SQL Data Warehouse and Data Lake Storage are explicitly designed for this purpose. Microsoft has also recently released Data Lake Storage Gen2, which combines the features of Blob Storage with Data Lake Storage, so the usage of Blob Storage for data warehousing is becoming unnecessary.

Azure SQL Data Warehouse

Azure SQL Data Warehouse is designed for storing big data that's in the form of relational data. Data stored in SQL Data Warehouse is in a form quite similar to tables in a SQL Server database, and in fact, when you analyze data in SQL Data Warehouse, you run complex SQL queries against the data.

SQL Data Warehouse provides secure authentication using both SQL Server Authentication in the connection string, which is username and password authentication, and Azure Active Directory. Once a user is authenticated, you can only perform actions that you've been authorized to perform, and that authorization is controlled via database permissions.

While your data is in SQL Data Warehouse, it's encrypted using Transparent Data Encryption (TDE) AES-256 encryption. Data is encrypted using a database encryption key, and this key is protected by a server certificate that's unique to each SQL Database server. These certificates are rotated by Microsoft at least every 90 days, so you can be assured that your data is safe.

SQL Data Warehouse uses several methods to control costs. In fact, in a recent study by GigaOm, SQL Data Warehouse was found to be 94% less expensive than Google BigQuery and up to 31% less expensive than Amazon AWS Redshift. (SQL Data Warehouse was also

much faster in benchmarks than either of these offerings.) One way SQL Data Warehouse reduces costs is by decoupling the data storage from compute resources. This allows you to easily scale to more compute resources when you need them, and then scale back down to save money when you no longer need the power.

There are two performance tiers available in SQL Data Warehouse, and both of them support scaling up or down and pausing resources so you don't pay for them. The Gen1 performance tier measures compute resources in Data Warehouse Units, or DWUs. When you scale Gen1 data warehouses for more power, you add DWUs. The Gen2 tier uses compute Data Warehouse Units, or cDWUs. The difference is that Gen2 uses a local disk-based cache in order to improve performance. As long as you don't scale or pause the data warehouse, the cache will substantially improve performance. If you do scale or pause, when the data warehouse is restarted, the cache will have to be refreshed, and you won't experience the same performance improvement until that refresh is complete.

To use SQL Data Warehouse, you create an instance of it in Azure and then you load data into it using either queries against the database or by using tools like ADF Copy, SQL Server Integration Services, or the command

line. You can then run complex queries against your data. Because of the power of SQL Data Warehouse, queries that would otherwise take several minutes to run can run in seconds, and a query that might take days to complete can finish in hours. Finally, you can use Microsoft's Power BI to gain important insight into your data in an easy-to-use web browser-based environment.

More Info Migrating Data To Sql Data Warehouse

For more information on migrating data to SQL Data Warehouse, see: <https://docs.microsoft.com/azure/sql-data-warehouse/sql-data-warehouse-migrate-data>.

Azure Data Lake Storage

Like SQL Data Warehouse, Azure Data Lake Storage is designed for storing large amounts of data that you'd like to analyze, but Data Lake Storage is designed for a wide array of data instead of relational data. In a data lake, data is stored in *containers*. Each container typically contains related data.

Note Not Just Azure

The terms data lake and data warehouse aren't specific to Azure. They are generic terms. A data lake refers to a repository of unordered data, and a data warehouse refers to a repository of ordered data.

The two common modes of accessing data are object-

based (such as Azure Blob Storage) and file-based. In an object-based mode, there isn't a hierarchy of objects. You simply store the object in a flat model. Traditional data lakes use the object-based access mode, but using this mode isn't always efficient because it requires that you individually interact with each object.

With the introduction of Data Lake Storage Gen2, Microsoft introduced the concept of a hierarchical namespace to storage. This organizes the objects into a system of directories much like the structure of the files on your computer, and it allows for the use of both object-based and file-based models in the same data lake. Microsoft calls this capability *multi-modal storage*, and Data Lake Storage Gen2 is the first cloud-based solution to offer this capability.

Data Lake Storage is ideal for performing analysis against large amounts of data that aren't stored in a relational way. For example, the vast amounts of information that Google or Facebook might have stored on users can be held in a data lake for analysis. However, data from a data lake often isn't ideal for presentation to users in a way that's easy to understand. People work better with data that is relational, and for that reason, it's pretty common for data to be analyzed in a data lake and then structured and moved into a data warehouse for further analysis and presentation.

Like Azure Blob Storage, Data Lake Storage is available in Hot, Cool, and Archive tiers. The Hot tier has the highest storage costs, but the lowest access costs. The Archive tier has the lowest storage costs, but it has the highest access costs.

If you use the file-based storage in Data Lake Storage Gen2, there are some additional costs for the metadata necessary to implement the file-based hierarchy. There are also some additional costs associated with operations that require recursive calls against the hierarchy. Data Lake Storage supports several open source data analytics platforms, including HDInsight, Hadoop, Cloudera, Azure Databricks, and Hortonworks.

Once you have data available in SQL Data Warehouse or Data Lake Storage, you can use one of Azure's analytic services to analyze the data, including Azure HDInsight and Azure Databricks.

More Info Azure Databricks

Because Azure Databricks is a big data service most often used with machine learning, we'll discuss it later in this chapter when we cover artificial intelligence.

Azure HDInsight

HDInsight makes it possible to easily create and manage clusters of computers on a common framework designed to perform distributed processing of big data.

HDInsight is essentially Microsoft’s managed service that provides a cloud-based implementation of a popular data analytics platform called Hadoop, but it also supports many other cluster types as shown in [Table 2-6](#).

Table 2-6 HDInsight supported cluster types

Cluster Type	Description
Hadoop	Large-scale data processing that can incorporate additional Hadoop components such as Hive (for SQL-like queries), Pig (for using scripting languages), and Oozie (a workflow scheduling system.)
HBase	Extremely fast and scalable NoSQL database.
Storm	Fast and reliable processing of unbounded streams of data in real-time.
Spark	Extremely fast analytics using in-memory cache across multiple operations in parallel.
Interactive Query	In-memory analytics using Hive and LLAP (processes that execute fragments of Hive queries).
R Server	Enterprise-level analytics using R, a language that’s specialized for big data analytics.
Kafka	Extremely fast processing of huge numbers of synchronous data streams, often from IoT devices.

Building your own cluster is time-consuming and often difficult unless you have previous experience. With HDInsight, Microsoft does all of the heavy lifting

on their own infrastructure. You benefit from a secure environment, and one that is easily scalable to handle huge data processing tasks.

An HDInsight cluster performs analytics by breaking up large data blocks into segments that are then handed off to nodes within the cluster. The nodes then perform analytics on the data and reduce it down to a result set. All of this work happens in parallel so that operations are completed dramatically faster than they would be otherwise. By adding additional nodes to a cluster, you can increase the power of your analytics and process more data even faster.

When you create an HDInsight cluster, you specify the type of cluster you want to create and give your cluster a name as shown in [Figure 2-42](#). You will also specify a username and password for accessing the cluster and an SSH user for secure remote access.

Dashboard > New > Marketplace > Everything > HDInsight > HDInsight > Basics

HDInsight

by Microsoft

Quick create Custom (size, settings, apps)

- 1 Basics**
Configure basic settings >
- 2 Storage**
Set storage settings >
- 3 Summary**
Confirm configurations >

i This cluster may take up to 20 minutes to create.

Basics

* Cluster name
jwc ✓
.azurehdinsight.net

* Subscription
Jim's Personal Azure Account ▾

* Cluster type ⓘ
Hadoop 2.7 (HDI 3.6) >

* Cluster login username ⓘ
admin ✓

* Cluster login password ⓘ
..... ⓘ ✓

Secure Shell (SSH) username ⓘ
sshuser

☒ Use same password as cluster login ⓘ

* Resource group
(New) HDIrg ▾
[Create new](#)


* Location
East US 2 ▾

Next

Figure 2-42 Creating an HDInsight Hadoop cluster

After you click the Next button, you configure the storage account and Data Lake Storage access if desired. Notice in Figure 2-43 that you only see Data Lake Storage Gen1. To use Data Lake Storage Gen2, you must create the storage account first and complete some additional configuration as outlined at:

<https://azure.microsoft.com/blog/azure-hdinsight-integration-with-data-lake-storage-gen-2-preview-acl-and-security-update/>.



The screenshot shows the 'Storage' configuration window with the following settings:

- Storage Account Settings**
 - Primary storage type**: Azure Storage (selected)
 - Selection method**: My subscriptions (selected), Access key (unselected)
 - Create a new Storage account**: jchdistore (with a green checkmark), [Select existing](#)
 - Default container**: jwc-2019-02-10t22-36-23-422z
- Additional storage accounts**: Optional >
- Data Lake Storage Gen1 access**: Optional > (highlighted with a dashed blue border)

Figure 2-43 Configuring an HDInsight cluster's storage account

Note Quick And Custom Create

The quick create process in Figures 2-42 and 2-43 creates six Hadoop nodes with 40 cores. If you want a different configuration, you can click Custom (shown in Figure 2-42) to specify your own settings.

Once you start the creation of your Hadoop cluster, it may take up to 20 minutes to complete, depending on your configuration. Once your cluster is ready, you can start the analysis of data by writing queries against it. Even if your queries are analyzing millions of rows, HD Insight can handle it, and if you need more processing power, you can add additional nodes as needed.

HD Insight clusters are billed on a per hour basis, and you pay more per hour based on how powerful the machines are in your cluster. For full pricing details, see:

<https://azure.microsoft.com/pricing/details/hdinsight/>.

Artificial Intelligence

Let's circle back through what we've learned up to this point. We know that the number of IoT devices far surpasses the number of humans, and those IoT devices are generating enormous amounts of data. It's pretty clear that there is a mind-boggling amount of data being collected.

We learned about technologies that allow us to store

this tremendous amount of data and how we can keep it safe and access it quickly. What we haven't talked about is what we do with all of that data. That's where artificial intelligence (AI) comes into the picture.

Before we go too far into AI, let's first come to an agreement on what we mean by AI. When many people think about computer AI, the image that comes to mind is a human-killing android or some other hostile technology obsessed with ridding the world of humans. You'll be relieved to know that's actually not what AI means in this context.

The AI of today is called Artificial Narrow Intelligence (or sometimes weak AI), and it refers to an AI that is capable of performing one specific task much more efficiently than a human can perform that same task. All of the AI that we've developed so far is weak AI. On the other end of the AI spectrum is Artificial General Intelligence, or strong AI. This is the type of AI you see depicted in movies and science fiction books, and we don't currently have this kind of capability.

In many ways, it's a bit misleading to call existing AI technology weak. If you place it in the context of the imaginary strong AI, it certainly has limited capabilities, but weak AI can do extraordinary things, and you almost certainly benefit from its capabilities every day. For example, if you speak to your phone or your smart

speaker and it understands what you've said, you've benefitted from AI.

In the 1973 edition of *Profiles of the Future*, the famous science fiction writer Arthur C. Clarke said, "Any sufficiently advanced technology is indistinguishable from magic." While AI was not yet a thing when Clarke made this assertion, the capabilities that AI make possible are certainly applicable, but AI isn't magic. AI is actually mathematics, and as anyone familiar with computers will tell you, computers are very good at math.

In order to develop AI capabilities, computer engineers set out to give computers the ability to "learn" in the same way that the human brain learns. Our brain is made up of neurons and synapses. Each neuron communicates with all of the other neurons in the brain, and together, they form what's known as a neural network. While each neuron on its own can't do much, the entire network is capable of extraordinary things.

AI works by creating a digital neural network. Each part of that neural network can communicate and share information with every other part of the network. Just like our brains, a computer neural network takes input, processes it, and provides output. AI can use many methods for processing the input, and each method is a subset of AI. The two most common are natural

language understanding and machine learning.

Natural language understanding is AI that is designed to understand human speech. If we were to try and program a computer to understand the spoken word by traditional computing means, it would take an army of programmers decades to come anywhere close to usable recognition. Not only would they have to account for accents and vocabulary differences that occur in different geographic regions, but they'd have to account for the fact that individuals often pronounce words differently even in the same regions. People also have difference speech cadences, and that causes some words to run together. The computer has to know how to distinguish individual words when that might not be easy to do. In addition to all of this complexity, the computer has to account for the fact that language is an ever-changing thing.

Given this complexity, how did Amazon ever develop the Echo? How does Siri ever understand what you're saying? How does Cortana know to crack a clever joke when we ask her about Siri? The answer in all of these cases is AI. We have millions of hours of audio recordings, and we have millions more hours in videos that include audio. There's so much data available that no human being could ever process all of it, but a computer processes data much more quickly. Not only does it have more analytical pathways than humans do,

but it also processes information much more quickly.

More Info Computers Are Fast

When I say that computers can process information more quickly than humans, I really mean it! Information in a human brain travels between neurons at a speed that's just under the speed of sound. While that's plenty fast for our needs, it's nothing compared to computers. The information in an AI neural network travels at the speed of light, and that's what enables computers to process enormous amounts of data. In fact, a computer's AI system can process 20,000 years of human-level learning in just one week.

If we feed all of those recordings into a natural language understanding engine, it has plenty of examples in order to determine what words we're speaking when we say something to a smart speaker or smart phone, and determining the meaning of these words is simply pattern recognition. As Apple, Amazon, and Microsoft were working on this technology, they fine-tuned it by getting your feedback. Sometimes they might actually ask you whether they got it right, and other times, they might assume they got something wrong if you just bowed out of the conversation early. Over time, the system gets better and better as it gets more data.

Machine learning (ML) is similar in that it uses a neural network to accomplish a task, but the task is different than understanding speech. In fact, machine learning can be used in many applications. One of the

common uses of machine learning is image recognition. As it turns out, AI neural networks are particularly good at recognizing patterns in images, and just like audio, we have an enormous amount of data to work with.



Exam Tip

In ML, the process of decision making at several points along the AI neural network is known as the ML pipeline. It's a series of decisions made by the ML model that eventually results in a particular output.

Many examples of ML relate to image processing because ML is well-suited to doing that kind of work. However, a lot of ML focuses on using existing data to make a prediction about what will happen in the future, and to do that with a high degree of reliability.

We're likely all aware that satellites have been photographing the surface of the earth for quite some time. We have detailed imagery from just about every square inch of our planet, and those images are valuable in many ways. For example, scientists who are working

on conservation efforts benefit by knowing how our planet is changing over time. Forest engineers need to know about the health of our forests. Wildlife conservationists need to know where to focus efforts on where animals are most at risk. By applying an ML model to all of these images, Microsoft is able to serve all of these needs.

More Info Microsoft AI For Earth

For more information on all the ways Microsoft is using AI for conservation and earth sciences, see: <http://aka.ms/aiforearth>.

Image analyzing AI isn't limited to the planetary scale. It's also helpful when we want to search through our own pictures. Perhaps you want to find all of the pictures you've taken of a particular person, or maybe you're interested in finding all of your pictures of flowers. Your phone can likely do this kind of thing, and it does it by using AI and ML. In fact, Google Photos is even able to identify specific people in photos when the time between two photos is decades apart. All of this uses ML.

ML uses a learning algorithm that is the basis for the AI. Once the algorithm is developed, you feed test data to it and examine the result. Based upon that result, you may determine that you need to tweak the algorithm. Once the algorithm is suitable to your task, you typically

deploy it to an environment where it has a vast array of compute resources that you can allocate to it. You can then feed huge amounts of data to it for processing. As the algorithm deals with more data, it can improve itself by recognizing patterns.

When you're testing an ML model, you typically set up a scenario where only a portion of your complete dataset is sent to your model for training. Once your model is trained, you send the rest of your data through your model in order to score the results. Since you're dealing with a historical dataset, you already know that which your model is attempting to figure out, so you can accurately determine the accuracy of your model. Once you have achieved the desired accuracy of your model, you can deploy it and begin using it against production data.

Even with careful modeling and scoring, ML algorithms can make mistakes. In a paper on ML published in 2016, Marco Ribeiro, Sameer Singh, and Carlos Guestrin wrote about an ML experiment that was designed to look at pictures and differentiate between dogs and wolves. As it turns out, the algorithm was making plenty of mistakes, but the humans couldn't figure out why.

When they went back and tested the ML algorithm to determine how it was making these incorrect decisions,

they found that the algorithm had come to the conclusion that pictures with wolves in them had a snowy background and pictures with dogs had grass in the background. Therefore, every picture with a dog-like creature that was taken on a snowy background was immediately classified (sometimes incorrectly) as a wolf.

More Info AI And Trust

The wolf analogy illustrates one of the primary concerns of AI, and that is how to determine when to trust an AI model. If you want to dig into this concept, check out this paper referenced at: <https://arxiv.org/pdf/1602.04938.pdf>.

When you're dealing with developing and using AI with the enormous amount of data available today, the cloud offers some distinct advantages. You can take advantage of the enormous computing resources that cloud providers make available, and you can use those resources in time slices only when you need to do work. This makes it possible to use more powerful resources than you'd have available on-premises, and it also makes it possible to control your costs by scaling your usage.

Microsoft offers many technologies in Azure to help you with your AI and ML needs. You can even get started without doing any of your own AI and ML work by using some of the provided services that Microsoft

itself uses. These services are part of Azure Cognitive Services, and they include:

- **Computer Vision** Analyze images and recognize faces, text, and handwriting.
- **Microsoft Speech** Recognize, transcribe, and translate speech.
- **Language Understanding Intelligent Service (LUIS)** Natural language service that uses ML to understand speech and take action on it.
- **Azure Search and Bing Search** Search for specific data in order to build complex data sets.

These offerings allow you to fast-track your ML capabilities by taking advantage of work Microsoft has done to support its own services like Bing, Office 365, and more. Microsoft also provides resources you can use to build your own offerings using many of the tools that engineers are already familiar with. They even provide a feature-rich development environment called Visual Studio Code that runs cross-platform and allows for rapid development of ML models.

Microsoft also supports numerous ML frameworks that are commonly used by developers of AI solutions. These include ONNX (Open Neural Network Exchange), Pytorch, TensorFlow, and Sci-Kit Learn. This allows AI programmers (known as *data scientists*) to get started in Azure without having to learn new frameworks and techniques.

Azure services aimed at data scientists run the frameworks mentioned above. These services include Azure Databricks, Azure Machine Learning Service, and Azure Machine Learning Studio. Powering these services are infrastructure components especially well-suited to AI and ML.

Azure Databricks

We've looked at some of the Azure services for storing big data such as SQL Data Warehouse and Azure Data Lake Storage. Data that gets stored in these services is typically raw data that is often unstructured and difficult to consume in order to build a ML model. We also may need data for our ML model that comes from multiple sources, some of which may even be outside of Azure. Azure Databricks is an ideal solution for accumulating data and for forming the data (called *data modeling*) so that it's optimal for ML models.

Figure 2-44 shows a new instance of an Azure Databricks resource. All of your interactivity with Databricks is via the Databricks workspace, a web-based portal for interacting with your data, and to access the workspace, click on the **Launch Workspace** button shown in [Figure 2-44](#).




Figure 2-44 An instance of Azure Databricks in the Azure portal

When clicking on **Launch Workspace**, you're taken

to the Databricks workspace. Azure will automatically log you in when you do this using your Azure account. My Databricks instance is completely empty at this point. Along the left side of the page (as shown in [Figure 2-45](#)) are links to access all of the Databricks entities such as workspaces, tables, and jobs. There's also a Common Tasks section, which allows you to access these entities, as well as to create new notebooks, which are detailed soon.




Figure 2-45 The Azure Databricks portal

Let's now create a cluster. Databricks does all of its work using clusters, which are the compute resources. To create a cluster, you can click on **New Cluster** under **Common Tasks**. You'll now see the Create Cluster screen shown in [Figure 2-46](#), where the new cluster has been named "jcCluster," and all other options are the default.

Create Cluster

New Cluster

Cancel

Create Cluster

2-8 Workers: 28.0-112.0 GB Memory, 8-32 Cores, 1.5-6 DBU
1 Driver: 14.0 GB Memory, 4 Cores, 0.75 DBU Cost \$0.40 per DBU

Cluster Name

jcCluster

Cluster Mode

Standard

Databricks Runtime Version

Runtime: 5.2 (Scala 2.11, Spark 2.4.0)

Python Version

3

Autopilot Options

☒ Enable autoscaling

☒ Terminate after 120 minutes of inactivity

Worker Type	Min Workers	Max Workers
Standard_DS3_v2 14.0 GB Memory, 4 Cores, 0.75 DBU	2	8

Driver Type

Same as worker 14.0 GB Memory, 4 Cores, 0.75 DBU

Advanced Options

Figure 2-46 Creating a Databricks cluster

Next, we'll create a notebook. Notebooks are a powerful way to present and interact with data that is related. Each notebook contains not only data, but also visualizations and documentation of that data to help us

understand the data. Once your data is in your notebook, you can run commands against ML frameworks in order to build your ML model right inside of your notebook.

Clicking the Azure Databricks button in the menu on the left (shown in [Figure 2-45](#)) allows you to then click on **New Notebook** to create a notebook. In [Figure 2-47](#), we create a new notebook that uses SQL as the primary language. Databricks will assume that the code written in this notebook will be SQL code unless specified. You can also choose to specify Python, Scala, or R as the language.




Figure 2-47 Creating a notebook

After you create a new notebook, you'll see an empty notebook with one cell. Inside of that cell, you can enter any data that you wish. For example, you might want to have some documentation that defines what this notebook contains. Documentation in notebooks is

entered using *markdown*, a language that's well-suited to writing documentation. [Figure 2-48](#) shows the new notebook with some markdown that documents what's in the notebook. Notice that the markdown starts with “%md.” This tells Databricks how the content that follows is in markdown and not in the primary language of SQL.




Figure 2-48 Documenting a notebook using markdown

If you click outside of this cell, the markdown code will be rendered in HTML format. In order to add some data to this notebook, you need to create a new cell by pressing “B” on your keyboard or by hovering over the existing cell and clicking the “+” button to add a new cell.

Note Keyboard Shortcuts

Keyboard shortcuts are by far the fastest way of working in Databricks. You can find the entire list of keyboard shortcuts by clicking the “Shortcuts” link shown in [Figure 2-48](#).

After pressing “B” on your keyboard, a new cell is inserted at the end of your notebook. You can enter some SQL code in this cell in order to populate a table with some data as shown in [Figure 2-49](#). (This code was taken from the Databricks quick start tutorial at <https://docs.azuredatabricks.net/getting-started/index.html>.) After entering your code, you can run it by clicking on the **Run** button.




Figure 2-49 Adding code and running a command

More Info Where Data Comes From

Notice that the path entered for the data starts with/databricks-datasets. When creating a cluster you gain access to a collection of datasets called Azure Databricks Datasets. Included in these datasets is some sample data in a comma-separated values format, and the specified path points to that data. When this command runs, it pulls that data into your notebook.

You can run a query against the data that was added using the command shown in [Figure 2-49](#) by writing a SQL query in a new cell. [Figure 2-50](#) shows the results of a query against the data.

The screenshot shows the Databricks SQL interface. The top section shows a SQL command to drop and create a table. The bottom section shows a SQL query that groups diamonds by color and calculates the average price.

Command 2:

```
1 DROP TABLE IF EXISTS diamonds;
2
3 CREATE TABLE diamonds USING CSV OPTIONS (path "/databricks-datasets/Rdatasets/data-001/csv/ggplot2/diamonds.csv", header "true")
```

Command took 9.09 seconds -- by janesche@live.com at 2/16/2019, 3:04:24 PM on jcCluster

Command 3:

```
1 SELECT color, avg(price) AS price FROM diamonds GROUP BY color ORDER BY COLOR
```

Command took 2.58 seconds -- by janesche@live.com at 2/16/2019, 3:12:00 PM on jcCluster

color	price
D	3169.9540959409596
E	3076.7524752475247
F	3724.886396981765
G	3999.135671271697
H	4486.669195568401
I	5091.874953891553
J	5323.81801994302

Figure 2-50 Querying my data

When you run commands in a cell, Databricks creates a job that runs on the compute resources you allocated

to your cluster. Databricks uses a serverless model of computing. That means that when you're not running any jobs, you don't have any VMs or compute resources assigned to you. When you run a job, Azure will allocate VMs to your cluster temporarily in order to process that job. Once the job is complete, it releases those resources.

This example is quite simple, but how does all of this relate to ML? Azure Databricks includes the Databricks Runtime for Machine Learning (Databricks Runtime ML) so that you can use data in Databricks for training ML algorithms. The Databricks Runtime ML includes several popular libraries for ML, including: Keras, PyTorch, TensorFlow, and XGBoost. It also makes it possible to use Horovod for distributed deep learning algorithms. You can use these components without using Databricks Runtime ML. They're open source and freely available, but the Databricks Runtime ML saves you from the hassle of learning how to install and configure them.



Exam Tip

A discussion of how you program ML models is far outside of the scope of the AZ-900 exam and we won't discuss it here. The important point to remember

is that Databricks works with third-party ML frameworks to allow you to build ML models.

To use the Databricks Runtime ML, you'll need to either specify it when you create your cluster, or edit your existing cluster to use it. You do that by choosing one of the ML runtimes as shown in [Figure 2-51](#).

Clusters / jcCluster

0
?

jcCluster

Cancel
Confirm and Resize

2-8 Workers: 28.0-112.0 GB Memory, 8-32 Cores, 1.5-6 DBU
1 Driver: 14.0 GB Memory, 4 Cores, 0.75 DBU Cost: \$0.40 per DBU

Switch to old style

jcCluster

Cluster Mode

Standard

Databricks Runtime Version

Runtime: 5.2 (Scala 2.11, Spark 2.4.0)

Databricks Runtime:

5.2	Scala 2.11, Spark 2.4.0
5.2	GPU, Scala 2.11, Spark 2.4.0
5.2 ML Beta	GPU, Scala 2.11, Spark 2.4.0
5.2 ML Beta	Scala 2.11, Spark 2.4.0
5.1	Scala 2.11, Spark 2.4.0
5.1	GPU, Scala 2.11, Spark 2.4.0
5.1 ML Beta	GPU, Scala 2.11, Spark 2.4.0
5.1 ML Beta	Scala 2.11, Spark 2.4.0
5.0	Scala 2.11, Spark 2.4.0
5.0 ML Beta	Scala 2.11, Spark 2.4.0
5.0	GPU, Scala 2.11, Spark 2.4.0
5.0 ML Beta	GPU, Scala 2.11, Spark 2.4.0
4.3	Scala 2.11, Spark 2.3.1
4.3	GPU, Scala 2.11, Spark 2.3.1
4.2	Scala 2.11, Spark 2.3.1
4.2	GPU, Scala 2.11, Spark 2.3.1
3.5 LTS	Scala 2.11, Spark 2.2.1

Min Workers
Max Workers

2
8

Figure 2-51 Databricks Runtime ML in cluster configuration

You're not limited to the libraries included with

Databricks Runtime ML. You can configure most any third-party ML tools in Azure Databricks, and Microsoft provides some pointers on doing that in their documentation located at:

<https://docs.azuredatabricks.net/spark/latest/mllib/index.html#third-party-libraries>.



Exam Tip

You might have noticed several references to Spark in Databricks. That's because Databricks is based on Apache Spark, an open source system for doing computer work in a clustered environment.

Once you've built your ML model in Databricks, you can export it for use in an external ML system. This process is referred to as *productionalizing* the ML pipeline, and Databricks allows you to productionalize using two different methods: MLeap and Databricks ML Model Export.

MLeap is a system that can execute an ML model and make predictions based on that model. Databricks allows you to export your model into what's called an MLeap bundle. You can then use that bundle in MLeap

to run your model against new data.

Databricks ML Model Export is designed to export your ML models and pipeline so that they can be used in other ML platforms. It's specifically designed to export Apache Spark-based ML models and pipelines.

Azure Machine Learning Service


The Azure Machine Learning Service provides a cloud-based solution for building ML models. The Machine Learning Service uses a programming language called Python, so you'll need to be familiar with Python to use the service.

The main purpose of the Azure Machine Learning Service is to use cloud-based resources to run the complex computations necessary to build ML models. Unlike Databricks where everything is in the cloud, with Machine Learning Service, you can build your data sets on-premises and then upload your data to the cloud to do ML modeling.

Like Databricks, Machine Learning Service uses notebooks. You can use Jupyter Notebooks on-premises, but you can also use Azure Notebooks, a cloud-based Jupyter Notebook offering from Microsoft. Whether you use a local notebook or Azure Notebooks, you'll typically start things off by training your model locally in order to save on compute costs. Once you're

ready to train your model in Machine Learning Services, you can move the data to the cloud, create a cloud-based script for your model, and start training your model, all within your notebook.

Figure 2-52 shows output of training a ML model tested on a local machine. This model looks at images of handwritten numbers and attempts to correctly identify the numbers that were written. It took three minutes to train on a local machine, providing us with a 92% accuracy level.



```
In [8]: %time
from sklearn.linear_model import LogisticRegression

clf = LogisticRegression()
clf.fit(X_train, y_train)

CPU times: user 2min 42s, sys: 4.45 s, total: 2min 46s
Wall time: 2min 46s

In [9]: y_hat = clf.predict(X_test)
print(np.average(y_hat == y_test))


0.9201
```

Figure 2-52 Training a simple ML model and prediction run locally using Azure Notebook

When you submit your model to your Machine Learning Service cluster for training, it will prepare the model and then queue it for training in the cluster. Like Databricks, the Machine Learning Service is a serverless service, meaning you only use compute resources when

your using the cluster. When you submit a job, it's queued until compute resources are available, usually taking only a few seconds. Once your job completes, those resources are released.

In [Figure 2-53](#), we've sent the model to a cluster in Machine Learning Services running an experiment on it to test for accuracy. If you click the link to the Azure portal, you can see additional information about the run.



```
In [14]: run = exp.submit(config=est)
run
```

Out[14]:

Experiment	Id	Type	Status	Details Page	Docs Page
sklearn-mnist	sklearn-mnist_1550362047_20a6d675	azureml.scriptrun	Starting	Link to Azure Portal	Link to Documentation

Figure 2-53 Running a script to train a model in the cloud

In [Figure 2-54](#), you can see the node in the cluster where this script is running. In this test, we only have one node in the cluster, but you can add additional compute resources if needed. If you were training a complex model, you might want to add compute resources in order to more quickly train that model.




Figure 2-54 Machine Learning Services node running an experiment

When you train models in Machine Learning Services, a Docker container is created, and your model actually runs inside of that container. A Docker container is a zipped copy of everything that's necessary

to run your model. That zipped copy is called a Docker image, and it can be run on any computer that is running the Docker runtime, including the VM that makes up your Machine Learning Services cluster.

When you want to export your model so that you can use it in a production workload, you can export it as a Docker image. By using Docker images, Machine Learning Services is able to make your model portable so that it can run just about anywhere. In addition to that, you can use powerful container clustering services like Azure Kubernetes Service to run your models at large scale.

More Info Docker Containers

A discussion of Docker containers is outside of the scope of this guide, but if you're interested in learning more about Docker, see: <https://www.docker.com>.

Machine Learning Services can also export your model as an FPGA image. FPGA stands for field-programmable gate array, and it's similar to a microprocessor except that it can be programmed by a user after manufacturing. FPGAs are extremely fast because they can be programmed explicitly for the task at hand. The only thing faster for AI processing is the application-specific integrated circuit, or ASIC, but an ASIC must be manufactured for its end purpose. It

cannot be reprogrammed later.

Microsoft has invested heavily in an FPGA infrastructure for AI, and FPGAs are available today in every Azure data center. In fact, Microsoft powers its own cognitive services for Bing search and more using FPGAs.

Azure Machine Learning Studio

The AZ-900 exam is not a technical exam, and it's pretty tough to tackle the concept of ML and AI without getting technical. Up until this point, we've tried to keep things at a high level and not get too technical, and because of that, some of the concepts might be a little hard to grasp. Thankfully, there's a way to deal with ML concepts in a visual way. Azure Machine Learning Studio allows people who aren't data scientists to delve into ML and gain a better understanding of the concepts we've discussed up to this point.

Machine Learning Studio is SaaS for ML. It provides an easy-to-use drag and drop interface for creating, testing, and deploying ML models. Instead of having to write your own models, Machine Learning Studio includes a large collection of pre-written models that you can apply to data. The best way to get a handle on Machine Learning Studio is a hands-on approach, so let's use Machine Learning Studio to build an ML model and test it.

To launch Machine Learning Studio, open a web browser and browse to <https://studio.azureml.net>. Click **Sign-In** in the upper right corner and sign in with your Azure subscription username and password.

Once Machine Learning Studio opens, you'll be taken to your default workspace. A workspace is a logical container for your experiments, datasets, models, and so on. Machine Learning Studio assigns your workspace a default name, but you can change it by clicking on **Settings** in the lower left corner as shown in [Figure 2-55](#).




Figure 2-55 Changing setting of our Machine Learning Studio workspace

Notice in [Figure 2-55](#) that the workspace type is Free. There are two tiers in Machine Learning Studio: Free and Standard. The Free tier is for experimentation while the Standard tier is what you'd want to use if you are using your ML model in a production scenario. There are additional capabilities in the Standard tier, and you have to pay for workspaces that use the Standard tier.

More Info Pricing Of Machine Learning Studio Tiers

For more information on the features and pricing of Machine Learning Studio tiers, click on the [Learn More](#) link next to **Workspace Type** as shown in [Figure 2-55](#). This will take you to the pricing page for Machine Learning Studio.

When you create a workspace by browsing directly to Machine Learning Studio, it will always be in the Free tier. If you want to create a workspace in the Standard tier, you will need to use the Azure portal to create a Machine Learning Studio Workspace. You can then choose your tier.

The Free tier is fine for our purposes because we're just running some tests. Change the workspace name to "AZ-900-Workspace" and add a useful description. You can then click **Save** at the bottom of the screen to save your new name and description.

We're now ready to create our ML model, but before we do, let's review exactly what we're going to do. We

will:

- Create an experiment so that we can test and train the ML model.
- Add data we will use to train the ML model.
- Add a pre-existing ML algorithm from Machine Learning Studio.
- Configure Machine Learning Studio to train the model based on the dataset.
- Run an experiment to see how reliable the ML algorithm is.

For this experiment, we're going to use data that's included with Machine Learning Studio. The data shows arrival and departure on-time data for various airlines over a one-year period. We will use this data to build a model that will predict the likelihood of a particular flight arriving on-time at its destination.

Step 1: Create an Experiment

The first step to building a ML model for our flight prediction is to create a new experiment. This is where we will create and test the ML model, and it's called an experiment for a reason. After we test the model, we'll change some things to try and increase the reliability of the model.

To create an experiment in Machine Learning Studio, click on **Experiments** on the menu on the left, and then click the **New** button at the bottom of the screen, as shown in [Figure 2-56](#).



Figure 2-56 Creating a new experiment in Machine Learning Studio

When you do this, you'll see a collection of templates that Microsoft provides for experiments. These are all pre-built ML experiments, and you can learn a lot by choosing one of them to see how they work, but for our purposes, we're going to start with a blank experiment. You then click on **Blank Experiment** as shown in [Figure 2-57](#).



Figure 2-57 Creating a blank experiment

Once you've created your experiment, you'll see the screen shown in [Figure 2-58](#). On the left side, you'll see a list of all the items you can add to your experiment. You'll see a list of sample data, but if you scroll down, you'll see all kinds of items you can use to build a model.

The main part of the screen is where you'll build your model, and you'll do that by dragging items from the list on the left and dropping them onto the main screen. You'll then connect items together to build your model.

Before we do that, let's rename this experiment so we'll be able to easily identify it. Your experiment name appears at the top of the screen. Click on that and enter a new name for your experiment as shown in [Figure 2-58](#).




Figure 2-58 A blank experiment in Machine Learning Studio

Step 2: Add Data

In order to train a ML algorithm, you need to feed data into it. ML uses historical data to learn how to predict a particular result in the future, and the more data you use to train your model, the more reliable your model will be.

Machine Learning Studio makes it easy to import data from Azure Blob Storage, Azure SQL Database, Hive queries, Azure Cosmos DB, and more. For our ML model, however, we're going to use some sample data that's included with Machine Learning Studio.

To find the data we want to use, enter "flight" in the search box on the left. When you do, you'll see "Flight Delays Data." This is the data we want to use, so click on it and drag it to the main screen on the right as shown in [Figure 2-59](#).




Figure 2-59 Adding data to our experiment

Before you do any work building a ML model, you need to have a good understanding of the data you will use to train that model. Only by having a good understanding of your data will you be able to build a reliable model, and Machine Learning Studio makes it easy to learn about your data. If you right-click on the **Flight Delays Data** item you just dragged onto your experiment, you can click on **Dataset** and then **Visualize**, as shown in Figure 2-60, to see the data contained in the dataset.




Figure 2-60 Using Machine Learning Studio to visualize data

Once your dataset opens in Machine Learning Studio, you will see that we have a little over 2.7 million rows of data to work with. If you click on the Month column header, you'll see that we have 7 unique values for the month as shown in [Figure 2-61](#). That means we have data here for 7 months out of the year, and while not perfect, that will suffice for what we're doing. If you were going to use our model in a real scenario, you'd likely want more data for additional months.




Figure 2-61 Visualizing a dataset in Machine Learning Studio

One important thing to pay attention to with ML modeling data is the Missing Values field shown in Figure 2-61. A missing value means that data is either missing completely or you have a 0 in a numeric field. If you are missing data, your model will be flawed, so

you'll want to try and ensure you're not missing important data.

In this particular dataset, we have some columns that have missing values. Machine Learning Studio includes items you can add to your model to account for missing values. Click the X in the upper-right corner of your dataset to close it. Enter “missing” in the search box and you'll see an item under Data Transformation called Clean Missing Data. Drag that into your model as shown in Figure 2-62.




Figure 2-62 The Clean Missing Data item helps to account for missing values

In order for Clean Missing Data to do anything to the dataset, we need to feed the dataset into it. To do that, click on **Flight Delays Data** and you'll see a circle with a "1" in it. That circle represents the output of the dataset, and we need to connect that to the input of the Clean Missing Data transformation. Click and hold on the "1" and drag it down to the small circle at the top of Clean Missing Data. You'll see the small circle at the top of Clean Missing Data turn green as shown in [Figure 2-63](#). Once the two nodes are connected, release your mouse button.




Figure 2-63 Connecting nodes in Machine Learning Studio

Now we need to tell Machine Learning Studio what we want to do about missing values. Click on Clean Missing Data and you'll see some properties you can set in the pane on the right. The first thing we need to do is select the columns that have missing values.

We can discover those columns by visualizing the dataset as shown in [Figure 2-61](#). Based on that, we know

that the DepDelay, DepDel15, and ArrDelay columns all have missing values. However, in the model, we aren't concerned with the DepDelay or ArrDelay columns. These columns contain the number of minutes a flight was delayed in departure or arrival, but the columns we want to use are DepDel15 and ArrDel15. Those columns will contain a 0 if a flight departed or arrived within 15 minutes of the schedule and a 1 if it didn't. Therefore, we only need to clean the DepDel15 column of missing values.

- With Clean Missing Data selected, click on Launch Column Selector in the Properties pane as shown in [Figure 2-64](#).




Figure 2-64 Properties of Clean Missing Data

- In the Select Columns screen, click on **No Columns** under Begin With.
- Change **Column Indices** to **Column Names** and enter DepDel15 in the input box. When you do, a list of columns will appear.
- Click on DepDel15 to select that column. Your Select Columns screen should now look like the one shown in [Figure 2-65](#).




Figure 2-65 Selecting columns to clean

- Once you've entered the DepDel15 column, click the check mark button in the lower right to save those settings.

We're going to completely remove any rows that have missing values in the DepDel15 column.

- Click on the **Cleaning Mode** dropdown (shown in [Figure 2-64](#)) and select **Remove Entire Row**. We can now perform our data transformation and check the results.
- Click the **Run** button at the bottom of Machine Learning Studio. This will queue our transformation job and you'll see a small

clock appear on the Clean Missing Data node.

When the transformation process starts, you'll see a green spinning circle. The transformation is going to clean almost 3 million rows, so it will take a minute or so. When it's complete, a green check will appear.

We can now look at our cleaned dataset to check the results of our data cleaning.

- Right-click on **Clean Missing Data**.
- Point to **Cleaned Dataset**.
- Click on **Visualize**, as shown in [Figure 2-66](#).




Figure 2-66 Visualizing a cleaned dataset

Use the same technique you used earlier when visualizing the dataset. You'll notice that we have fewer records than we had before because our data

transformation removed rows where data was missing. If you click on the DepDel15 column, you'll see that it now has no missing values.

Now that we have clean data, we need to tell our model which columns in this data are of interest to our model. We want to create a model that will predict the likelihood that a particular flight will arrive late, so we only want columns in the dataset that are important for us to look at to predict that. Anything else is just noise that the model doesn't need to consider.

To tell the model which columns to look at, we'll use the Select Columns in Dataset item.

- Click in the search box and enter **Select**.
- Locate the Select Columns in Dataset transformation. Drag it from the list, and drop it directly under Clean Missing Data.
- To connect the cleaned dataset to Select Columns in Dataset, click on **Clean Missing Data** and drag the “1” node to the top node of Select Columns in Dataset as shown in [Figure 2-67](#).




Figure 2-67 Connecting a clean dataset to Select Columns in Dataset

Notice that Select Columns in Dataset is displaying a red circle with an exclamation point. That tells us that we need to configure something for it to work. In this case, we need to tell it which columns we want to select.

- Click on **Select Columns in Dataset** and click **Launch Column Selector** from the Properties pane.
- Select **Year** from the list on the left and click the right arrow button to move it to the list of selected columns.
- Do the same thing for all columns except DepDelay, ArrDelay, and Cancelled. You should see a screen like the one shown in [Figure 2-68](#).




Figure 2-68 Selecting relevant columns for the model

The columns that we selected are all columns that contain information that might impact the arrival time of any particular flight. We're now ready to start training the ML model with the new dataset.

Note Spend Time With Your Data

We've spent a lot of time looking at the data and working to clean it. This step is extremely important when it comes to ML. You need to not only fully understand your data, but you also need to ensure your data is as clean as possible and that you're only sending relevant data to your model.

One of the great things about Machine Learning Studio is that you can always go back and redo things easily if you make mistakes.

Step 3: Train the Model

When you train a machine language model, you don't give it all of your data. Instead, you split your data and send a percentage of the data to the model for training. Once the model is trained, you use the remaining data to test your model. The process of testing a model is called *scoring* the model. By using data with known values to score the model, you can see how many times your trained model got the prediction right before you throw real data with unknown results at it.

Machine Learning Studio makes it easy to split your data for training.

- Enter **split** in the search box and locate the Split Data item.
- Drag it under the Select Columns from Dataset node.
- Connect the Select Columns in Dataset output node (the circle on the bottom) to the top node of Split Data as shown in [Figure 2-69](#).




Figure 2-69 Split Data to train your ML model

As a general rule of thumb, it's a good idea to send about 80% of your data to your model for training and 20% of your data to score the trained model. To do that, click on **Split Data** and set the fraction of rows to include in the first output dataset to .8 as shown in [Figure 2-70](#).




Figure 2-70 Setting a fraction for our first dataset

To train a model in Machine Learning Studio, use the Train Model item.

- Click in the search box and enter **train model**.
- Drag Train Model to your screen under Split Data.
- To make room for our additional items, click on the button circled in [Figure 2-71](#) to enable pan mode and move your items up.
- Once you've moved them where you want them, click the pan mode button again to turn it off.

- Connect the left node of Split Data to the top right node of the Train Model item as shown in [Figure 2-71](#). This sends 80% of our dataset to the Train Model item.




Figure 2-71 Connecting the first dataset from Split Data to Train Model

There are a couple of things we need to tell Train Model before we can start training. First of all, we need to tell it what we want our model to predict. We want

this model to predict whether a particular flight will arrive at its destination on time, so we want it to predict the value of the ArrDel15 column. (Remember, this value will be 0 if the flight was within 15 minutes of scheduled arrival time and 1 if it's not.)

- Click on **Train Model** and click **Launch Column Selector** in the Properties pane.
- Click on the ArrDel15 column to select it and then click the check button.

Next, we need to tell Train Model what kind of ML model we want it to use. If you were writing this ML model yourself, you'd have to write the model using a programming language like Python or R, but Machine Learning Studio contains a large number of ML models you can use without any programming.

More Info ML Models

There are specific ML algorithms that are well-suited to particular situations. A great way to determine the best algorithm in Machine Learning Studio is to use the Machine Learning Cheat Sheet. You can find it at: <https://aka.ms/mlcheatsheet>.

The model is going to predict a Boolean value based on input data, and for this kind of model, the Two-Class Boosted Decision Tree algorithm is ideal.

- Click in the search box and type **two-class boosted**.

- Drag the Two-Class Boosted Decision Tree item to your screen and drop it to the left of Train Model.
- Connect the node on Two-Class Boosted Decision Tree to the left node of Train Model as shown in Figure 2-72.




Figure 2-72 Adding an algorithm to train our model

We now have everything in place to actually train the ML model, but we also want to score the model to see how accurate it is. Before we actually train it, let's configure the workspace so that we can see how well the model did after training.

Step 4: Score the Model

In order to tell how successful this model is, we need to be able to score it with the remaining 20% of the data. We'll use the Score Model item in Machine Learning Studio to do that.

- Click in the search box and type **score model**.
- Drag the Score Model item to your workspace and drop it under Train Model.
- Connect the output node of Train Model to the top left node of Score Model. This sends the trained model to Score Model.
- In order to send the remaining 20% of the data to Score Model for scoring, connect the right output node of **Split Data** to the top right node of Score Model as shown in [Figure 2-73](#).




Figure 2-73 Scoring our model with Score Model

You have just completed all of the steps necessary to create an ML model, send millions of rows of data to it for training, and then test your trained model to see how well it succeeds. The only thing left is to run the training and scoring and see the result. We still need to add a way for us to evaluate that result.

- Click in the search box and type **evaluate**.
- Drag the **Evaluate Model** item so that it's under **Score Model**.

- Connect the bottom node of **Score Model** to the top left node of **Evaluate Model** as shown in [Figure 2-74](#).




Figure 2-74 Evaluate a model with Evaluate Model

The Evaluate Model item will allow you to see an easily-readable report showing us the accuracy of the model.

Step 5: Train and Score the Model

To train the model and then score it with the remaining

20% of the data, click the **Run** button at the bottom of the workspace. It will take a while to run, so be patient.

When the model has been trained and scored, you'll see a green check on each item in your workspace. To view the results of the experiment, right-click on **Evaluate Model**, point to **Evaluation Results**, and click on **Visualize**, as shown in [Figure 2-75](#).




Figure 2-75 Visualizing evaluation results

After clicking on Visualize, Machine Learning Studio

will show you an ROC curve as shown in [Figure 2-76](#). ROC stands for receiver operating characteristic, and it's a typical graph for determining the effectiveness of an ML model. The further the line is to the left of the graph, the better the ML model. In this case, we see the model did well. Scroll down to see details on how well we did.




Figure 2-76 An ROC curve showing the results of scoring

In [Figure 2-77](#), you can see the full results of the experiment. You can see that we scored a 91.3% accuracy rate. We can improve this rate by feeding more data to the model, or possibly by using a different algorithm. Feel free to experiment on your own, and if you want to dig into a much more complicated version of this experiment, check out the [Binary Classification: Flight Delay Prediction](#) experiment template in Machine Learning Studio.




Figure 2-77 Accuracy of the ML model

Once you're satisfied with the result of the ML model, you can publish it to the Internet so that users can call it. Machine Learning Studio uses a Web service (which is a website) to do this, and you can easily deploy your ML model to a Web service by clicking the **Set Up Web**

Service button at the bottom. When you do, Machine Learning Studio will add some nodes to a new tab called Predictive Experiment to your workspace, so you'll need to click the Run button to test the new configuration before you can use the Web service.

After you click Run and your new predictive experiment complete successfully, click the **Deploy To Web Service** button to finish creating your Web service. Machine Learning Studio will display your Web service as shown in [Figure 2-78](#), and you can even click on **Test** to test your new Web service.



Figure 2-78 A new Web service for the ML model

Serverless computing

As you’ve already learned, one of the great advantages of moving to the cloud is that you can take advantage of the large amounts of infrastructure that cloud providers have invested in. You can create VMs in the cloud and pay for them only when they’re running. Sometimes you just need to “borrow” a computer in order to run a

computation or perform a quick task. In those situations, a serverless environment is ideal. In a serverless situation, you pay only when your code is running on a VM. When your code's not running, you don't pay anything.

The concept of serverless computing came about because cloud providers had unused VMs in their data centers and they wanted to monetize them. All cloud providers need surplus capacity so they can meet the needs of customers, but when VMs are sitting there waiting for a customer who might want to use it, it's lost revenue for the cloud provider. To solve that problem, cloud providers created consumption-based plans that allow you to run your code on these surplus VMs and you pay only for your use while your code is running.



Exam Tip

It's important to understand that “serverless” doesn't mean that no VMs are involved. It simply means that the VM that's running your code isn't explicitly allocated to you. Your code is moved to the VM, it's executed, and then it's moved off.

Because your serverless code is running on surplus capacity, cloud providers usually offer steep discounts on consumption-based plans. In fact, for small workloads, you may not pay anything at all.

Azure has many serverless services. We've already discussed that Azure Databricks and Azure Machine Learning Service are serverless. However, there are other serverless services that don't fit into the categories we've already discussed. They are Azure Functions for serverless compute, Azure Logic Apps for serverless workflows, and Azure Event Grid for serverless event routing.

Azure Functions

Azure Functions is the compute component of Azure's serverless offerings. That means that you can use Functions to write code without having to worry about deploying that code or creating VMs to run your code. Apps that use Azure Functions are often referred to as Function Apps.

More Info Function Apps Use App Service

Function Apps are serverless, but under the hood, they run on Azure App Service. In fact, you can choose to create your Function App in an App Service plan, in which case you don't benefit from the consumption model of paying only when your code runs. We'll cover that in more detail later in this chapter.

Functions can be created in many different ways. You can create a Function App using:

- Microsoft Visual Studio
- Microsoft Visual Studio Code
- Maven for Java Function Apps
- Python command line for Python Function Apps
- Azure command line interface (CLI) on Windows or Linux
- The Azure portal

Assuming you aren't creating your Function App using a method specific for a particular language, you can choose between .NET (for C# and F# Function Apps), Java, and JavaScript (for Node Function Apps.) In [Figure 2-79](#), we're creating a Function App in the Azure portal, and selected .NET as the Function App runtime so that you can use the C# language to write functions.




Figure 2-79 Creating a new Function App in the Azure portal

Once your Function App is ready, you can open it in the portal to begin creating functions. Figure 2-80 shows the new Function App in the Azure portal.




Figure 2-80 A new Function App in the Azure portal

From here, you can create a new function, a new proxy, or a new slot. A function is code that runs when something triggers it. (We'll look at triggers soon.) A proxy allows you to configure multiple endpoints for your Function App, but expose all of them via a single URL. Slots allow you to create a copy of your Function App that isn't public-facing. You can write code and test this copy, and when you are satisfied that it's ready for production, you can swap it over to production with the click of a button. This feature in App Service is called Deployment Slots.

If you click on **Function App Settings** under Configured Features (shown in [Figure 2-80](#)), we can change some settings for the Function App, as shown in [Figure 2-81](#).

OverviewPlatform featuresFunction app settings x

Daily Usage Quota (GB-Sec) ⓘ

Enter value in GB-sec

Set quota

Application settings

[Manage application settings](#)

Runtime version

Runtime version: 2.0.12427.0 (~2)

~1

~2

Function app edit mode

Change the edit mode of your function app

Read/Write

Read Only

Host Keys (All functions)

NAME	VALUE	ACTIONS
_master	Click to show	Copy Renew
default	Click to show	Copy Renew Revoke

Add new host key

host.json

1

```
{ }
```

Figure 2-81 Function App settings

From this screen, you can configure a daily quota for your Function App. Once you reach the quota, Azure will stop the Function App until the next day. You can

also change the Function App runtime version. This is the runtime version of Azure Functions, and while it's generally advised to use the latest version, if your functions were written in an earlier version, you won't be able to upgrade them by simply changing the version here. Changing major versions can cause your app to break, so Microsoft will prevent you from changing the version if you have existing functions in your Function App.

You can also change your Function App to read-only mode to prevent any changes to it. This is helpful if you have multiple developers writing code for your app and you don't want someone changing something without your knowledge. Finally, you can view, renew, revoke, and add new host keys. A host key is used to control access to your functions. When you create a function, you can specify whether anyone can use it or whether a key is required.



Exam Tip

Although a key can help protect your functions, they're not designed to offer complete security of Function Apps. If you want to protect your Function App from unauthorized use, you should use authentication features available in

**App Service to require authentication.
You can also use Microsoft API
Management to add security
requirements to your Function App.**

If you click on Application Settings (shown in [Figure 2-80](#)), you can configure the settings for the Function App. These are settings specific to App Service. [Figure 2-82](#) shows some of these settings, including whether the app runs in 32-bit or 64-bit, the HTTP version, how you can access your files using FTP, and more. You can also configure database connection strings from this page.

Save

Discard

General settings

PHP version ⓘ

5.6

Platform ⓘ

32-bit

64-bit

HTTP Version

1.1

2.0

ⓘ

Auto swap destinations cannot be configured from production slot

Auto Swap

Off

On

Auto Swap Slot

ⓘ

FTP based deployment can be disabled or configured to accept FTP (plain text) or FTPS (secure) connections. Click to learn more.

FTP access

FTP + FTPS

FTPS Only

Disable

Debugging

Remote debugging

Off

On

Remote Visual Studio version

2015

2017

Figure 2-82 Some of the Function App settings

Finally, if you click on the Platform Features tab, you can see all of the features available to you in the App Service platform, as shown in [Figure 2-83](#). From here, you can configure things such as SSL certificates,

custom domain names for your Function App, turn-key authentication, and more.




Figure 2-83 App Service platform features available to your Function App

More Info Azure App Service

A full discussion of Azure App Service is outside of the scope of this book, but if you want to learn more, check out:
<https://azure.microsoft.com/services/app-service>.

To create a new function, click on the + sign as shown in [Figure 2-84](#). You can then choose your development environment. You can choose Visual Studio, Visual Studio Code, a development environment right inside the Azure portal, or you can use a code editor of your choice alongside the Azure Functions Core Tools.

contosopharm

Function Apps

Search

Jim's Personal Azure Account

Function Apps

▼ contosopharm

► Functions

► Proxies

► Slots (preview)

Overview

Platform features

Quickstart



Azure Functions for .NET - getting started

Follow our Quickstart guidance to author and publish a function [Learn more](#)

1

CHOOSE A DEVELOPMENT
ENVIRONMENT

2

CREATE A FUNCTION



Visual Studio

Use Visual Studio to author, build, and run .NET functions



VS Code

Use Visual Studio Code to author your functions



Any editor + Core Tools

Write functions using your favorite editor and the Azure Functions Core Tools



In-portal

Author functions quickly in the portal

Figure 2-84 Creating a function

If you choose any option other than In-Portal, you'll need to specify how you want to deploy your function to App Service. Your options depend on which development environment you choose, but typically it will involve either using features of your environment to send the function directly to App Service, or you'll need to use App Service Deployment Center. Either way, deployment is quick and easy.

Depending on which development environment you choose, you will likely have to complete some prerequisite steps in order to develop your function. You'll see a screen telling you exactly what to do so that everything will work correctly. In [Figure 2-85](#), you can see what's required to use VS Code to develop functions. In most cases, it will require you to install the Azure Functions Core Tools.




Figure 2-85 Creating a function using Visual Studio Code and Azure Functions Core Tools

Functions work using a trigger-based system. When you create your function, you choose a trigger that will kick off your function. When it's triggered, your

function code will run. You will typically want your function code to do something simple. If you need a more complex function that performs many things, you can use Function Proxies to create several functions that work together to complete a task. This kind of development is referred to as *microservices*, and it allows you to quickly swap out functionality by simply changing a single function.

After your function is triggered and the code runs, you can choose what happens using what's called an *output binding*. The type of bindings you can use are dependent on the type of function you create. [Figure 2-86](#) shows some of the different output bindings available when using an HttpTrigger for a function. This function will run as soon as a particular URL is requested.




Figure 2-86 Output bindings in Azure Functions

More Info Httptrigger Functions

HttpTrigger functions are incredibly powerful because they can be called as a webhook. Many online services support webhooks. In a webhook scenario, you can configure a service to make a request to a

particular URL in response to events. If you configure that webhook to call your Azure function's URL, you can easily add powerful functionality to your workflow.

You can configure multiple outputs for your function as well. For more complex workflows, however, Logic Apps is often a better choice, and you can integrate Logic Apps directly with Azure Functions.

Azure Logic Apps

Logic Apps are similar to Function Apps in that they are kicked off by a trigger, but what happens after that is completely different. Unlike Function Apps, you don't have to write code to create some powerful workflows with Logic Apps.

A workflow simply means that a Logic App reacts to something happening and responds by performing a series of tasks such as sending an email, transferring data to a database, and so on. It can do these things in order, but it can also do two things at once. As an example, you might have an e-commerce site and when a customer orders a product you might want to:

- Update your inventory count of the product
- Generate an invoice for the item
- Email the invoice to the customer
- Sign the customer up for your newsletter
- Generate a shipping label for the item

Logic Apps allows you to create these kinds of complex workflows easily, and because Logic Apps integrates with over a hundred other services (both Azure services and third-party services), you can do just about anything in a Logic Apps workflow.

There are three components in Logic Apps that make workflows possible: connectors, triggers, and actions. A connector is a component that connects your Logic App to something. That could be another Azure service, but it could also be a third-party service, an FTP server, and so forth. Each connector will have one or more triggers and actions specific to that connector. A trigger is a specific action that will cause your Logic App workflow to run, and an action is what your Logic App will do as an output. You can combine multiple actions for a connector, and you can also combine multiple connectors to create complex and powerful workflows.

You create Logic Apps in the Azure portal. Once you create it, the Logic Apps designer is shown by default. From the designer, you can choose the trigger for your Logic App as shown in [Figure 2-87](#). The list shown is a brief list of common triggers, but there are many more to choose from. In fact, there's a trigger for Azure Functions as well, so you can trigger a Logic Apps workflow when your function runs.



Figure 2-87 Common Logic App triggers



Exam Tip

It's important to understand the difference between connectors and triggers. All of the items shown in Figure 2-87 are triggers that are associated with specific connectors. For example, When A New File Is

Created On OneDrive is a trigger for the OneDrive connector. There are other OneDrive triggers available as well, including When A File Is Modified and When A File Is Deleted.

If you scroll down, you'll see a large number of templates you can use to create a Logic App as shown in [Figure 2-88](#). These templates will automatically configure a Logic App that contains a full workflow that you can modify for your own purposes. This is the fastest way to get started, but the included templates might not be exactly what you want, so you can also create a blank Logic App and start from scratch.




Figure 2-88 Logic App templates

After you create your blank Logic App, you can choose from several ways to start building your workflow. You can select a trigger from the list, search for a trigger or connector, or you can just select a connector from the list and see what triggers are available. As shown in [Figure 2-89](#), there are many options available to get

started.




Figure 2-89 Adding triggers to your Logic App

In [Figure 2-90](#), we've configured the OneDrive connector to monitor a folder in OneDrive. When a file is modified in that folder, it will start the Logic App

workflow. In order to do something when a file is modified, click on **New Step** to add an action.




Figure 2-90 Using the OneDrive connector

When you click on **New Step**, you'll see the same kind of screen that shows when the Logic App starts. Since we added a step to a workflow that already has a trigger, Logic Apps shows the actions you can take when the app is triggered. There are many actions to choose from, as shown in [Figure 2-91](#).




Figure 2-91 Adding an action to the Logic App

In [Figure 2-92](#), we configured the Logic App to call the Function App when a file is modified in the OneDrive folder. You can pass the filename that was modified to the Function App so that it will know what has changed, which you can do using dynamic content. Just click on **File Name** from the list. Of course, you

can only pass one dynamic content item in your action.




Figure 2-92 Configuring a Function App action

More Info Passing Parameters To Function Apps

When using a Logic App to call a Function App, make sure the Function

App was designed to accept the data the Logic App is passing to it. Otherwise, the Function App will encounter an error when it's triggered by the Logic App.



Exam Tip

As you're configuring triggers and actions in the Logic Apps designer, Logic Apps is writing code for you under the hood that will implement your workflow. Logic App workflows are defined using JSON files, and the designer generates this JSON code as you are configuring your app.

You now have a functioning Logic App. You can test the workflow by clicking **Save** at the top of the designer. The OneDrive connector was configured to check for a modified file every three minutes (see [Figure 2-90](#)), so you may need to wait a few minutes before the workflow is triggered. You can also click on **Run Trigger** at the top of the designer to manually run the trigger.

You can monitor your Logic Apps using the Azure portal. Open the app and click on **Overview** to see when your trigger was activated, and whether or not it ran your workflow as shown in [Figure 2-93](#).




Figure 2-93 The Azure portal displaying when my Logic App flow ran

If you click on **See Trigger History**, you can see an entire history of when your trigger was evaluated and when it fired the workflow for your Logic App.

In this case, we've used a Logic App to call an Azure Function, but you could have written a log file to Azure

Storage or stored some information in an Azure SQL Database. If you want your Logic App to integrate specifically with other Azure services such as this, you can integrate your Logic App with Azure Event Grid for a more optimal experience.

Azure Event Grid

The concept of different Azure services interacting with each other should be pretty familiar to you by now. There are many ways that you can integrate services such as this, and in some cases, you need one Azure resource to know about a change in another Azure resource. You could use a polling method for this, similar to the Logic App checking once against OneDrive every three minutes looking for a change. It's more efficient, however, to enable an Azure service to trigger an event when something specific happens, and configure another Azure service to listen for that event so it can react to it. Event Grid provides that functionality.

Note Event Grid And Serverless Computing

Event Grid has many capabilities that aren't related to serverless computing, but in the scope of this chapter, we only cover serverless capabilities and Event Grid.

Both Azure Functions and Azure Logic Apps are integrated with Event Grid. You can configure a

function to run when an Event Grid event occurs. In Figure 2-94, you can see the list of Azure resources that you can trigger Event Grid events. Not all Azure services are represented in Event Grid, but more services are being added over time.




Figure 2-94 Resources available in Event Grid

Once you've selected the resource type, configure the event you want to listen for. The events that are available may differ depending on the resource you selected. In [Figure 2-95](#), we are creating an event for an Azure subscription.

The screenshot shows the Azure Event Grid console interface for configuring a new event subscription. The title bar reads "When a resource event occurs (Preview)". The configuration is as follows:

- * Subscription:** A dropdown menu showing "Jim's Personal Azure Account".
- * Resource Type:** A dropdown menu showing "Microsoft.Resources.Subscriptions".
- * Resource Name:** A dropdown menu showing "Jim's Personal Azure Account".
- Event Type Item - 1:** A dropdown menu with a list of available event types. The list includes:
 - Microsoft.Resources.ResourceActionCancel
 - Microsoft.Resources.ResourceActionFailure
 - Microsoft.Resources.ResourceActionSuccess
 - Microsoft.Resources.ResourceDeleteCancel
 - Microsoft.Resources.ResourceDeleteFailure
 - Microsoft.Resources.ResourceDeleteSuccess
 - Microsoft.Resources.ResourceWriteCancel
 - Microsoft.Resources.ResourceWriteFailure
 - Microsoft.Resources.ResourceWriteSuccessA "Enter custom value" link is visible at the bottom of the dropdown.

Figure 2-95 Events for an Azure subscription

More Info Events

For full details on all events and what they mean, see:
<https://docs.microsoft.com/azure/event-grid/event-schema>.

When an event occurs, you can take an action against an Azure resource using the Azure Resource Manager connector in a Logic App. You can also run a script that interacts with the Azure resource to do something like tag a resource, or configure it in a way that is specific to your organization.

The primary benefit of using Event Grid in this way is the rapid development of solutions. You also benefit from Event Grid reliably triggering your events. If an Event Grid event fails to trigger for any reason, Event Grid will continue to retry triggering the event for up to 24 hours. Event Grid is also extremely cost effective. The first 100,000 operations per month are free, and after that point, you pay 60 cents for every million operations.

SKILL 2.4: UNDERSTAND AZURE MANAGEMENT TOOLS

You now have experience using Azure portal and Azure Resource Manager (ARM). Although using the Azure portal is a common way to interact with Azure services, it's sometimes not the most efficient way, especially if you are doing a lot of things at the same time. For those

more complex situations, Microsoft offers PowerShell cmdlets that you can use to interact with Azure resources, and they also offer the Command Line Interface (CLI) for cross-platform users.

More Info Rest API And Azure App

Microsoft also offers a REST API for interacting with Azure, but we won't cover that in this book because it's not covered in the AZ-900 exam.

This section covers:

- The Azure portal
- Azure and PowerShell
- Azure CLI
- Azure Advisor

The Azure portal

The Azure portal that is in use today is the third iteration of the Azure portal, and it came about when Microsoft moved to ARM. Everything that you do in the Azure portal calls ARM on the back-end.



Exam Tip

For the AZ-900 exam, you probably don't need to know that the Azure

portal is just making calls to ARM on the back end, but it doesn't hurt to know it. For the rest of this section, however, we'll cover only the different parts of the portal and how to navigate and customize it. That information is on the AZ-900 exam.

The first time you open the Azure portal, you'll be prompted to take a tour of the portal. If you're completely unfamiliar with the portal, taking a tour will help you to get a feel for how it works. If you choose not to, and change your mind later, you can click the question mark in the top toolbar to access the guided tour at any time.

The default view in the portal is Home, as shown in [Figure 2-96](#). From here, you can see icons for various Azure services, and if you click on one of those icons, it will show you any resources of that type that you've created. The menu on the left side includes these same icons, and more.

 Help + support[See all \(+100\) >](#)

Kubernetes
services

Cognitive Services

[Security Center >](#)[Cloud Shell >](#)

Figure 2-96 The Home screen in the Azure portal

The bottom portion of the screen includes large tiles designed to help you learn more about Azure. If you click on one of the links provided in these tiles, a new tab will open in your browser so you don't lose your place in the portal.

At the top of the screen, you can choose your default view for the portal. You can select between Home and Dashboard. The Dashboard is a fully-customizable screen that we'll look at a little later. Once you've made your choice, click **Save** and the portal will always open in the screen of your choice. However, you can always access the Home screen or the Dashboard by clicking the relevant links in the menu on the left side of the Azure portal.

Along the top colored bar, you'll find a search bar where you can search for Azure services, docs, or your Azure resources. To the right of the search box is a button that will launch Azure Cloud Shell. Cloud Shell is a web-based command shell where you can interact with Azure from the command-line. You can create Azure resources and more. As you're reading through Azure documentation, you may see a Try It button, and those buttons use Cloud Shell to help you test out different services and features.

To the right of the Cloud Shell button is a filter

button that allows you to configure the portal to only show resources in a certain Azure subscription or Azure Active Directory. To the right of that is the Notification button. This is where you'll see notifications from Azure related to your services and subscription. In [Figure 2-96](#), you can see the number 1 on the button. That indicates that you have one unread notification.

To the right of the notifications button is the Settings button. Clicking on that brings up a panel where you can alter portal settings as shown in [Figure 9-97](#).

Figure 2-97 Portal settings

From Settings, you can change your default view, you can alter the color scheme of the portal, you can disable toast notifications, or pop up notifications that Microsoft may display from time to time. Other settings that appear here may change as Microsoft adds new features. For example, in [Figure 2-97](#), you can see how if you choose to you can change the portal to the new browse experience for your resources.

If you click on your name in the upper-right corner (shown in [Figure 2-96](#)), you can log out or switch to other Azure accounts. You can also change the Azure Active Directory to access resources in another directory. This is helpful if your company has a corporate directory and you also have a personal directory.

The menu along the left side of the portal contains a default list of Azure resources. Clicking on one of those will display all resources of that type. If you don't find a service on that list that you'd like to add to the list, click **All Services**, locate the service you want to add to the list, and click on the star to the right of the service to mark it as a favorite, as shown in [Figure 2-98](#).




Figure 2-98 Marking a favorite service

Note Moving Menu Items

You can reorder items on the menu. Click and hold on an item and drag it to a new location in the menu.

In Figure 2-99, we clicked **App Service Plans** on the

menu to see all of the App Service plans. From this list, you can click on a resource to see that resource. You can also click on a column header to sort by that column, assuming you have more than one resource of that type. Click on **Edit Columns** to edit the columns that are displayed here. To create a new resource of this type, click on **Add**. Finally, you can click the three dots on the far right side of the resource to delete the resource.




Figure 2-99 Viewing a list of resources

When you click on a particular resource, it will open that resource in the portal. Along the left side will be a menu that's specific to the type of resource you opened. In the main window, you'll see different items based on the type of resource you're viewing. These window areas

in the portal are often referred to as blades.

In [Figure 2-100](#), you'll see an App Service Web App in the portal. The Overview blade is a blade that's common to most Azure resources, but the information that appears there will differ based upon the resource. In a Web App, you can see the resource group it's in, the status, the region, and more. We also have various tiles related to Web Apps such as the Http 5xx tile and Data In tile. In the upper right of these tiles is a pin button. If you click on that pin, it will add that tile to the portal dashboard.

Home > contosojc

contosojc
App Service

Search (Ctrl+F)

Overview

Activity log

Access control (IAM)

Tags

Security (Preview)

Diagnose and solve problems

Deployment

Quickstart

Deployment credentials

Deployment slots

Deployment Center

Settings

Application settings

Configuration (Preview)

Authentication / Authorization

Application Insights

Identity

Backups

Browse Stop Swap Restart Delete Get publish profile Reset publish profile

Click here to access our Quickstart guide for deploying code to your app →

Resource group (change)
az900_rg

Status
Running

Location
Central US

Subscription (change)
Jim's Personal Azure Account

Subscription ID
2ed0e6a0-aaf7-4f5f-aff6-bf25ffdef188

Tags (change)
Click here to add tags

URL
<https://contosojc.azurewebsites.net>

App Service Plan
ServicePlan6af6380-b214 (Standard: 1 Small)

FTP/deployment username
contosojcjamesche

FTP hostname
<ftp://waws-prod-dm1-109.ft.azurewebsites.windows.net>

FTPS hostname
<https://waws-prod-dm1-109.ft.azurewebsites.windows.net>

Diagnose and solve problems
Our self-service diagnostic and troubleshooting experience helps you identify and resolve issues with your web app.

Application Insights
Application Insights helps you detect and diagnose quality issues in your apps, and helps you understand what your users actually do with it.

App Service Advisor
App Service Advisor provides insights for improving app experience on the App Service platform. Recommendations are sorted by freshness, priority and impact to your app.

Http 5xx

Data In

Data Out

Figure 2-100 Viewing a Web App in the portal

Along the top of the blade for the Web App are several buttons for interacting with the resource. For a Web App, you have a Browse button that will open the app in a browser, a Stop button to stop the Web App, a Swap button to swap deployment slots, and so on. Each resource type will have different buttons available to you so you can easily interact with the resource from the Overview blade.

If you click on an item in the menu at the left, the content from the Overview blade is replaced with the selected new item. In [Figure 2-101](#), we have clicked on **Diagnose And Solve Problems**, which replaces the Overview blade with new content from the Diagnose And Solve Problems blade.




Figure 2-101 A new blade

As you use the portal, you'll find that there is

inconsistency between different services. Each team at Microsoft has their own portal development sub-team, and they tend to design portal interfaces that make sense for their own team. For that reason, you may see buttons on the top in some blades and buttons on the bottom in other blades.

You can customize your portal experience using the dashboard. If you click on Dashboard from the portal home screen, you'll see your default dashboard. As you're managing your resources, click on pins (as shown in [Figure 2-100](#)) to pin tiles to your dashboard. You can then move these tiles around and customize them in other ways to create a view that's unique to your needs.

To customize your dashboard, click **Dashboard** in the menu to show the dashboard and then click on **Edit** as shown in [Figure 2-102](#).




Figure 2-102 Editing a dashboard

From the customize screen shown in [Figure 2-103](#), you can change the name of your dashboard by clicking inside the current name and changing it to a new name. You can add tiles to the dashboard by choosing from one of the hundreds of tiles available in the Tile Gallery on the left side of the portal, and you can search and filter the list if necessary. If you hover over an existing tile, you'll see a Delete button and a menu button represented by three dots. Click on the **Delete** button to remove the tile from the dashboard. Click the menu button to access a context menu where you can resize the tile.

Add, pin, move, and resize your tiles. Done customizing

« Jim's Dashboard

Tile Gallery

Search tiles

All categories All resource types

207 tiles You can drag any tile to the dashboard

All resources Add

Resource groups Add

Clock Add

Service Health Add

Markdown Add

User Activity Summary Add

Users and groups Add

CPU percentage Virtual machines (classic) Add

Activity Log Analytics Overview Add

Azure portal Add

ARM Data Add

Service Health

Personalized guidance and support when issues in Azure services affect you. [Learn more](#)

Marketplace

Subscriptions

Portal settings

What's new

Feedback

Azure classic

Help + support

1 x 1

2 x 1

2 x 2

4 x 2

2 x 4

4 x 4

Remove from dashboard

Done customizing

Technet24

Figure 2-103 Customizing a dashboard

When you're satisfied with your dashboard, click on **Done Customizing** to close the customization screen.

You can create new dashboards for specific purposes by clicking the plus sign (shown in [Figure 2-102](#)) next to your dashboard name. This takes you into a customization screen for your new dashboard just like the one shown in [Figure 2-103](#).

In [Figure 2-104](#), we've created a dashboard specific to Web Apps. You can easily switch between this dashboard and the default dashboard by clicking the down arrow next to the dashboard name.




Figure 2-104 Switching between dashboards

Azure and PowerShell

If you're a PowerShell user, you can take advantage of that knowledge to manage your Azure resources using

the Azure PowerShell Az module. This module offers cross-platform support, so whether you're using Windows, Linux, or macOS, you can use the PowerShell Az module.

More Info Azurerm And Az

The PowerShell Az module is relatively new. Prior to it, all PowerShell commands used the AzureRm module. The commands that you use with both are identical. The only difference is the module name.

More Info INSTALL POWERSHELL ON LINUX OR MACOS

If you're running Linux, you can find details on installing PowerShell at <https://docs.microsoft.com/powershell/scripting/install/installing-powershell-core-on-linux?view=powershell-6>. MacOS users can find steps at <https://docs.microsoft.com/powershell/scripting/install/installing-powershell-core-on-macos?view=powershell-6>.



Exam Tip

The PowerShell Az module uses the .NET Standard library for functionality, which means it will run with PowerShell version 5.x or 6.x. PowerShell 6.x is cross-platform and can run on Windows, Linux, or macOS.

If you're running Windows 7 or later

and you have PowerShell 5.x, you'll also need to install .NET Framework 4.7.2.

Before you can use the PowerShell Az module, you'll need to install it. To do that, you first need to run PowerShell elevated. In Windows, that means running it as an Administrator. In Linux and macOS, you'll need to run it with superuser privileges using sudo.

To install the module, run the following command.

[Click here to view code image](#)

```
Install-Module -Name Az -AllowClobber
```

When you install a new PowerShell module, PowerShell checks all existing modules to see if they have any command names that are the same as a command name in the module you're installing. If they do, the installation of the new module fails. By specifying -AllowClobber, you are telling PowerShell that it's okay for the Az module to take precedence for any commands that also exist in another module.

If you are unable to run PowerShell elevated, you can install the module for your user ID only by using the following command.

[Click here to view code image](#)

```
Install-Module -Name Az -AllowClobber -Scope CurrentU
```

Once you've installed the module, you need to sign in with your Azure account. To do that, run the following command.

```
Connect-AzAccount
```

This command will display a token in the PowerShell window. You'll need to browse to <https://microsoft.com/devicelogin> and enter the code in order to authenticate your PowerShell session. If you close PowerShell, you'll have to run the command again in your next session.

More Info Persisting Credentials

It is possible to configure PowerShell to persist your credentials. For more information on doing that, see: <https://docs.microsoft.com/powershell/azure/context-persistence>.

If you have more than one Azure subscription, you'll want to set the active subscription so that commands you enter will impact the desired subscription. You can do that using the following command.

[Click here to view code image](#)


```
Set-AzContext -Subscription "subscription"
```



Replace **subscription** with the name or subscription ID of your Azure subscription you want to use with the Az module.

All Az module commands will have a common syntax that starts with a verb and an object. Verbs are things like **New**, **Get**, **Move**, or **Remove**. The object is the thing that you want the verb to impact. For example, the following command will create a resource group called MyRG in the South Central US region.


[Click here to view code image](#)



```
New-AzResourceGroup -Name MyRG -Location "South Central US"
```


If this succeeds, you'll see a message letting you know that. If it fails, you'll see an error. To remove the resource group, run the following command.

[Click here to view code image](#)



```
Remove-AzResourceGroup -Name MyRG
```

When this command is entered, you'll be asked to confirm whether you want to delete the resource group. Type a **y** and the resource group will be removed as shown in [Figure 2-105](#).

A screenshot of a PowerShell terminal window titled "jincheshire — pwsh — 125x34". The terminal shows the execution of two PowerShell commands. The first command is "New-AzResourceGroup -Name MyRG -Location 'South Central US'", which outputs the resource group details: "ResourceGroupName : MyRG", "Location : southcentralus", "ProvisioningState : Succeeded", "Tags :", and "ResourceId : /subscriptions/2ed0e6a0-aaf7-4f5f-aff6-bf25ffdef183/resourceGroups/MyRG". The second command is "Remove-AzResourceGroup -Name MyRG", which prompts for confirmation: "Confirm", "Are you sure you want to remove resource group 'MyRG'", and "[Y] Yes [N] No [S] Suspend [?] Help (default is 'Y'): y". The user enters "y", and the command completes successfully, returning to the prompt "PS /Users/jincheshire>".

```
PS /Users/jincheshire> New-AzResourceGroup -Name MyRG -Location "South Central US"

ResourceGroupName : MyRG
Location           : southcentralus
ProvisioningState  : Succeeded
Tags               :
ResourceId          : /subscriptions/2ed0e6a0-aaf7-4f5f-aff6-bf25ffdef183/resourceGroups/MyRG

PS /Users/jincheshire> Remove-AzResourceGroup -Name MyRG

Confirm
Are you sure you want to remove resource group 'MyRG'
[Y] Yes [N] No [S] Suspend [?] Help (default is 'Y'): y
True
PS /Users/jincheshire>
```

Figure 2-105 Creating and deleting a resource group with PowerShell Az

In many situations, you will be including PowerShell commands in a script so that you can perform a number of operations at once. In that case, you won't be able to confirm a command by typing **y**, so you can use the **-Force** parameter to bypass the prompt. For example, you can delete the resource group using the following command and you won't be prompted.

[Click here to view code image](#)

```
Remove-AzResourceGroup -Name MyRG -Force
```

You can find all of the commands available with the PowerShell Az module by browsing to:

<https://docs.microsoft.com/powershell/module/?view=azps-1.3.0>.

Azure CLI

As I pointed out earlier, one of the main benefits of PowerShell is the ability to script interactions with Azure resources. If you want to script with PowerShell, however, you'll need someone who knows PowerShell development. If you don't have anyone who can do that, the Azure command-line interface (Azure CLI) is a great choice. Azure CLI can be scripted using shell scripts in various languages like Python, Ruby, and so on.

Like the PowerShell Az module, the Azure CLI is cross-platform and works on Windows, Linux, and macOS as long as you use the 2.0 version. Installation steps are different depending on your platform. You can find steps for all operating systems at:

<https://docs.microsoft.com/cli/azure/install-azure-cli?view=azure-cli-latest>.

Once you install the Azure CLI, you'll need to login to your Azure account. To do that, run the following

command.

```
az login
```


When you run this command, the CLI will open a browser automatically for you to login. Once you login, if you have multiple Azure subscriptions, you can set the default one by entering the following command.

[Click here to view code image](#)

```
az account set --subscription "subscription"
```

Replace **subscription** with the name or subscription ID you want to use.

To find a list of commands you can run with the CLI, type **az** and press Enter. You'll see a list of all the commands you can run. You can find detailed help on any command by entering the command and adding a **--help** parameter. [Figure 2-106](#) shows the help for **az resource**.



```
PS /Users/jimcheshire> az resource --help

Group
  az resource : Manage Azure resources.

Subgroups:
  link       : Manage links between resources.
  lock       : Manage Azure resource level locks.

Commands:
  create      : Create a resource.
  delete      : Delete a resource.
  invoke-action : Invoke an action on the resource.
  list        : List resources.
  move        : Moves resources from one resource group to another(can be under different
               subscription).
  show        : Get the details of a resource.
  tag         : Tag a resource.
  update      : Update a resource.
  wait        : Place the CLI in a waiting state until a condition of a resources is met.

PS /Users/jimcheshire>
```

Figure 2-106 Azure CLI help

You can take this a step further if you aren't sure what the commands do. You can, for example, run the following command to get help on the syntax for **az resource create**.



```
az resource create --help
```

This provides you with help and example commands to understand the syntax.



Exam Tip

Like PowerShell, most commands in the Azure CLI have a `--force` parameter that you can include so that no prompts are displayed. When scripting PowerShell or the CLI, you need to include this parameter or your script won't work. Watch out for examples in the AZ-900 exam that test for this kind of knowledge.

An even easier way to learn the CLI is to switch into interactive mode. This provides you with auto-complete, the scoping of commands, and more. To switch into interactive mode enter **az interactive** at the command prompt. The CLI will install an extension to add this functionality. [Figure 2-107](#) shows the Azure CLI with interactive mode active. You've typed **we** at the command prompt, and it's displaying the rest of the command in dimmed text. You can press the right arrow key to enter the dimmed text in one keystroke.

The image shows a terminal window titled "jimcheshire - Python - az interactive - 128x38". The prompt is "az>>". The user has entered "webapp create". The terminal shows a list of commands and their descriptions, separated by dashed lines. The commands are: #[cmd] (use commands outside the application), [cmd] + [param] + '??[query]' (Inject jmespath query from previous command), '??[query]' (Jmespath query of the previous command), [cmd] :: [num] (do a step by step tutorial of example), \$ (get the exit code of the previous command), %[cmd] (set a scope, and scopes can be chained with spaces), and %% .. (go back a scope). At the bottom, there are keyboard shortcuts: [F1]Layout, [F2]Defaults, [F3]Keys, [Ctrl+D]Quit, and a subscription note: Subscription: Jim's Personal Azure Account.

```
az>> webapp create

#[cmd]                : use commands outside the application
[cmd] + [param] + '??[query]': Inject jmespath query from previous command
'??[query]'           : Jmespath query of the previous command
[cmd] :: [num]        : do a step by step tutorial of example
$                     : get the exit code of the previous command
%[cmd]                : set a scope, and scopes can be chained with spaces
%% ..                : go back a scope

[F1]Layout    [F2]Defaults    [F3]Keys    [Ctrl+D]Quit    Subscription: Jim's Personal Azure Account
```

Figure 2-107 CLI interactive mode

You can install additional extensions for added functionality. Because the CLI uses an extension architecture, Azure teams are able to provide support for new functionality without having to wait for a new CLI release. You can find a list of all available extensions that Microsoft provides by running the

following command.

[Click here to view code image](#)

```
az extension list-available --output table
```

This will not only show you available extensions, but it will show you if you already have the extension installed and if there's an update you should install. To install an extension, run the following command.

[Click here to view code image](#)

```
az extension add --name extension_name
```

Replace **extension_name** with the name of the extension you want to install.

Azure Advisor

Managing your Azure resources doesn't just include creating and deleting resources. It also means ensuring that your resources are configured correctly for high-availability and efficiency. Figuring out exactly how to do that can be a daunting task. Entire books have been written on best-practices for cloud deployments. Fortunately, Azure can notify you about problems in your configuration so that you can avoid problems. It does this via the Azure Advisor.

Azure Advisor can offer advice in the area of high availability, security, performance, and cost. While the documentation states that Azure Advisor is available only for Azure VMs, availability sets, application gateways, App Service applications, SQL Server, and Azure Redis Cache, many more services are onboarded to Azure Advisor and you will get recommendations for just about all of your Azure services.

To access Azure Advisor, log into the Azure portal and click Advisor in the menu on the left. [Figure 2-108](#) shows Azure Advisor with 1 low-impact recommendation for high availability and 2 high impact recommendations for security.




Figure 2-108 Azure Advisor

To review details on a recommendation, click the tile. In [Figure 2-109](#), we have clicked the high availability tile and you can see a recommendation to create an Azure service health alert.




Figure 2-109 Advisor recommendations

You don't have to do what Azure Advisor recommends. If you click on the description, you can decide to postpone or dismiss the alert as shown in Figure 2-110. If you choose to postpone the alert, you have the option of being reminded in 1 day, 1 week, 1 month, or 3 months.



Figure 2-110 Acting on a recommendation

If you have a large number of recommendations, or if

you're not the right person to take action on the recommendations, you can download Azure Advisor recommendations as either a comma-separated values file or a PDF. Click **Download As CSV** or **Download As PDF**, as shown in [Figure 2-108](#). You can also download a file with specific recommendations by clicking the appropriate download button while reviewing details as shown in both [Figure 2-109](#) and [Figure 2-110](#).

THOUGHT EXPERIMENT

Now that you've learned about core Azure services, let's apply that knowledge. You can find the answers to this thought experiment in the section that follows.

ContosoPharm has contacted you for assistance in setting up some Azure virtual machines for hosting their Azure services. They want to ensure that their services experience high-availability and are protected against disasters that might occur in a datacenter at a particular Azure region. In addition to that, they want to ensure that a power outage at a particular datacenter doesn't impact their service in that region. They also want to be certain that their application doesn't go down in case a VM has to be rebooted for any reason.

ContosoPharm's VMs will also be using specific configurations for virtual networks, and they want to ensure that they can easily deploy these resources into

new Azure regions, if necessary, at a later time. It's critical to them that the later deployments have the exact same configuration as all other deployments because any differences can cause application incompatibilities.

Some of the VMs they are deploying are under the cost center for research and development. Other VMs are going to be used for marketing to track pharmaceutical orders. For cost reporting, it's important that they be able to report on Azure expenses for each cost center separately.

During some periods of time, ContosoPharm has noticed that their applications can cause extreme CPU spikes. They'd like a system that will account for that and possibly add additional VMs during these peak times, but they want to control costs and don't want to pay for these additional VMs when they aren't experiencing a usage spike. Any advice you can offer for that would be a bonus.

The marketing application uses a website for orders, and keeping accurate inventory in real-time is critical. ContosoPharm has sales people all over the globe, and they want to implement a system where a user who accesses the site in one particular geographic region is directed to a website running in an Azure datacenter close to them.

In addition to that, they want to ensure that they keep a copy of each order invoice. These invoices are uploaded to the website as a PDF, and they want to keep them in the cloud. They don't need to be able to run any kind of reporting on these invoices, but they do need them in case regulators ask for them at some point in the future.

All of ContosoPharm's chemicals and pharmaceuticals are kept in a large research facility. They'd like to integrate a database in that facility with their Azure services, and they need that connection to be encrypted and secure. They also need to be able to carefully track the temperature of that facility where the on-premises database is stored. They've added Internet-enabled thermostat devices in the building, but they currently have no way to ensure that they can be notified if something is out of the ordinary with the temperature.

Because of the sensitivity of on-premises inventory, they'd like to store all of the telemetry from all of the devices that monitor temperature. They currently have over 500,000 sensors that record the temperature every two seconds. The CTO of the company has told you that he believes they should be able to take all of those historical readings and set up some kind of system that will be able to predict when an anomaly is happening before it becomes a problem and puts their assets at

risk. He'd like a recommendation on how we can implement that.

The last requirement that they have is the ability to easily tell if there are any opportunities for them to reduce cost based on their Azure resource usage over time. They have invested a large amount of money in the planning of this system, and they want to ensure that additional expenses are controlled wherever possible.

Provide a recommendation to ContosoPharm that meets all their requirements. You don't need to give them specific technical details on how to implement everything, but you should point them in the right direction if you don't have specifics.

THOUGHT EXPERIMENT ANSWERS

In this section, we'll go over the answers to the thought experiment.

To ensure that their VMs are protected against disasters at a datacenter within a particular Azure region, you should recommend that ContosoPharm use availability zones. By deploying VMs in availability zones, they can ensure that VMs are distributed into different physical buildings within the same Azure region. Each building will have separate power, water, cooling system, and network.

To protect their application when a VM has to be rebooted, they should use an availability set. An availability set would provide them with multiple fault domains and update domains so that if a VM has to be rebooted, they'd still have an operational VM in another update domain.

In order to ensure consistent deployments now and in the future, ContosoPharm can create an ARM template for their deployment. By using an ARM template, they can ensure that every deployment of their resources will be identical.

To separate invoice tracking for the R&D department and marketing department, Contoso Pharm can use resource tags for each of their resources. Their Azure invoice can be filtered on these tags so they can track expenses.

To ensure that they always have enough VMs to handle load when CPU spikes, they should use scale sets. They can then configure auto-scale rules to scale out when load requires it and scale back in to control costs.

To ensure that sales people using the marketing website are directed to a datacenter that's geographically close to them, ContosoPharm should use Azure Traffic Manager with Geographic rules. This will ensure that the traffic goes to a datacenter closest to the

DNS server that made the request.

To store their invoices in the cloud, ContosoPharm can use Azure Blob Storage. They could store them in a database as binary blobs, but since they don't need to run any kind of reporting or queries against them, Azure Blob Storage will be cheaper.

To connect their on-premises database to Azure resources, they can use a VPN with VPN Gateway. This allows them to set up an encrypted tunnel between their on-premises resources and their Azure virtual network.

To monitor their on-premises thermostat devices, ContosoPharm can use IoT Hub. They can set up alerts to notify someone when temperatures are outside of normal range. They can even use the device twin to configure tags so that they can set up different rules for different groups of IoT devices. Because they will need to add over 500,000 devices, they can use IoT Hub Device Provisioning Service to provision all of those devices.

You can advise the CTO to route IoT data from IoT Hub to Azure Data Lake Storage. You can then use Azure Databricks to clean that data and feed it into Azure Machine Learning Service. If their developers can develop a ML model that can be trained to discover anomalies, they can score that model to determine if they can reliably predict a problem before it happens. If

they don't have anyone with the expertise to develop a model, they can likely get that work done without programming using Machine Learning Studio. The model can then be exposed as a web service that can be called by another application.

Finally, to ensure they are taking action to reduce costs as much as possible, you can advise them to make use of Azure Advisor to take action on any cost recommendations.

CHAPTER SUMMARY

This chapter covered a lot of ground! Not only did you learn some of the basics of Azure related to regions and resource groups, but you learned about a lot of the core services Azure provides. You also learned about some of the hottest topics in technology today: IoT, machine learning, and serverless computing. We wrapped it up with information on how you can use some of the management tools Azure provides.

Here's a summary of what this chapter covered.

- An Azure region is an area within a specific geographical boundary, and each region is typically hundreds of miles apart.
- A geography is usually a country, and each geography contains at least two regions.
- A datacenter is a physical building within a region, and each datacenter has its own power, cooling supply, water support, generators, and network.

- Round-trip latency between two regions must be no greater than 2ms, and this is why regions are sometimes defined as a “latency boundary.”
- Customers should deploy Azure resources to multiple regions to ensure availability.
- Availability zones ensure that your resources are deployed into separate datacenters in a region. There are at least three availability zones in every region.
- Azure Resource Manager (ARM) is how Azure management tools create and manage Azure resources.
- ARM uses resource providers to create and manage resources.
- An ARM template allows you to ensure consistency of large Azure deployments.
- Resource groups allow you to separate Azure resources in a logical way, and you can tag resources for easier management.
- Azure Virtual Machines are an IaaS offering where you manage the operating system and configuration.
- Availability sets protect your VMs with fault domains and update domains. Fault domains protect your VM from a hardware failure in a hardware rack. You are protected from VM reboots by update domains.
- Scale sets allow you to set up auto-scale rules to scale horizontally when needed.
- Containers allow you to create an image of an application and everything needed to run it. You can then deploy this image to Azure Container Instances, Azure Kubernetes Service, or Web App for Containers.
- An Azure virtual network (VNET) allows Azure services to communicate with each other and the Internet.
- You can add a public IP address to a VNET for inbound Internet connectivity. This is useful if a website is running in your VNET

and you want to allow people to access it.

- Azure Load Balancer can distribute traffic from the Internet across multiple VMs in your VNET.
- Azure Application Gateway is a load balancer well-suited to HTTP traffic and is a good choice for websites.
- VPN Gateway allows you to configure secure VPN tunnels in your VNET. This can be used to connect across Azure regions or even to on-premises machines.
- Azure Content Delivery Network caches resources so that users can get a faster experience across the globe.
- Azure Traffic Manager is a DNS-based solution that can help to load balance web requests, send traffic to a new region in an outage, or send users to a particular region that's closest to them.
- Azure Blob Storage is a good storage option for unstructured data such as binary files.
- If you need to move a large amount of data to Blob Storage, Azure Data Box is a good option. You can have hard drives of numerous sizes shipped to you. Add your data to them and ship them back to Microsoft where they'll be added to your storage account.
- Azure Queue Storage stores messages from applications in a queue so they can be processed securely.
- Azure Disk Storage is virtual disk storage for Azure VMs. Managed Disks allow you to remove the management burden of disks.
- Azure Files allows you to have disk space in the cloud that you can map to a drive on-premises.
- Azure SQL Database is a relational database system in the cloud that is completely managed by Microsoft.
- Azure Cosmos DB is a NoSQL database in the cloud for unstructured data.

- The Azure Marketplace is a source of templates for creating Azure resources. Some are provided by Microsoft and some are provided by third-parties.
- The Internet of Things (IoT) refers to devices with sensors that communicate with each other and with the Internet.
- Azure IoT Hub allows you to manage IoT devices and route message to and from those devices.
- Azure IoT Hub Provisioning Service makes it easy to provision a large number of devices into IoT Hub.
- Azure IoT Central is a SaaS offering for monitoring IoT devices.
- Big data refers to more data that you can analyze through conventional means within a desired time-frame.
- Big data is stored in a data warehouse. In Azure, that can be Azure SQL Data Warehouse or Azure Data Lake Storage. SQL Data Warehouse is good for relational data. Data Lake Storage is good for any type of data.
- HDInsight is Microsoft's solution for clustered Hadoop processing of big data.
- The process of AI decision making at several points along the neural network is referred to as the ML pipeline.
- Azure Databricks is a good solution for modeling data from a data warehouse so that it can be effectively used in ML modeling.
- Databricks clusters are made up of notebooks that can store all types of information.
- Azure Machine Learning Service uses cloud-based resources to train ML models much faster.
- Azure Machine Learning Studio allows you to build, train, and score ML models in a drag-and-drop interface.
- Serverless computing refers to using surplus VMs in Azure to run your code on-demand. You pay only for when your code runs.

- Azure Functions is the compute component of serverless in Azure.
- Azure Logic Apps is a workflow serverless solution that uses connectors, triggers, and actions.
- Azure Event Grid makes it possible to raise and handle events as you interact with your Azure resources.
- The Azure portal is a web-based interface for interacting with your Azure services. It uses ARM API calls under the hood to talk to Azure Resource Manager.
- Azure PowerShell Az is a cross-platform PowerShell module that makes it easy to manage Azure resources in PowerShell.
- The Azure CLI is a command-line tool that is cross-platform and can be scripted in multiple languages.
- Azure Advisor provides best practice recommendations in the area of high-availability, security, performance, and cost.

Chapter 3. Understand security, privacy, compliance, and trust

As businesses move to the cloud, one of the most important concerns is the security of applications and data in the cloud. The term security, however, is a broad term with many meanings. Businesses want to be sure that someone with malicious intentions can't impact an application's availability or access application data, but they also want to ensure that their own employees' access to cloud resources is controlled. In addition, businesses also have many legal standards and policies they must comply with, and there must be trust that a cloud provider meets those standards.

Azure has services and features that focus on all of these concerns. Applications that are exposed to the public Internet are protected by network features such as Azure Firewall, Network Security Groups, and DDoS Protection. Your Azure account is protected with identity offerings such as Azure Active Directory and multi-factor authentication. Data and other assets, such as documents and emails, are protected by features such

as Azure Key Vault and Azure Information Protection.

Even with these protections in place, Azure can constantly monitor your cloud resources for signs of attack using Azure Advanced Threat Protection. Azure Security Center provides recommendations based on your resources, and how they're configured, so that you can proactively prevent threats in the first place.

Features like Azure Policies and Role-Based Access Control ensure that users with legitimate access to your subscription are able to access only those resources you desire, and you can even lock resources so they can't be reconfigured or deleted by mistake. Azure Monitor and Azure Service Health ensure you always know what's happening with your Azure resources.

Finally, Azure has many services and features that focus on compliance and data protection standards. Microsoft Trust Center is a website that teaches you how Azure secures and protects your data in the cloud, and the Service Trust Portal is a website offering tools such as Azure Compliance Manager that helps ensure you're compliant with standards that are important to you.

With that as your framework, here are the skills that we'll cover in this chapter.

Skills covered in this chapter:

- Understand securing network connectivity in Azure

- Describe core Azure Identity services
- Describe security tools and features of Azure
- Describe Azure governance methodologies
- Understand monitoring and reporting options in Azure
- Understand privacy, compliance, and data protection standards in Azure

SKILL 3.1: UNDERSTAND SECURING NETWORK CONNECTIVITY IN AZURE

The primary attack vector for applications and data in the cloud is the network, and if your application is exposed to the public Internet, the threat is much greater. Web applications are often the target of attacks that are intended to either bring an application down or to obtain unauthorized access to data. Threats from the outside, however, aren't the only source of vulnerabilities. In order to keep your application and data secure, it's also important that you control the traffic inside your virtual network in Azure. In this skill section, you'll learn about several Azure services and features designed to address these security concerns.

This section covers:

- Azure Firewall
- DDoS Protection
- Network Security Groups
- Choosing an appropriate Azure security solution

Azure Firewall

In computing parlance, a firewall is an appliance through which network traffic into and out of a particular network travels. The purpose of a firewall is to allow only desired traffic on the network, and to reject any traffic that might be malicious or that comes from an unknown origin. A firewall imposes control on the network using rules that specify a source and destination IP address range and port combination.

In a typical firewall configuration, all traffic is denied by default. In order for the firewall to allow traffic to pass through it, a rule must match that traffic. For example, if you want to allow someone on the public Internet to access a web application you have running on a particular server, create a firewall rule that allows communication to ports 80 and 443 (the ports for HTTP and HTTPS traffic). You then configure the rule to send that traffic to your web server.

There are several firewalls available from third-parties in the Azure Marketplace, but Microsoft also offers its own firewall called Azure Firewall. Azure Firewall is a PaaS offering in Azure, and it's easily managed and offers a 99.95% SLA. Azure Firewall scales according to your networking needs, so you don't have to worry about traffic spikes causing latency or

downtime for your applications.

Note Azure Firewall Is A Stateful Firewall

Azure Firewall is a stateful firewall. That means that it stores data in its memory about the state of network connections that flow through it. When new network packets for an existing connection hit the firewall, it's able to tell if the state of that connection represents a security threat.

For example, if someone spoofs your IP address and attempts to gain access to your virtual network in Azure, the firewall would recognize that the hardware address of the computer being used has changed and reject the connection.

A typical setup for Azure Firewall consists of the following:

- A centralized hub network that contains the Azure Firewall and a VM that operates as a *jumpbox*. The firewall exposes a public IP address, but the jumpbox VM does not.
- One or more additional networks (called *spoke* networks) that don't expose a public IP address. These networks contain your various Azure resources.

The jumpbox is a VM that you can remote into in order to manage other VMs in your networks. All other VMs are configured to only allow remote access from the jumpbox VM's IP address. If you want to access a VM in a spoke network, you first remote into the jumpbox VM, and then you remote into the spoke network VM from the jumpbox. This setup is referred to as a *hub and spoke* configuration, and it provides

additional security for your network resources.

Note Other Network Configurations Are Possible

A hub and spoke configuration isn't the only configuration where Azure Firewall can be used. For example, you might have a single virtual network and Azure Firewall in that network to filter traffic from the Internet. A hub and spoke network configuration is the most common in real-world business applications.

Figure 3-1 is an illustration of a typical hub and spoke configuration that also includes Azure Firewall. Traffic that comes from the Internet over port 443 (HTTPS traffic) is directed by the firewall to a web server running in Spoke VNet 1. Traffic that comes in over the remote desktop port is directed to the jumpbox VM, and users can then RDP from the jumpbox VM to a VM in Spoke VNet 2.




Figure 3-1 An example of a hub and spoke network configuration with Azure Firewall

Before you can configure a firewall to handle network traffic, you'll need to create an instance of Azure Firewall. You can choose to include Azure Firewall when you create your virtual network in Azure, or you can create a firewall and add it to an existing virtual network. Figure 3-2 shows Azure Firewall being created during the creation of a new virtual network.

The image shows a screenshot of the Azure Firewall creation interface. At the top, there is a 'Firewall' section with a toggle switch set to 'Enabled'. Below this, there are four required fields marked with a red asterisk: 'Firewall name' with the value 'FW01', 'Firewall subnet address space' with the value '10.0.1.0/26' (which also shows '10.0.1.0 - 10.0.1.63 (64 addresses)' below it), 'Public IP address' with radio buttons for 'Create new' (selected) and 'Use existing', and 'Public IP address name' with the value 'azureFirewalls-ip'. Below these fields, the 'Public IP address SKU' is set to 'Standard'. At the bottom, there is a blue 'Create' button and a link for 'Automation options'.

Figure 3-2 Creating Azure Firewall

When you create a firewall during the creation of a virtual network, Azure creates a subnet in the virtual network called AzureFirewallSubnet, and it uses the

address space you specify for that subnet. A public IP address is also created for the firewall so that it can be accessed from the Internet.

While the PaaS nature of Azure Firewall does remove much of the complexity, using a firewall isn't as simple as enabling it in your virtual network. You will also need to tell Azure to send traffic to the firewall, and then you'll need to configure rules in the firewall so that it knows what to do with that traffic.

To send traffic to your firewall, you need to create a route table. A route table is an Azure resource that is associated with a subnet, and it contains rules (called *routes*) that define how network traffic in the subnet is handled.

A route table is created using the Route Table item in the Azure Marketplace. Once you create a new route table, you must associate it with one or more subnets. To do that, click on **Subnet** and then click **Associate** as shown in [Figure 3-3](#).




Figure 3-3 Associating a route table with a subnet

After you click Associate, select the virtual network and the subnet as shown in [Figure 3-4](#).




Figure 3-4 Choosing a subnet to associate

In our particular setup, we want to associate both the JumpboxSubnet and the ServerSubnet with the route table. This will ensure that the firewall will handle all network traffic to the jumpbox VM and all traffic out of the ServerSubnet.



Exam Tip

It's important to understand that a firewall can (and should) be used to filter traffic flowing into and out of a network. For example, you want the

firewall to handle traffic into your jumpbox, but you also want to ensure that traffic flowing from the subnet where other servers are located is secure and not inappropriately sending data out of your network.

Once we've associated the route table with the subnets, we create a user-defined route so that traffic is directed through Azure Firewall. To do that, click on **Routes** and then click **Add** as shown in Figure 3-5.





Figure 3-5 Adding a new user-defined route to the route table

Figure 3-6 shows the configuration of a new user-defined route named ToFirewall. This route is configured for 0.0.0.0/0, the notation for all traffic. It's then sending that traffic to a virtual appliance (Azure Firewall in this case) located at IP address 10.1.1.4, which is the internal IP address of this firewall. Once this route is configured, it will immediately apply to all devices on the subnets associated with the route table.



Add route □ ×

AZ900RouteToFirewall

* Route name

ToFirewall ✓

* Address prefix ⓘ

0.0.0.0/0 ✓

Next hop type ⓘ

Virtual appliance ▾

* Next hop address ⓘ

10.1.1.4 ✓

ⓘ Ensure you have IP forwarding enabled on your virtual appliance. You can enable this by navigating to the respective network interface's IP address settings.

OK

Figure 3-6 Adding a user-defined route

Remember that Azure Firewall blocks all traffic by default, so at this point, there's no way to reach the

jumpbox VM that's in the JumpboxSubnet. In order to access that VM, you must configure a firewall rule in Azure Firewall that will forward the appropriate traffic to the jumpbox VM.

To add a firewall rule, open Azure Firewall in the Azure portal and click on **Rules**, select the type of rule, and click the **Add** button to add a new rule collection as shown in Figure 3-7.




Figure 3-7 Azure Firewall rule collections in the Azure portal

There are three types of rule collections available in Azure Firewall.

- **Network address translation (NAT) rules** These rules are used to forward traffic from the firewall to another device on the

network.

- **Network rules** These are rules that allow traffic on specific IP address ranges and ports that you specify.
- **Application rules** Application rules are used to allow applications such as Windows Update to communicate across your network. They can also be used to allow particular domain names such as azure.com and microsoft.com.

Azure Firewall combines all of the rules of a specific type and priority into a rule collection. The priority is a number between 100 and 65,000. Lower numbers represent a higher rule priority and are processed first. In other words, if you want to ensure that a rule is always applied before all other rules, include that rule in a rule collection with a priority of 100.

When network traffic enters the firewall, NAT rules are applied first. If the traffic matches a NAT rule, Azure Firewall applies an implicit network rule so that the traffic can be routed appropriately, and all further rule processing stops.

If there isn't a NAT rule that matches the traffic, network rules are applied. If a network rule matches the traffic, all further rule processing is stopped. If there isn't a network rule that applies to the traffic, the application rules are applied. If none of the application rules match the traffic, the traffic is rejected by the firewall.

To allow access to remote into the jumpbox VM, you might configure a NAT rule that forwards any traffic on port 55000 to port 3389 (the port for remote desktop) on the internal IP of the jumpbox VM as shown in Figure 3-8. Because port 55000 is a general port that wouldn't normally be used for remote desktop, someone with malicious intent would likely never discover that it's being used for that purpose.

Add NAT rule collection

* Name: ToJumpbox

* Priority: 100

* Action: Destination Network Address Translation (DNAT)

Rules

NAME	PROTOCOL	SOURCE ADDRESSES	DESTINATION ADDR...	DESTINATION PORTS	TRANSLATED ADDRE...	TRANSLATED PORT
JumpboxRDP	TCP	*	20.45.5.10	55000	10.1.0.4	3389
	0 selected	* 192.168.10.1, 192...	192.168.10.0	8080	192.168.10.0	8080

Figure 3-8 Adding a NAT rule

In addition to rules that you configure, the threat intelligence feature (currently in preview) in Azure Firewall can protect you from known malicious IP addresses and domain names. Microsoft constantly updates their list of known bad actors, and the data

collected is provided in the Microsoft Threat Intelligence feed.

When you enable threat intelligence, you can choose to have Azure alert you if traffic from a known malicious IP address or domain name attempts to enter your network. You can also choose to have the traffic denied by the firewall automatically as shown in [Figure 3-9](#).




Figure 3-9 Threat intelligence can help protect your Azure virtual network

For every hour that you have an Azure Firewall

deployed, Microsoft will bill you \$1.25. You are also billed \$0.03 for every gigabyte of data processed by the firewall.

DDoS Protection

Cloud applications that are accessible from the Internet over a public IP address are susceptible to *distributed denial of service* (DDoS) attacks. DDoS attacks can overwhelm an application's resources and can often make the application completely unavailable until the attack is mitigated. DDoS attacks can also be used to exploit security flaws in an application and attack systems that an application connects to.

Azure can help protect against DDoS attacks with DDoS protection. DDoS protection is a feature of Azure Virtual Networks. There are two tiers of DDoS protection; Basic and Standard. Basic protects you from volume-based DDoS attacks by distributing large amounts of volume across all of Azure's network infrastructure. Basic DDoS protection applies to both IPv4 and IPv6 public IP addresses. With the Basic tier, you have no logging or reporting of any DDoS mitigation, and there's no way to configure alerts, so that you're notified if a problem is detected. However, the Basic tier is free and provides basic protection.

The DDoS Standard tier offers protection not only from volume-based DDoS attacks, but when used in

combination with Azure Application Gateway, it also protects from attacks designed to target the security of your applications. It offers logging and alerting of DDoS events and mitigations, and if you need help during a DDoS attack, Microsoft provides access to experts who can help you. The DDoS Standard tier applies only to IPv6 public IP addresses.

The Standard tier is targeted at enterprise customers and is billed at \$2,994.00 per month, plus a small fee per gigabyte for data that is processed. The fixed monthly price covers up to 100 resources. If you need to cover additional resources, you pay an additional \$30.00 per resource per month.

To enable the DDoS Standard tier, click **DDoS Protection** in your virtual network in the Azure portal and select **Standard** as shown in [Figure 3-10](#).




Figure 3-10 DDoS protection in the Azure portal

To enable Standard tier, you'll need a DDoS protection plan. If you don't currently have one, click **Create A DDoS Protection Plan** to create one in the Azure portal. You can then apply that DDoS protection plan to your virtual network and to other virtual networks that you have access to in Azure. Virtual networks that use the DDoS protection plan aren't

required to be in the same subscription, so in most cases, an organization will only need a single DDoS protection plan to protect all of their virtual networks.



Exam Tip

The fact that you can add virtual networks from multiple Azure subscriptions to the same DDoS protection plan is an important concept. You are billed a large monthly charge for the DDoS protection plan, and if you create two DDoS protection plans, you have just doubled your costs.

DDoS protection Standard tier monitors your network traffic 24/7, and it uses machine learning to profile your traffic over time and adjust itself to accommodate your network's traffic profile. During a DDoS event, Standard tier allows you to stream logs to a *security information and event management* (SIEM) system. SIEM systems are designed to allow for the aggregation of data from a large number of sources for the purpose of analysis, and to comply with data retention policies and standards.

Once you've configured any alerts and monitoring for

DDoS protection, you can simulate a DDoS event using a BreakingPoint Cloud account available at:

<https://www.ixiacom.com/products/breakingpoint-cloud>. This allows you to ensure that your DDoS protection is protecting you from DDoS attacks.

Network Security Groups

A Network Security Group (NSG) allows you to filter traffic on your network and apply rules on that traffic. An NSG contains several built-in rules provided by Azure that are designed to allow your resources in the virtual network to communicate with each other. You can then add your own rules to the NSG to control traffic into and out of the network, and also between resources in the network.

Figure 3-11 shows the multi-tier application first shown in Chapter 2, Understanding core Azure services.




Figure 3-11 A multi-tier application

Here's the traffic flow of this application.

- Subnet 1 receives data from another virtual network running Azure Firewall.
- Subnet 1 communicates with Subnet 2 to process requests.
- Subnet 2 communicates with a database server in Subnet 3 in order to access data.

If you want to ensure a secure environment, Subnet 1 should not be able to directly communicate with resources in Subnet 3. Likewise, Subnet 3 should not be able to directly communicate with resources in Subnet 1. Finally, only Subnet 1 should be able to communicate with the other virtual network running Azure Firewall.

You can use NSGs to implement rules that will enforce these policies.

NSGs can be associated with a subnet or to a network interface attached to a VM. Each network interface or subnet can only have one NSG associated with it, but you can create up to 1,000 rules in a single NSG, so you should be able to easily apply all of the rule logic necessary for any task. If you associate an NSG to both a subnet and to one or more network interfaces inside of that subnet, the rules for the NSG associated with the network interfaces are first applied, and then the subnet's NSG's rules are applied.



Exam Tip

An NSG that's associated with a subnet impacts all VMs inside of that subnet, as well as traffic to and from the subnet. For example, if you configure an NSG to prevent all traffic except traffic from the Internet, and you then associate that NSG with a subnet containing two VMs, those two VMs will no longer be able to communicate with each other because only traffic from the Internet is allowed by the NSG.

To prevent rules from interfering with each other, each rule you create in an NSG has a priority between 100 and 4,096. Rules with a lower priority take precedence over rules with a higher priority. Network traffic is applied against the rule with the lowest priority number first. If the traffic matches that rule, the rule is applied, and processing of the rule stops. If the traffic doesn't match the rule, it is evaluated against the next lowest priority rule. This continues until the traffic has either matched a rule, or there are no additional rules.

More Info Priority Of Default Rules

The default rules that Azure applies to all NSGs have a priority in the 65,000 range. This prevents the default rules from ever overriding an explicit rule that you create, and it makes it easier for you to override the default rules if needed.

To create an NSG, search for Network Security Group in the Azure Marketplace. When you create an NSG, give it a name, choose or create a resource group, and specify the location for the NSG, as shown in [Figure 3-12](#).


The image shows a web form titled "Create network security g..." with standard window controls. It contains four required fields, each marked with a red asterisk: "Name" with the value "Subnet-1-NSG" and a green checkmark; "Subscription" with a dropdown menu showing "Jim's Personal Azure Account"; "Resource group" with a dropdown menu showing "SecurityRG" and a link "Create new" below it; and "Location" with a dropdown menu showing "South Central US". At the bottom, there is a blue "Create" button and a link "Automation options".

Figure 3-12 Creating an NSG

After you create an NSG, you can then add inbound and outbound rules for the NSG. Once you open the NSG in the Azure portal, click **Inbound Security Rules** to add new inbound rules, and **Outbound Security Rules** to add outbound rules.

In [Figure 3-13](#), we click on Inbound security rules to add a new rule that allows traffic from the virtual network running Azure Firewall. After that, the NSG will be associated with Subnet-1. Note that you can associate the NSG with a subnet or network interface before adding rules.




Figure 3-13 Inbound security rules for an NSG

Click Add to add a new NSG rule. [Figure 3-14](#) shows a new rule being added that allows traffic into this subnet from the address space of another virtual network that's running Azure Firewall.

The rule being configured in [Figure 3-14](#) uses CIDR notation for the source IP addresses, but you can also enter a specific IP address or change the Source dropdown to **Any** if you want the rule to apply to all IP addresses.

Click **Subnets** to associate an NSG with a subnet, or **Network Interfaces** to associate it with a network interface used by a VM. Then click **Associate**, as shown in [Figure 3-15](#).




Figure 3-15 Associating an NSG

Figure 3-16 shows the blade where an NSG is associated with a subnet.




Figure 3-16 Associating an NSG with a subnet

Outbound security rules are created in the same way that inbound rules are created. You aren't required, however, to create a corresponding outbound rule for every inbound rule. NSGs maintain what's called a *flow record* that stores the state of a connection, and the NSG will allow traffic that corresponds to that flow record without an explicit rule. If a security rule allows inbound traffic to port 80 from IP addresses in the range of 10.1.0.0/16, such as the rule configured in Figure 3-14, the NSG will also allow outbound traffic on port 80 to addresses in that same range using the flow record. Only once traffic stops flowing for a few

minutes, will the flow record no longer be in effect.

There are some cases where you won't know the specific IP address range. For example, if you want to configure an NSG rule on a virtual network that allows all traffic from the Internet, you wouldn't specify an exact address range. To deal with that, NSGs allow you to use *service tags* when configuring rules.

A service tag is a special identifier created by Microsoft that applies to the Internet or to a specific service type within Azure. For example, if you have some web apps running in Azure App Service, and you want to allow them to communicate with your subnet, you can use the AppService service tag in your inbound rule to allow that. Azure services also have region-specific service tags so that you can allow or deny traffic only from specific regions.

To use a service tag, set the Source of your rule to Service Tag. You can then select a service tag from the Source service tag dropdown. In [Figure 3-17](#), the AppService.CentralUS service tag is being used to allow traffic from Azure App Service resources in the Central US region.

Add inbound security rule Subnet-1-NSG

Basic

* Source ⓘ
Service Tag

* Source service tag ⓘ
AppService.CentralUS

* Source port ranges ⓘ
*

* Destination ⓘ
Any

* Destination port ranges ⓘ
443

* Protocol
Any TCP UDP

* Action
Allow Deny

Figure 3-17 Using a service tag in an NSG rule

Choosing an appropriate Azure security solution

Depending on your application and its configuration, you may or may not need all of the security solutions we've discussed. NSGs should be considered in almost all scenarios to ensure that only the desired network traffic flows between your resources. If your application is exposed to the Internet, NSGs can ensure that traffic

from the Internet is only allowed into specific subnets or VMs. Even if you aren't using a public-facing IP address, NSGs can help you enforce secure communication between layers of an application.

Azure Firewall is a powerful way to ensure that traffic into your virtual networks is tightly controlled. Unlike an NSG, Azure Firewall is a stateful solution that understands the makeup of a network connection and can identify if an attack is being attempted on your network. Azure Firewall is not necessary, however, if your application doesn't expose a public IP address.

DDoS protection is an effective means of protecting your network from attacks designed to impact your applications with a large volume of seemingly legitimate traffic. You can add another layer of protection by using Application Gateway alongside of DDoS protection Standard tier to protect against security threats to the application.

SKILL 3.2: DESCRIBE CORE AZURE IDENTITY SERVICES

Security isn't only about controlling network traffic. In order to provide a secure environment, you must have some means of identifying who's accessing your application. Once you know the identity of a user, you need to ensure that they aren't allowed access to data or other resources that they shouldn't access.

Authentication is the process by which a user's identity is confirmed. Once a user is authenticated and begins interacting with an application, additional checks may take place to confirm which actions the user is and isn't allowed to perform. That process is called authorization, and authorization checks are performed against a user who is already authenticated.

Azure offers a service called Azure Active Directory that provides both authentication and authorization capabilities for resources and applications, both in the cloud and on-premises.

This section covers:

- Azure Active Directory
- Multi-factor authentication

Azure Active Directory

If you have any experience with on-premises Windows Active Directory, you might find understanding Azure Active Directory (Azure AD) to be a challenge. That's because Azure AD isn't the cloud-equivalent of Windows Active Directory. It's entirely different.

Azure AD is a cloud-based identity service in Azure that can help you to authenticate and authorize users. You can use Azure AD to give users access to Azure

resources. You can also give users access to third-party resources used by your company and on-premises resources, all using the same username and password.

More Info Granting Access To Azure Resources

You'll learn about how you can give other users access to your Azure resources when we cover role-based access control in Skill 3.4, "Describe Azure governance methodologies."

The core of Azure AD is a directory of users. Each user has an *identity* that's comprised of a user ID, a password, and other properties. Users also have one or more *directory roles* assigned to them. The user ID and password are used to authenticate the user, and the roles are used for authorization to perform certain activities in Azure AD.

When you sign up for an Azure subscription, an Azure AD resource is automatically created for you, and it's used to control access to Azure resources you create under your subscription. [Figure 3-18](#) shows Azure AD in the Azure portal.




Figure 3-18 Azure AD in the Azure portal

To view or manage users in Azure AD, click on **Users** in the menu on the left side of the page. This opens the All Users blade shown in [Figure 3-19](#).




Figure 3-19 The All Users blade in the Azure portal

The Azure AD shown in [Figure 3-19](#) contains two users. The first user's source is Microsoft Account, meaning this user is tied to a Microsoft Account email address. The other user is an Azure Active Directory user that was manually added.

To add a new user from your company to your Azure

AD, click on **New User** to display the blade shown in Figure 3-20.

The screenshot shows a management console with two side-by-side blades. The left blade is titled 'User' and contains fields for 'Name' (James Taylor), 'User name' (james@contosopharm.com), 'Profile' (Not configured), 'Properties' (Default), 'Groups' (0 groups selected), and 'Directory role' (User). It also has a 'Password' section with a masked input field and a 'Show Password' checkbox. A 'Create' button is at the bottom. The right blade is titled 'Groups' and has a search bar labeled 'Select' with the placeholder 'Search by name or email address'. Below the search bar is a list of groups, with 'Sales Excs' (SE) visible. A 'Selected groups' section shows 'No groups selected'. A 'Select' button is at the bottom.

Home > Jim's Directory > Users - All users > User > Groups

User

Jim's Directory

* Name ⓘ
James Taylor ✓

* User name ⓘ
james@contosopharm.com ✓

Profile ⓘ
Not configured >

Properties ⓘ
Default >

Groups ⓘ
0 groups selected >

Directory role ⓘ
User >

Password
[Masked Password]
☐ Show Password

Create

Groups

Select groups in which this user is to be a member

Select ⓘ
Search by name or email address ✓

SE Sales Excs

Selected groups:
No groups selected

Select

Figure 3-20 Adding a new Azure AD user

The specified user name is used to log into Azure AD. The domain name you use must be one that you own and that is associated with your Azure AD. You can also click on **Groups** to pick a group for this user. Groups makes it easier to manage larger groups of similar users.

Azure AD offers a feature called Azure AD B2B (business-to-business) collaboration that allows you to add users who don't belong to your company. So, you can invite other users from outside of your company to be members of your Azure AD. Those users can then be given access to your resources. Users who are not part of your company are called *guest users*. To add a guest user, click New Guest User shown in [Figure 3-19](#). This will open the New Guest User blade shown in [Figure 3-21](#).

New Guest User
Jim's Directory

 This user will be added as a Guest. Click here to learn more. 

* Email address 
chris@contoso.com 

Include a personal message with the invitation

Hey, Chris. We'd like for you to help manage our social media presence. Accept this invite to get access to our social media accounts so you can help us out!

Invite

Figure 3-21 Adding a new guest user

When you invite a guest user, an invitation to join your Azure AD is sent to the email address you specify. In order to accept the invite, the user's email address must be associated with a Microsoft Account. If the user doesn't have a Microsoft Account, they'll be given the option to create one so they can join your Azure AD.

The user in Figure 3-21 can be given access to the

corporate social media accounts by adding those applications to Azure AD. Applications to add include, not only social media applications such as Facebook and Twitter, but also thousands of others. To add an application, open Azure AD in the Azure portal, click **Enterprise Applications**, and click **New Application**, as shown in Figure 3-22.




Figure 3-22 Enterprise applications in Azure AD

After you click on New Application, you can choose from a list of included applications, as shown in [Figure 3-23](#). You can also add your own application, add an application that exists in your on-premises environment, or integrate any other application. The application that you add needs to expose a login page that you can point Azure AD to in order to integrate it.




Figure 3-23 Enterprise application gallery

After you add an application, you can configure Azure AD so that users with access to that application can authenticate to it using the same credentials they use to log into Azure AD. This kind of authentication is known as *single sign-on* (or SSO), and it's one of the key benefits to using Azure AD.



Exam Tip

Azure AD B2B allows you to invite guest users to your Azure AD from other businesses. Another AD feature called Azure AD B2C allows you to give users access to Azure AD applications by signing in with existing accounts such as a Facebook or Google account.

Another important benefit to using Azure AD for managing user access to other applications is that you can easily revoke that access from a single interface. For example, if you give a user the username and password of your social media account so they can post to your account, when you no longer want that user to have access, you'd have to change the username and password on your social media account. If, however, you grant them access using Azure AD with SSO configured, you can remove that access easily within the Azure

portal. The user never has to know the username and password you use for the social media account.

Multi-factor authentication

All the Azure AD features we've covered so far are included in the free version of Azure AD that everyone with an Azure subscription gets. Azure AD has three other pricing tiers that aren't free: Basic, Premium P1, and Premium P2. If you upgrade to one of the Premium plans, you have the ability to enable multi-factor authentication for your users.



Exam Tip

If you're using the free Azure AD plan, you have a subset of MFA features for global administrators only. These features allow you to enable MFA for global administrators when accessing the Azure portal and the Microsoft 365 admin center.

More Info Azure Active Directory Pricing

For more information on Azure AD pricing plans and what's included in each of them, see: <https://aka.ms/aadpricing>.

By default, users are able to log into your Azure AD

using only a username and password. Even if you require your users to use strong passwords, allowing access to your resources with only a username and password is risky. If a hacker obtains the password by using software that guesses passwords, or by stealing it through phishing or some other means, your resources are no longer secure.

Multi-factor authentication solves this problem. The concept behind multi-factor authentication is that you must authenticate using a combination of:

- Something you know, such as a username and password.
- Something you have, such as a phone or mobile device.
- Something you are, such as facial recognition or a fingerprint.

If multi-factor authentication requires all three of these, it's referred to as three-factor authentication, or sometimes 3FA. If only the first two are required, it's referred to as two-factor authentication, or sometimes 2FA. (Microsoft actually calls it *two-step verification*.) Azure multi-factor authentication is two-factor authentication.

Note Biometrics In Mobile Devices

Even though Azure multi-factor authentication is two-factor, if you are using a mobile device that includes biometric features, you may be authenticating using three-factor authentication. However, the third factor is enforced by your mobile device and not by Azure. Azure multi-factor authentication doesn't require three-factor authentication.

To enable multi-factor authentication for one or more users of your Azure AD, open the All Users blade and click on **Multi-Factor Authentication** as shown in Figure 3-24.




Figure 3-24 Enabling multi-factor authentication

When you click on Multi-Factor Authentication, a new browser window opens and displays the Azure AD user management site. Select one or more users you

want to enable multi-factor authentication for and click **Enable** as shown in Figure 3-25.




Figure 3-25 Enabling multi-factor authentication

You can't enable multi-factor authentication for guest users using this method. If you want to enforce multi-factor authentication for guest users, you will need to

set up conditional access to your Azure AD. To do that, open your Azure AD in the Azure portal and click on **Conditional Access** as shown in [Figure 3-26](#).




Figure 3-26 Setting up conditional access

In the Conditional Access blade, click **New Policy** as shown in [Figure 3-27](#).




Figure 3-27 Adding a conditional access policy

Enter a name for the new policy and click **Users and Groups** under Assignments. Click the **Select Users and Groups** radio button and check the **All Guest Users** checkbox. Then click **Done** as shown in [Figure 3-28](#).

Home > Jim's Directory > Conditional Access - Policies > New > Users and groups

New

Info

* Name

MFA for Guest Users ✓

Assignments

Users and groups ⓘ

0 users and groups selected >

Cloud apps ⓘ

0 cloud apps selected >

Conditions ⓘ

0 conditions selected >

Access controls

Grant ⓘ

0 controls selected >

Session ⓘ

0 controls selected >

Enable policy

On

Off

Create

Users and groups

Include

Exclude

☐ None

☐ All users

☒ Select users and groups

☒ All guest users (preview) ⓘ

☐ Directory roles (preview) ⓘ

☐ Users and groups

Done

Figure 3-28 Configuring a policy for guest users

You now need to configure which app this policy applies to. To do that, follow these steps as shown in Figure 3-29.

1. Click **Cloud Apps**.
2. Click the **Select Apps** radio button.
3. Click **Select**.
4. Select the **Microsoft Azure Management app**.
5. Click **Select**.




Figure 3-29 Adding the Microsoft Azure Management app to the policy



Exam Tip

Conditional access requires a Premium plan for Azure AD.

To add a multi-factor authentication requirement to the policy, click on **Grant** under Access Controls as shown in [Figure 3-30](#). Click **Grant Access** and check the **Require Multi-Factor Authentication** checkbox. You'll then need to click **Select** to add the access control. Finally, enable the policy and click **Create** to finish the process.

Home > Jim's Directory > Conditional Access - Policies > New > Grant

New

Info

* Name

MFA for Guest Users

Assignments

Users and groups

Specific users included

Cloud apps

1 app included

Conditions

0 conditions selected

Access controls

Grant

0 controls selected

Session

0 controls selected

Enable policy

On

Off

Create

Grant

Select the controls to be enforced.

☐ Block access

☒ Grant access

☒ Require multi-factor authentication

☐ Require device to be marked as compliant

☐ Require Hybrid Azure AD joined device

☐ Require approved client app

For multiple controls

☒ Require all the selected controls

☐ Require one of the selected controls

Select

Figure 3-30 Configuring a policy to require multi-factor authentication

SKILL 3.3: DESCRIBE SECURITY TOOLS AND FEATURES OF AZURE

Threats to your data and resources can originate from anywhere. Some threats are external, such as hackers attempting to remote into your virtual machines by guessing administrator passwords. Other threats, such as employees not abiding by security best-practices, can come from within. Ensuring the security of your cloud resources can be challenging, and as the number of Azure resources you have deployed grows over time, that challenge can grow exponentially.

This section covers:

- Azure Security Center
- Azure Key Vault
- Azure Information Protection
- Azure Advanced Threat Protection

Azure Security Center

Azure Security Center is a service in Azure that offers you a single portal for monitoring and managing the security of your Azure resources. You can also add on-premises resources to Security Center by installing a Security Center agent on your on-premises resources.

Security Center offers two tiers of service. The free

tier provides general assessment and recommendations for securing your Azure resources and covers only Azure virtual machines and Azure App Service. The Standard tier adds coverage of your Azure SQL Databases, MySQL databases, PostgreSQL, and Azure blob storage, as well as additional features such as advanced threat detection, analysis from Microsoft Threat Intelligence, and the ability to manage the regulatory compliance of your Azure resources. The Standard tier is billed by the hour, and full details on pricing can be found at <https://azure.microsoft.com/en-us/pricing/details/security-center>.

To get started with Security Center, click **Security Center** in the menu in the Azure portal. This will take you to the Overview blade where you can see an overview of all your resources being protected by Security Center as shown in [Figure 3-31](#).




Figure 3-31 Azure Security Center

There are three primary areas of coverage in Security Center.

- **Policy & Compliance** Provides a secure and overall score of how secure your resources are. This area also covers your compliance with regulatory standards.
- **Resource Security Hygiene** Provides a high-level overview of the health of your resources from a security perspective. Security issues are categorized as high, medium, or low severity.
- **Threat protection** Shows you any active or past attacks or threats on your resources.

Information for the first two areas is provided by the service being protected. This information is often related to best-practices. Threat protection, on the other hand, is specifically targeted at analyzing both the network traffic and the behavior of users of your resources. If anything looks suspicious, it's reported by Security Center.

Microsoft Threat Intelligence is used to identify security threats. Threat Intelligence uses Microsoft's historical data and machine learning to identify possible threats. These threats could be a hacker attempting to gain access to a resource, or they could be related to suspicious activity performed by a user. For example, if a user elevates his privileges on a VM and runs an unknown process, that would likely be flagged as an incident that should be investigated.

More Info Advanced Threat Protection

Threat protection information is obtained using Azure Advanced Threat Protection. You'll learn more about Advanced Threat Protection

later in this skill.

By clicking on an item in the Overview blade, you can drill down into more details. In [Figure 3-32](#), we've clicked on **Compute & Apps** in the Overview blade. You can see all of the recommendations for the VMs, App Service, cloud services, and container resources. You can also see how much your secure score will improve if you address each recommendation.




Figure 3-32 Overview of all Azure Compute recommendations in Security Center

Clicking on one of the recommendations will provide additional information. In most cases, you'll see a link to instructions on how you can address the recommendation, but Security Center has the ability to automatically take care of recommendations. For example, clicking on the recommendation to Install Endpoint Protection On My VMs, takes you to the screen shown in Figure 3-33.




Figure 3-33 Details on a Security Center recommendation

From here, you can click Install On 1 VMs to automatically install endpoint protection right from inside Security Center. In this scenario we only have one VM, but you can imagine how helpful this capability would be if you had hundreds of VMs running in Azure. With the single click of a button, you can install endpoint protection on them all.

One of the greatest security threats to your cloud resources is open network ports on your VMs. Accessing

your VMs using the remote desktop for Windows VMs, or SSH for Linux VMs, is a necessary part of managing those resources, but hackers commonly use the network ports used for remote management to break into VMs. Security Center provides a feature called just-in-time (JIT) access that helps to protect your VMs from attacks on management ports.

When JIT access is enabled, users must request access to a VM in order to remote into it. Until someone is given JIT access, management ports on the VM are closed so they can't be accessed. Once JIT access is given to a user, the ports are open for a specific period of time as requested by the user. Once that time period has elapsed, the management ports are closed again.

To enable JIT access on a VM, click on **Just In Time VM Access** in Security Center, as shown in [Figure 3-34](#). Click on the Recommended tab to see VMs that are currently not configured for JIT access. Select one or more VMs and click **Enable JIT** to turn on the feature.




Figure 3-34 Enabling JIT access on a VM

When enabling JIT access, you can choose which ports you want to protect as shown in [Figure 3-35](#). The recommended ports for management are listed, but you can add your own ports. For example, if you've changed

your VM configuration so that management happens on a non-typical port, you can add that port for JIT access.

Home > Security Center - Just in time VM access > JIT VM access configuration > Add port configuration

JIT VM access configuration

Server1

+ Add Save X Discard

Configure the ports for which the just in time VM access will be applicable

PORT	PROT...	ALLOWED SOUR...	IP RANGE	TIME RANGE (H...
22 (Recommended)	Any	Per request	N/A	3 hours ...
3389 (Recommended)	Any	Per request	N/A	3 hours ...
5985 (Recommended)	Any	Per request	N/A	3 hours ...
5986 (Recommended)	Any	Per request	N/A	3 hours ...

Add port configuration

* Port: 3389

Protocol: Any TCP UDP

Allowed source IPs: Per request CIDR block

IP addresses:

Max request time: 3 (hours)

Figure 3-35 JIT access configuration

In addition to specifying the port, you can also control which protocols are allowed over the port and which IP addresses are allowed. If the allowed IPs are set to **Per Request**, the user who requests access will have the option of specifying an IP address or a CIDR block. Otherwise, you can specify a CIDR block yourself in order to allow access only from a specific IP address range.

When a user requests access, the number of hours that access is given can be specified up to the maximum number of hours you specify in the port configuration. Maximum request time can be configured anywhere from 1 to 24 hours.

Once a VM is configured for JIT access, users request access from inside of Security Center. After clicking on **Just in Time VM Access**, select the VM and click **Request Access** as shown in Figure 3-36.




Figure 3-36 Requesting JIT access

As shown in [Figure 3-37](#), users requesting access must specify which ports to open, the IP addresses that are allowed (assuming they weren't specified when JIT access was enabled for the VM), and how long access is needed, up to the maximum time configured. Once **Open Ports** is clicked, the requested ports will remain open for the period specified.

Home > Security Center - Just in time VM access > Request access

Request access

Server1

Please select the ports that you would like to open per virtual machine.

PORT	TOGGLE	ALLOWED SOURCE IP	IP RANGE	TIME RANGE (HOURS)
▼ Server1				
22	On Off	My IP IP Range	No range	3
3389	On Off	My IP IP Range	No range	6
5985	On Off	My IP IP Range	No range	3
5986	On Off	My IP IP Range	No range	3

Open ports

Figure 3-37 Details of a JIT access request

Azure Key Vault

Most applications use sensitive or secret information. For example, an application that uses a database needs to know how to connect to that database, and that connection information is stored in a connection string. The connection string may contain a username and password that protects the database, and storing that username and password in a clear text file would be an

obvious security risk.

Azure Key Vault provides a secure way to store secrets, keys, and certificates. Once an item is stored in Key Vault, you can apply security policies that define which users and applications can access it. Key Vault is encrypted using encryption keys, but Microsoft has no visibility into the encryption keys or the encrypted data.

Key Vaults are created in the Azure portal as shown in Figure 3-38.

Home > Key vaults > Create key vault > Access policies

Create key vault

* Name ⓘ
AZ900Vault ✓

* Subscription
Jim's Personal Azure Account ▾

* Resource Group
SecurityRG ▾
[Create new](#)

* Location
South Central US ▾

Pricing tier
Standard >

Access policies
1 principal selected >

Virtual Network Access
All networks can access. >

Create

Automation options

Access policies

[Click to hide advanced access policies](#)

☐ Enable access to Azure Virtual Machines for deployment ⓘ

☐ Enable access to Azure Resource Manager for template deployment ⓘ

☐ Enable access to Azure Disk Encryption for volume encryption ⓘ

+ Add new ...

 **Jim Cheshire**
USER (Directory ID: f1a...

...

OK

Figure 3-38 Creating a Key Vault

There are two pricing tiers available in Key Vault: Standard and Premium. The only difference between the two is that keys are stored in hardware security modules (HSMs) in the Premium tier. An HSM is a separate piece of hardware that is designed for securely storing encrypted content, and it's also specialized for processing cryptographic data.



Exam Tip

Keeping encryption keys in an HSM boundary is required for Federal Information Processing Standard (FIPS) 140-2, so companies that need to maintain compliance with FIPS 140-2 can do so by using the Premium tier of Key Vault.

You can import a key, secret, or certificate into Key Vault, but Key Vault can also generate security keys and certificates for you. For example, you may want to generate a security key that your company can use to sign certificates. If you want to generate a 4,096-bit security key for this purpose and store it in Key Vault, click on **Keys** and then click **Generate/Import** as shown in [Figure 3-39](#).




Figure 3-39 Adding a key to Key Vault

In [Figure 3-40](#), a 4,096-bit RSA key is being generated and stored in Key Vault.

Home > AZ900Vault - Keys > Create a key

Create a key

Options

Generate

* Name ⓘ

SecureKey900 ✓

Key Type ⓘ

RSA EC

RSA Key Size

2048 3072 4096

Set activation date? ⓘ ☐


Set expiration date? ⓘ ☐

Enabled? Yes No

Create

Figure 3-40 Generating an RSA key

As shown in [Figure 3-41](#), once the key has been stored, you can view the entry to get the key identifier, a URL that can be used to retrieve the key by users or applications that are authorized. However, you cannot view the key because it's encrypted and not available except through the key identifier.



The screenshot shows the 'Properties' and 'Settings' tabs for an RSA key in Azure Key Vault. The key's identifier is 79b2bca07a3e4037b47ce220a94a2760. The key type is RSA with a size of 4096. It was created and updated on 3/16/2019 at 11:38:44 AM. The key identifier is displayed as a URL: https://az900vault.vault.azure.net/keys/SecureKey900/79b2bca07a3e4037b47ce220a94a27... The 'Settings' tab shows options for setting activation and expiration dates, and a toggle for 'Enabled?' set to 'Yes'. There are no tags. The 'Permitted operations' section shows all operations (Encrypt, Decrypt, Sign, Verify, Wrap Key, Unwrap Key) are checked.

79b2bca07a3e4037b47ce220a94a2760
Key Version

Save Discard

Properties

Key Type RSA

RSA Key Size 4096

Created 3/16/2019, 11:38:44 AM

Updated 3/16/2019, 11:38:44 AM

Key Identifier

<https://az900vault.vault.azure.net/keys/SecureKey900/79b2bca07a3e4037b47ce220a94a27...>

Settings

Set activation date? ☐

Set expiration date? ☐

Enabled? Yes No

Tags
0 tags

Permitted operations

☒ Encrypt	☒ Sign	☒ Wrap Key
☒ Decrypt	☒ Verify	☒ Unwrap Key

Figure 3-41 Details on a key



Exam Tip

A key stored in Azure Key Vault would

typically be accessed programmatically by an application. To protect the key, the application developers can retrieve the key each time it's needed instead of retrieving it once and storing it in memory. This ensures that the key remains secure.

Another common use scenario for Key Vault is to store encryption keys for Azure VMs. One of the security recommendations offered by Security Center is to encrypt VM disks. A VM disk is stored as a VHD file, and when it's encrypted, the host operating system that runs the VM must be able to access the security key in order to decrypt the VHD and run the VM. Key Vault offers capabilities that are specifically targeted at this kind of scenario.

In order to use Key Vault for disk encryption keys, the access policies must be configured to allow the vault for disk encryption. If this wasn't done when the vault was created, you can change it by clicking **Access Policies**, and checking the option to enable access to Azure Disk Encryption as shown in [Figure 3-42](#).




Figure 3-42 Setting access policies to allow access to Azure Disk Encryption

Azure Disk Encryption is enabled on your VMs using either Azure PowerShell or the Azure command-line interface (CLI).

More Info **ENABLING ENCRYPTION**

In order to enable encryption and store the keys in Key Vault, your VMs and Key Vault must be in the same Azure subscription, and they must be in the same Azure region. For more details on disk encryption requirements and steps to enable encryption, see: <https://docs.microsoft.com/en-us/azure/security/azure-security-disk->

Azure Information Protection

Solutions such as Azure Key Vault are effective at ensuring that information that never leaves your control is safe and secure. There are often situations where sensitive or confidential information is shared outside of your company. For those situations, Azure Information Protection (or AIP) can help you to keep information within your control.

AIP protects emails and Microsoft Office documents from reaching the wrong hands. By configuring different classifications for emails and other documents, and then specifying restrictions that apply to each classification, a company can ensure that information isn't over-shared or that sensitive information doesn't leave the company.

In Figure 3-43, an email is being sent that contains a credit card number. In order to prevent my credit card from being used by someone inappropriately, you can use AIP to classify the email so that only the recipient you specified can read it.




Figure 3-43 A sensitive email that contains confidential information

The **Protect** button in Microsoft Outlook allows you to easily categorize and protect your email. By clicking on the **Protect** button, as shown in Figure 3-44, you can mark this email so that it can only be read by the recipient that it's addressed to.




Figure 3-44 A sensitive email that contains confidential information

This message can be read by a recipient who's using Microsoft Outlook. If the recipient isn't using Microsoft Outlook, a link to read the message will be available, as shown in [Figure 3-45](#). Clicking that link will send the user a single-use passcode that they can enter to read

the email message in Office 365 in a web browser. This passcode will work even if the user isn't an Office 365 subscriber.




Figure 3-45 Opening a protected message



Exam Tip

As previously mentioned, AIP can also be used to protect Office documents in the same way. You can even control whether users can edit the document or just read its contents. If a user doesn't have the Office app used to create the document, it will open in Office Online in a web browser.

Azure Advanced Threat Protection

You've learned about some of the Azure services designed to protect your Azure resources from network attacks, and those that are designed to protect against attacks on your VMs and applications running in the cloud. One of the most common attack vectors (and one that is often most difficult to detect), however, is an attack on a user's identity within your on-premises environment, or via other devices employees use to connect to your network. Attacks on your cloud resources are often begun on-premises, or on mobile devices where the target may not be as hardened.

These types of attacks are difficult to detect because

they often look like legitimate traffic. Hackers will enter your environment on one machine using stolen credentials, and they'll move laterally through your infrastructure attempting to gain access to additional systems and sensitive data. Mobile devices are also a common attack vector because they can connect to insecure networks.

Azure Advanced Threat Protection (or ATP) is available as part of the Enterprise Mobility + Security 5 suite from Microsoft. You can also purchase it using a standalone license. ATP is designed to identify and mitigate identity threats in your on-premises environment and on devices that connect to your environment.

Deploying ATP in your environment is a multi-step process.

Step 1: Determine Capacity

ATP collects information on your network, servers, and environment using software that you install called an *ATP sensor*. In order to plan your ATP capacity, you need to determine how many sensors you need and what size those sensors should be. (Sensor size is based on the volume of network traffic in your environment.)

Microsoft has a guide on planning your ATP capacity at: <https://docs.microsoft.com/azure-advanced-threat-protection/atp-capacity-planning>.

Step 2: Create an Instance of Azure ATP

The ATP sensors that you install on-premises will connect to an Azure ATP instance in the cloud. That's where they will store all the data that's collected in order to do threat analysis and detection.

Your Azure ATP instance is created by browsing to the Azure ATP portal at <https://portal.atp.azure.com>. You'll need an Azure ATP license in order to access this site.

Step 3: Connect ATP To On-Premises Active Directory

Azure ATP connects to your on-premises Active Directory in order to get information about your environment and your users. You can also connect to a multi-forest Active Directory.

Full details on how to connect ATP to your Active Directory are available at:

<https://docs.microsoft.com/en-us/azure-advanced-threat-protection/install-atp-step2>.

Step 4: Download, Install, and Configure the ATP Sensors

The ATP sensors are available in a package from Microsoft. Once you download and install the ATP sensor, you'll need to configure it before you'll start seeing data.

You can find details on the ATP sensor installation and configuration at: <https://docs.microsoft.com/en-us/azure-advanced-threat-protection/install-atp-step3>.

Once installed and configured, ATP is able to use its

analytics and machine learning to identify what's normal and what's not, and it does this in real-time. If an attack occurs, ATP can provide tools to your IT department to investigate the nature of the attack and take appropriate actions.



Exam Tip

ATP represents an entire ecosystem at Microsoft across various offerings. Office 365 ATP is designed to offer Office 365 users with protection from threats related to email and Office documents. Windows Defender ATP helps protect Windows computers from threats and attacks. Azure ATP protects your on-premises identities and servers. It also protects mobile devices that are used to connect to your on-premises environment.

SKILL 3.4: DESCRIBE AZURE GOVERNANCE METHODOLOGIES

As your cloud presence grows, you'll likely end up with a large number of Azure resources that span many different Azure services. Unless you have some control

over how those resources are created and managed, costs can spiral out of control. In addition to cost control, you may have other restrictions you'd like in place as well, such as which regions certain resources should be created in, or how certain resources are tagged, and so on.

The traditional way of handling such governance issues would be to send out a memo to everyone explaining what the requirements were, and then crossing your fingers that people adhere to them. Fortunately, Azure Policy can ensure your requirements and policies are adhered to.

This section covers:

- Azure Policy
- Role-based access control
- Locks
- Azure Advisor

Azure Policy

Azure Policy allows you to define rules that are applied when Azure resources are created and managed. For example, you can create a policy that specifies that only a certain size VM can be created and that the VMs must be created in the South Central US region. Azure will take care of enforcing this policy so you remain in

accordance with your corporate policies.

To access Azure Policy, type **policy** in the search box in the Azure portal and click on **Policy**. Alternatively, you can click on **All Services** and search for “policy” in the list. This will display the Policy blade as shown in Figure 3-46.




Figure 3-46 Azure Policy in the Azure portal

By default, Azure Policy shows your compliance with policies defined on an Azure subscription. If you want to, you can scope this view to a different subscription or to a resource group by clicking the ... button next to scope, and selecting the new scope as shown in [Figure 3-47](#).




Figure 3-47 Changing the scope of the Policy blade in the portal

The non-compliance shown in Figure 3-46 is based on policies implemented by Azure Security Center. By clicking on the non-compliant item, you can see the full details of what is and isn't within policy as shown in Figure 3-48.




Figure 3-48 Details on compliance

Notice that the title of this item is ASC Default followed by a subscription ID. ASC Default is actually a collection of multiple policies that are defined by Azure Security Center. Azure Policy makes it easy to impose a full suite of policies by combining them into a group called an *initiative*. By defining an initiative, you can easily define complex rules that ensure governance of your company's policies.

You can assign a new policy either by selecting a policy from a list of included policies, or by creating your own policy. To assign a policy from the list of included policies, click on **Assignments** and then click **Assign Policy**, as shown in [Figure 3-49](#).




Figure 3-49 Assigning a policy

Click the ellipses next to Policy Definition as shown in [Figure 3-50](#) to select a policy.

Assign policy

SCOPE

* Scope ([Learn more about setting the scope](#))

Jim's Personal Azure Account

Exclusions

Optionally select resources to exempt from the policy assignment

BASICS

* Policy definition

Figure 3-50 Selecting a policy definition

In this case, you apply a policy that will ensure that any App Service app that gets created has diagnostic logging enabled. You can do that by entering **app service** in the search box, and selecting the built-in policy that applies to that policy, as shown in [Figure 3-51](#).




Figure 3-51 Adding a built-in policy definition

As shown in [Figure 3-52](#), the effect of this particular policy is Audit, meaning that if the policy isn't in compliance, a warning will appear in the portal, but it will still allow the resource to be created.

* Assignment name ⓘ

Audit enabling of diagnostic logs in App Services ✓

Description

Assigned by

Jim Cheshire

PARAMETERS

* Effect ⓘ

Audit ▼

MANAGED IDENTITY

Policies with effect type deploy/IfNotExist need the ability to deploy resources. To do this, a managed identity will be created to deploy the resources for you.
[Learn more about Managed Identity.](#)

☐ Create a Managed Identity

* Managed Identity location

East US ▼

Figure 3-52 Completing the assignment

There are six different effects supported in Azure

policy. Not all effects, however, are available for built-in policies. The effects are:

- **Append** Adds additional properties to a resource. It can be used to add a tag with a specific value to resources.
- **Audit** Logs a warning if the policy is not complied with.
- **AuditIfNotExists** Allows you to specify an additional resource type that must exist along with the resource being created or updated. If that resource type does not exist, a warning is logged.
- **Deny** Denies the create or update operation.
- **DeployIfNotExists** Allows you to specify an additional resource type that you want deployed, along with the resource being created or updated. If that resource type is not included, it is automatically deployed.
- **Disabled** The policy is not in effect.

More Info More On Policy Effects

For more information on policy effects, including examples of each, see: <https://docs.microsoft.com/en-us/azure/governance/policy/concepts/effects>.

In addition to using the built-in policies, you can also define your own policies by creating a custom policy definition. Custom policy definitions are ARM templates that define the policy. For more information on creating a custom policy definition, see: <https://docs.microsoft.com/en-us/azure/governance/policy/tutorials/create-custom-policy-definition>.

Role-based access control

Role-based access control (RBAC) is a generic term that refers to the concept of authorizing users to a system based on defined roles to which the user belongs. Azure implements RBAC across all Azure resources so that you can control how users and applications can interact with your Azure resources.

You might want to allow users who administer your databases to have access to databases in a particular resource group, but you don't want to allow those people to create new databases or delete existing databases. You might also want some web developers to be able to deploy new code to your web applications, but you don't want them to be able to scale the app to a higher-priced plan. These are just two examples of what you can do with RBAC in Azure.

There are four elements to RBAC.

- **Security principal** The security principal represents an identity. It can be a user, a group, an application (which is called a service principal), or a special AAD entity called a *managed identity*. A managed identity is how you authorize another Azure service to access your Azure resource.
- **Role** A role (sometimes called a role definition) is what defines how the security principal can interact with an Azure resource. For example a role might define that a security principal can read the properties of a resource but cannot create new resources or delete existing resources.
- **Scope** The scope defines the level at which the role is applied,

and it controls how much control the security principal has. For example, if the scope is to a resource group, the role defines activities that can be performed on all matching resources in the resource group.

- **Role assignments** Roles are assigned to a security principal at a particular scope, and that's what ultimately defines the level of access for the security principal.

RBAC includes many built-in roles. Three of these built-in roles apply to all Azure resources.

- **Owner** Members of this role have full access to the resources.
- **Contributor** Members of this role can create resources and manage resources, but they cannot delegate that right to anyone else.
- **Reader** Members of this role can see Azure resources, but they cannot create, delete, or manage those resources.

All of the other built-in roles are specific to certain types of Azure resources.

To give someone access to a resource using RBAC, open the resource you want to give access to in the Azure portal. Click on **Access Control (IAM)** in the portal to configure RBAC. In [Figure 3-53](#), RBAC is being configured for a web app hosted in Azure App Service. Clicking on **Add** in the Add a Role Assignment box allows you to add a role.




Figure 3-53 Configuring RBAC for a web app



Exam Tip

The scope of RBAC is defined by where the RBAC role is assigned. For example, if you open a resource group in the portal and assign an RBAC role to a user, the scope is at the resource group level. On the other hand, if you open a web app within that resource group and assign the role, the scope is to that web app only. That user will not have access to other apps in the resource group unless you apply other role assignments to the user.

RBAC roles can be scoped to the management group, subscription, resource group, or resource level.

After clicking **Add**, choose the role you want to assign. The list of roles will differ depending on what type of resource this is. Choose who or what you want to assign the role to, and then click on **Save**, as shown in [Figure 3-54](#).

Add role assignment ✕

Role ⓘ
Website Contributor ▼

Assign access to ⓘ
Azure AD user, group, or service principal ▼

Select ⓘ
Search by name or email address ✓

JC

Jim Cheshire

SE

Sales Execs

Selected members:

CC

Christine Conrad
cconrad@...onmicrosoft.com

Remove

Save Discard

Figure 3-54 Adding a role assignment

Figure 3-54 shows a list of users in the AAD, because the Assign Access To dropdown is set to AAD objects. You can see a list of other types of objects by selecting a different type. For example, in Figure 3-55, we are selecting a built-in managed identity type that will add

Azure VMs to the Website Contributor role for this web app.




Figure 3-55 Assigning a managed identity to a role



Exam Tip

It's important to understand that role assignments are additive. Your RBAC abilities at any particular scope are the result of all role assignment up to that level. In other words, if I have the Owner role on a resource group and you assign me the Website Contributor

role on a web app within that resource group, the Website Contributor assignment will have no effect because I already have the Owner role on the entire resource group.

RBAC is enforced by Azure Resource Manager (ARM). When you attempt to interact with an Azure resource, whether in the Azure portal or by using a command line tool, you are authenticated by ARM and a token is generated for you. That token is a representation of your identity and all of your role assignments, and it's included with all operations you perform on the resource. ARM is able to determine if the action you are performing is allowed by the roles to which you are assigned. If it is, the call succeeds. If not, you are denied access.

You can ensure that someone has the rights you desire by checking access in the Azure portal. Open the resource and click on **Access Control (IAM)**. Click the **Check Access** tab and search for the user or resource you've granted access to as shown in [Figure 3-56](#).




Figure 3-56 Checking access

Click on the user or other object and the access to the resource will be displayed as shown in [Figure 3-57](#).




Figure 3-57 Viewing role assignments for a user

For a greater level of detail on what exact operations are and aren't allowed, click on the role that's displayed. This will allow you to see a detailed list of operations and the combination of read, write, delete, and other actions that a security principal can perform as shown in [Figure 3-58](#).

Microsoft Web Apps				
Permissions - Website Contributor (preview)				
 The permissions for Website Contributor refer to one or more resource types and/or actions that are not recognized by Microsoft Web Apps. Click to learn more.				
RESOURCE TYPE (MANAGEMENT)	READ	WRITE	DELETE	OTHER ACTIO...
▼ Microsoft Web Apps				
▼ Web App	✓	✓	✓	✓
Custom Hostname	✓			
▼ Function App				✓
Web Apps Hostruntime Functions Keys	✓			
▼ Web Apps Hostruntime Host	✓			
Function App	✓			
Recommendation	✓			✓
▼ Web App	✓	✓	✓	✓
Web Apps Config Snapshots	✓			
Web App	✓	✓	✓	

Figure 3-58 Detailed permissions applied to a user for a web app

Locks

RBAC is a great way to control access to an Azure resource, but in cases where you just want to prevent changes to a resource, or prevent that resource from being deleted, locks are a simpler solution. Unlike

RBAC, locks apply to everyone with access to the resource.



Exam Tip

In order to create a lock, you must either be in the Owner or the User Access Administrator role in RBAC. Alternatively, an administrator can create a custom role that grants the right to create a lock.

Locks can be applied at the resource level, the resource group level, or at the subscription level. To apply a lock to a resource, open the resource in the Azure portal and click **Locks** in the Settings section of the menu on the left, as shown in [Figure 3-59](#).



Figure 3-59 Locking a resource

To add a lock to the resource, click **Add**. (You can also review and add Locks to the resource group by clicking on **Resource Group**, or to the subscription by clicking on **Subscription**.) Provide a name for the lock, set the lock type, and add an optional note as shown in [Figure 3-60](#).




Figure 3-60 Adding a Read-Only Lock

A read-only lock is the most restrictive lock. It prevents changing properties of the resource or deleting the resource. A delete lock prevents the resource from being deleted, but properties can still be changed. The result of a read-only lock is often unpredictable, because of the way that locks are handled by Azure.

Locks only apply to operations that are handled by ARM, and some operations specific to a resource are handled internally by the resource instead of being handled by ARM. For example, the read-only lock

applied to the Azure Key Vault in [Figure 3-60](#) will prevent a user from changing the access policies on the vault, but users can still add and delete keys, secrets, and certificates because those operations are handled internally by Key Vault.

There are other situations where a read-only lock can prevent operations that occur unexpectedly. For example, if you place a read-only lock on a storage account, it will prevent all users from listing the access keys for the storage account, because the operation to list keys makes the keys available for write access.

If a lock is applied to a resource group, all resources in that resource group inherit the lock. Similarly, if a lock is applied at the subscription level, all resources in the subscription inherit the lock. It is possible to nest locks, and in such situations, the most restrictive lock is the effective lock. For example, if you have a read-only lock on a resource group and a delete lock on a resource in that resource group, the resource will actually have a read-only lock applied to it. The explicit delete lock is ineffective.



Exam Tip

Locks are also inherited by newly-created resources. If you apply a delete

lock to a resource group, and add a new resource to the resource group at a later time, the new resource will automatically inherit the delete lock.

When an operation is attempted in the portal and denied because of a lock, an error will display as shown in [Figure 3-61](#). Notice that there isn't any information about why the operation failed, so unless you know that a lock is applied on a resource, you might find this failure confusing.





Figure 3-61 Denied by a lock

Attempting to delete the Key Vault from the Azure command-line interface (CLI), however, clearly shows that the delete failed because of a lock as shown in [Figure 3-62](#).

A terminal window titled 'jimcheshire - bash - 112x50' on a Mac. The user runs the command 'az keyvault delete --name AZ900Vault'. The output shows an error: 'The scope '/subscriptions/.../resourceGroups/SecurityRG/providers/Microsoft.KeyVault/vaults/AZ900Vault' cannot perform delete operation because following scope(s) are locked: '/subscriptions/.../resourceGroups/SecurityRG/providers/Microsoft.KeyVault/vaults/AZ900Vault'. Please remove the lock and try again.' The prompt returns to 'jimcheshire\$'.

```
jimcheshire$ az keyvault delete --name AZ900Vault
The scope '/subscriptions/.../resourceGroups/SecurityRG/providers/Microsoft.KeyVault/vaults/AZ900Vault' cannot perform delete operation because following scope(s) are locked: '/subscriptions/.../resourceGroups/SecurityRG/providers/Microsoft.KeyVault/vaults/AZ900Vault'.
Please remove the lock and try again.
jimcheshire$
```

Figure 3-62 A failure in Azure CLI because of a lock

Some resource types are better about displaying detailed errors in the portal, but if you're using locks on resources, and you experience an unexpected error when working with a resource, it's always a good idea to try the operation in PowerShell or the Azure CLI to see if a lock might be impacting you.

Azure Advisor

Azure Advisor is a best-practices analyzer for Azure resources. The goal of Azure Advisor is to help you ensure high-availability and performance, control costs, and to secure your Azure resources. Security assistance

in Azure Advisor is fed directly from Azure Security Center, but Azure Advisor provides a single view of all recommendations, including the ability to take action on recommendation directly in the Azure Advisor blade.

To access Azure Advisor, click **Advisor** in the menu in the Azure portal, as shown in Figure 3-63.

Microsoft Azure

Home > Advisor

Advisor

Search (Ctrl+F)

Feedback Download as CSV Download as PDF

Subscriptions: Jim's Personal Azure Account - Don't see a subscription? [Open Directory](#) + [Subscription settings](#)

Active

All types

High Availability	Security
3 Recommendations	11 Recommendations
0 High impact 2 Medium impact 1 Low impact	11 High impact 0 Medium impact 0 Low impact
2 Impacted resources	9 Impacted resources

Performance	Cost
You are following all of our performance recommendations	You are following all of our cost recommendations

Is Advisor helpful?

Left sidebar menu:

- Create a resource
- Home
- Dashboard
- All services
- FAVORITES
- All resources
- Resource groups
- App Services
- Function Apps
- SQL databases
- Azure Cosmos DB
- Virtual machines
- Load balancers
- Storage accounts
- Virtual networks
- Azure Active Directory
- Monitor
- Advisor**
- Security Center
- Cost Management + Billing
- Help + support
- App Service plans

Figure 3-63 Azure Advisor

Clicking on the Security tile takes you to a view of all security recommendations as shown in [Figure 3-64](#).




Figure 3-64 Security recommendations from Azure Advisor

From this blade, I can open Security Center to view details on these recommendations, but I can also click on **Follow Security Center Recommendations** to view and take action directly from within Azure Advisor. When this is clicked, you are taken to a comprehensive list, as shown in [Figure 3-65](#). Click on any of these recommendations for more details and to take action on

the recommendation.



Figure 3-65 Viewing security recommendations

In Figure 3-66, an individual recommendation is shown. You can also see the steps to take to address the recommendation, and a list of Azure resources affected by the recommendation.

Home > Advisor - Security > Recommendations > Provision an Azure AD administrator for SQL server (Preview)

Provision an Azure AD administrator for SQL server (Preview) ✕

^ General Information

Recommendation score ⓘ 0/20

Recommendation impact +20

User impact Moderate

Implementation cost Moderate

^ Threats

- Malicious insider
- Account breach

^ Remediation steps

To provision an Azure AD administrator for SQL server, see [Configure and manage Azure Active Directory authentication with SQL Database, Managed Instance, or SQL Data Warehouse](#)

Unhealthy resources

1

Healthy resources

0

LEARN MORE

[Learn more about recommendations](#) 🔗

Unhealthy resources (1) [Healthy resources \(0\)](#) [Unscanned resources \(0\)](#)

NAME	SUBSCRIPTION
 jwc	Jim's Personal Azure Account

Figure 3-66 Viewing a specific recommendation

SKILL 3.5: UNDERSTAND MONITORING

AND REPORTING OPTIONS IN AZURE

Whether you have a thousand Azure resources or just a few, being able to monitor resource usage in Azure is important. It's also important to understand the health of your Azure resources and whether a problem in your configuration or a problem in Azure itself is impacting the health of your cloud applications. Azure Monitor and Azure Service Health are two Azure services that provide these features.

This section covers:

- Azure Monitor
- Azure Service Health

Azure Monitor

Azure Monitor aggregates metrics for Azure services and exposes them in a single interface. You can also create alerts that will notify you, or someone else, when there are concerns you might want to address.

To access Azure Monitor, click on **Monitor** in the Azure portal to display the Azure Monitor blade, as shown in [Figure 3-67](#). Azure Monitor is customizable, so you can see exactly what interests you the most. For that reason, it doesn't show any metrics until you configure them. To view metrics, click on **Metrics** and

then **Select A Resource.**

Microsoft Azure

JKM'S DIRECTORY

Create a resource

Home

Dashboard

All services

FAVORITES

All resources

Resource groups

App Services

Function Apps

SQL databases

Azure Cosmos DB

Virtual machines

Load balancers

Storage accounts

Virtual networks

Azure Active Directory

Monitor

Advisor

Security Center

Cost Management + Billing

Help + support

App Service plans

Home > Monitor > Metrics

Monitor - Metrics

Microsoft

Search (Ctrl+F)

Overview

Activity log

Alerts

Metrics

Logs

Service Health

Insights

Applications

Virtual Machines (preview)

Containers

Network

More

Settings

Diagnostics settings

Autoscale

Support + Troubleshooting

Usage and estimated costs

Advisor recommendations

New support request

+ New chart

Refresh

Export to Excel

Last 24 hours (Automatic)

Chart Title

Add metric

Add filter

Apply splitting

Line chart

New alert rule

Pin to dashboard

RESOURCE	METRIC_NAMESPACE	METRIC	AGGREGATION
+ Select a resource	Select namespace	Select metric	Select value(s)

100

90

80

70

60

50

40

30

20

10

0

Select a metric above to see data appear on this chart or learn more below:



Filter + Split

Apply filters and splits to identify outlying segments



Plot multiple metrics

Create charts with multiple metrics and resources



Build custom dashboards

Pin charts to your dashboards

12 PM

06 PM

Mar 17

06 AM

Figure 3-67 Azure Monitor

In Figure 3-68, we've selected a VM in the SecurityRG resource group so that you can view metrics from this VM in Monitor.

Select a resource

Browse

Recent

Subscriptions:

Jim's Personal Azure Account – Don't see a subscription?

Open Directory + Subscription settings

Subscription

All subscriptions

Resource group

SecurityRG

Resource type

All types

Search to filter items...

RESOURCE	RESOURCE GROUP
 az900	SecurityRG
 AZ900Vault	SecurityRG
 jwc/az900db	SecurityRG
 jwc/master	SecurityRG
 Server1	SecurityRG
 server1430	SecurityRG
 ServicePlan256ea972-862f	SecurityRG

Apply

Cancel

Figure 3-68 Selecting a resource to monitor

Once you select a resource, you are presented with a list of metrics related to that resource. Resources for VMs are shown in [Figure 3-69](#).




Figure 3-69 Metrics for VMs

As soon as you select a metric, the chart updates to show a graph of that metric. You can add additional metrics to your chart by clicking on **Add Metric**, as shown in [Figure 3-70](#).



Figure 3-70 Monitoring VM CPU usage

When adding multiple metrics, you'll want to include only those metrics that share a common metric. For example, if you were to add the Disk Read Bytes metric to the chart shown in [Figure 3-70](#), it wouldn't make a lot of sense because Disk Read Bytes is measured in bytes,

and Percentage CPU is measured as a percentage.

In Figure 3-71, we've added Disk Read Bytes and Disk Write Bytes to a chart. Azure Monitor color codes each metric automatically to distinguish between them. We've also selected **Area Chart** as the type of chart to more clearly see the patterns.




Figure 3-71 Chart showing disk usage

By default, charts are shown for the past 24-hour period, and the real-time value is shown at the right edge of the chart. However, you can customize the timeframe that is shown by clicking on the timeframe and adjusting it to what you want to see, as shown in Figure 3-72.



Figure 3-72 Changing the chart timeframe

Once you have a chart that you find useful, you can pin that chart to the portal dashboard by clicking on **Pin To Dashboard**. As shown in Figure 3-73, you can pin it to the current dashboard, or you can pin it to a specific dashboard to create a monitoring dashboard in the portal customized for a specific use.




Figure 3-73 Pinning a chart

Azure Monitor Alerts can notify you or others with email or SMS text message, run a Logic App flow, call a Function App, make a request to a webhook, and more, when a certain condition is met. Alerts are based on rules that you define, and when a rule's condition is

met, an alert performs the action you specify.

You can create an alert rule that is automatically configured for the metrics you've selected in your chart by clicking on **New Alert Rule** at the top of your chart. You can also start from scratch by clicking on **Alerts** in the menu for Azure Monitor, as shown in [Figure 3-74](#), and then clicking on **New Alert Rule**.




Figure 3-74 Creating an alert rule

To start your rule, click on **Select** and select the

resource you want to configure an alert for. In [Figure 3-75](#), we've selected my VM for a new alert rule.




Figure 3-75 Selecting a resource for an alert

Next, you'll need to specify the condition for your alert. Click on **Add Condition**, and then select the signal you want to monitor for your alert. In [Figure 3-76](#), we configure an alert based on the Percentage CPU signal of the VM.




Figure 3-76 Configuring a condition

Once you select a signal, the logic for the signal is configured. As shown in [Figure 3-77](#), Monitor displays an interactive graph of the signal you've chosen, so you can get a feel for how your resource has been

performing historically. This shows the last four hours by default, although you can adjust the chart period. You can specify an operator, aggregation type, threshold, and click **Done** to create the logic for the alert.




Figure 3-77 Alert rule logic

Note Multiple Conditions

An alert rule can consist of multiple conditions. For example, you can have a rule that only triggers if CPU averages above 70% and disk usage is also high. The choice is yours.

When an alert is triggered, it performs an action that you specify using an *action group*. An action group contains a list of actions to take when an alert is triggered. To create a new action group, click on **Create New**, as shown in [Figure 3-78](#).

Create rule

Rules management

 * RESOURCE

HIERARCHY

 Server1

 Jim's Personal Azure Account

>

 SecurityRG

Select

 * CONDITION

Monthly cost in USD (Estimated) ⓘ

 Whenever the Percentage CPU is Greater than 70 %

\$ 0.10



Total \$ 0.10

Add condition

 * ACTION GROUPS

Notify your team via email and text messages or automate actions using webhooks, runbooks, functions, logic apps or integrating with external ITSM solutions. Learn more [here](#)

ACTION GROUP NAME	ACTION GROUP TYPE
No action group selected	

Select existing

Create New

Figure 3-78 Creating an action group

In Figure 3-79, we are creating an action to notify the IT director. In this case, the action will send a text message to the IT director, and it will also send a push

notification using the Azure Mobile app.

Add action group X Email/SMS/Push/Voice □ X

* Action group name: Alert all IT on VM Problem ✓

* Short name: AlertIT ✓

* Subscription: Jim's Personal Azure Account ▼

* Resource group: Default-ActivityLogAlerts (to be created) ▼

Actions

ACTION NAME	ACTION TYPE	STATUS	DETAILS	ACTIONS
Notify IT Director ✓	Email/SMS/Push/V... ▼		Edit details	X

Please configure the action by clicking the link.

Unique name for the ac... ▼

[Privacy Statement](#)

[Pricing](#)

Name

Notify IT Director

☐ Email

email@example.com

☐ Email Azure Resource Manager Role

None ▼

☒ SMS

Country code: 1 ▼ Phone number: 2145551234 ✓

[Carrier charges may apply.](#)

☒ Azure app Push Notifications

[Learn about the connecting to your Azure resources using the Azure app.](#)

itdir@kontoso.com ✓

This is the email you use to log into your Azure account.

☐ Voice

Country code: 1 ▼ Phone number: 1234567890

OK

Figure 3-79 Creating an action

Action groups are designed to contain several actions that are executed by an alert being triggered. In [Figure](#)

3-80, we've added an additional action to the action group. This action calls a Function App that runs some code to reboot the VM.

Home > Monitor - Alerts > Create rule > Add action group > Azure Function

Add action group

* Action group name: Alert all IT on VM Problem ✓

* Short name: AlertIT ✓

* Subscription: Jim's Personal Azure Account ▼

* Resource group: Default-ActivityLogAlerts (to be created) ▼

Actions

ACTION NAME	ACTION TYPE	STATUS	DETAILS	ACTIONS
Notify IT Director	Email/SMS/Push/V...		Edit details	X
Reboot Server	Azure Function	v	Edit details	X

Please configure the action by clicking the link.

Unique name for the ac... ▼

[Privacy Statement](#)

[Pricing](#)

Azure Function

Subscription: Jim's Personal Azure Account ▼

Resource group: SecurityRG ▼

Function App: [How to create a Function app](#)

NOTE: Function apps with App Service Authentication enabled are not supported.

ManageVMFunc. ▼

Function: [How to create a Function](#)

RebootVM ▼

Figure 3-80 Adding another action

Azure Service Health

Microsoft operates an Azure Status web page where you can view the current status of Azure services in all

regions where Azure operates. While it is a helpful view of overall Azure health, the enormous scope of the web page doesn't make it the most effective way to get an overview of the health of your specific services. Azure Service Health can provide you with a view specific to your resources.

To access Service Health, click **All Services**, and then click on **Service Health** as shown in [Figure 3-81](#).




Figure 3-81 Azure Service Health

Figure 3-82 shows the Service Health blade showing the health and status of my resources. The map shown has two green dots on it representing the health of the two Azure regions where resources are deployed. This map is specific, and by clicking on **Pin Filtered World**

Map to Dashboard, you can have a quick reference of Azure health for just the regions where you have resources.




Figure 3-82 Service issues in Service Health

You can also view any upcoming planned maintenance that might impact you by clicking on Planned Maintenance. By clicking on Health Advisories, you can see health information that might be related to your own configuration and not a problem somewhere

in Azure.

When a service issue is impacting you, you'll see details on the issue as shown in [Figure 3-83](#). In addition to full details on the incident, you also have a link that refers to details on the incident. You can also download a PDF that contains an official Microsoft notice of the incident.

ISSUE NAME	TRACK...	EVENT T...	SERVICE...	REGION...	START TIME	UPDAT...
Service Management Errors - App Servi...	X7YK-B_G	Incident	App Service	Central US...	12:50 UTC, 03/03/2019	2 wk ago

[Summary](#)
[Issue updates](#)

Tracking ID

X7YK-B_G

Share the below link with your team or use it for reference in your problem management system

https://app.azure.com/h/X7YK-B_G/2ed188


Impacted service(s)

App Service

Impacted region(s)

Central US, South Central US

Last update (2 wk ago)

SUMMARY OF IMPACT: Between 12:50 and 16:45 UTC on 03 Mar 2019, you were identified as a customer using App Service who may have experienced periods of increased latency or received failure notifications when performing service management operations - such as create, update, or delete.

PRELIMINARY ROOT CAUSE: Engineers determined that a backend database supporting the App Service platform experienced a period of high CPU utilization, preventing service management requests from completing.

 [Download the issue summary as a PDF.](#)

 Track this issue on mobile:



 Quickly connect with our problem-solving experts.

[Tweet @AzureSupport](#)

Figure 3-83 Azure Service Health incident

Both Azure Monitor and Azure Service Health are critical to the overall view of your Azure resources. Azure Monitor is geared toward monitoring the cost and performance of your resources and alerting you and

others when conditions warrant. Azure Service Health, on the other hand, is the single-point-of-truth for information on the health of Azure itself and how Azure incidents are impacting your resources. The combination of these two services provides you with all of the tools you need to keep up with your Azure resources and how well they're performing.

SKILL 3.6: UNDERSTAND PRIVACY, COMPLIANCE, AND DATA PROTECTION STANDARDS IN AZURE

As you move to the cloud, you offload some of the responsibility for your services and data to your cloud provider. This includes some of the responsibility for compliance with data protection standards. Even though the cloud provider takes care of some of that burden for you, it's still vital that you have confidence in the cloud provider and that you trust them to maintain compliance.

There are many standards that business must comply with. For example, in 2016 the European Union passed the General Data Protection Regulation, or GDPR. GDPR regulates the way personal data is handled for individuals with the EU, but it also controls any personal data that is exported out of the EU. Companies doing business in EU countries are legally required to abide by the GDPR.

One way that organizations can ensure they are abiding by the GDPR and other regulations that regulate data, is to maintain compliance with industry-wide standards focused on helping organizations keep information secure. One of those standards is the International Organization of Standards (ISO) 27001 standard. Companies that comply with the ISO 27001 standard can be confident that they are maintaining the best practices necessary to keep information secure. In fact, many companies won't do business with a cloud provider unless they can prove ISO 27001 compliance.

Systems that deal with governmental data must maintain compliance with standards that are maintained by the National Institute of Standards and Technology, or NIST. The NIST SP 800-53 is a publication by NIST that outlines all the requirements for information systems dealing with government data. In order for any government agency to use a service, it must first prove compliance with NIST SP 800-53.

Microsoft addresses these compliance requirements across its infrastructure in many different ways.

This section covers:

- Microsoft Privacy Statement
- Trust Center
- Service Trust Portal
- Compliance Manager

- Azure Government
- Azure Germany

Microsoft Privacy Statement

The Microsoft privacy statement is a comprehensive statement from Microsoft that outlines the following as it relates to handling data and your personal information.

- Personal data Microsoft collects
- How Microsoft uses personal data
- Reasons Microsoft shares personal data
- How to access and control your personal data collected by Microsoft
- How Microsoft uses cookies and similar technologies
- What organizations providing Microsoft software to you can do with your data
- What data is shared when you use a Microsoft Account with a third-party
- Specifics about how Microsoft secures data, where it's processed, and retention policies

Microsoft links to the privacy statement in all official communications, and you can access the privacy statement online at: <https://aka.ms/privacystatement>.

Trust Center

The Trust Center is a web portal where you can learn all

about Microsoft's approach to security, privacy, and compliance. You can access Trust Center by browsing to:
<https://www.microsoft.com/en-us/trustcenter/default.aspx>.

Explore the Trust Center using the dropdown menus at the top of the page as shown in Figure 3-84. As standards and compliance evolve, you can always find the latest information on the Trust Center website.



Figure 3-84 Microsoft Trust Center

Service Trust Portal

The Service Trust Portal (STP) is a portal that provides access to various compliance tools Microsoft provides for you to track compliance in your applications running on Microsoft's various platforms. You can access the STP by browsing to: [*https://aka.ms/STP*](https://aka.ms/STP). Figure 3-85 shows the STP home page.



Figure 3-85 Azure Service Trust Portal

The STP is a launching point for the following resources.

- **Compliance Manager** A tool for managing your regulatory compliance in the cloud.
- **Audit Reports** Comprehensive reports and resources that allow you to see details on how Microsoft maintains compliance.
- **Data Protection Information** Full details on how Microsoft designs its cloud offerings to ensure that customer data is protected.
- **Privacy** Information related to how Microsoft helps you maintain compliance with GDPR.

Compliance Manager

Compliance Manager is a tool within the STP that makes it easy to visualize your compliance with industry standards. Compliance Manager also provides details on how you can improve compliance, and for those areas where compliance is Microsoft's responsibility, it provides full details on how Microsoft maintains compliance.

To access Compliance Manager, click Compliance Manager at the top of the STP page. Compliance Manager allows you to track your compliance with related applications by grouping them into groups that you can give a name of your choice. Each group you create is represented by a tile in Compliance Manager, and you can see at a glance how far you've progressed at

compliance in each group as shown in [Figure 3-86](#).




Figure 3-86 Compliance Manager

You can add a new assessment by clicking on **Add Assessment**. You can add an assessment to an existing group or you can create a new group as shown in [Figure 3-87](#).




Figure 3-87 Adding a new assessment

You can choose to assess Azure, Office 365, Microsoft Dynamics, and more. You can also choose the standard you want to evaluate against. In [Figure 3-88](#), we've chosen to evaluate Azure resources against GDPR.




Figure 3-88 Evaluating for GDPR compliance

Based on the products we're evaluating, Compliance Manager knows which parts of compliance are Microsoft's responsibility, and which are the customer's responsibility. In [Figure 3-89](#), a newly-created assessment is shown in Compliance Manager. Microsoft's responsibilities are displayed at the top of the list, and they are already completed. My responsibilities appear in the list as well, and they outline all the requirements we must meet in order to

be compliant with GDPR.

< Back To Dashboard
[Export to Excel](#)

My Azure Apps Azure GDPR

Group Name	Product	Assessment	Assessed Controls	% Assessed	Status	Last Modified	Compliance Score

45/130

In Progress

3/17/2019

N/A

- Azure in-Scope Cloud Services ▼
- Microsoft Managed Controls ▼
- Customer Managed Controls ▲

Conditions for collection and processing	0/12 Assessed ▼
Data protection by design and by default	0/10 Assessed ▼
Personal data sharing, transfer, and disclosure	0/12 Assessed ▼

Feedback

Figure 3-89 A GDPR assessment in Compliance Manager

If you click on one of the Microsoft controls, you'll see details on what Microsoft has done to ensure compliance. In the control shown in [Figure 3-90](#), you can see how Microsoft complied with specific GDPR articles, when it was tested, and how it was tested.

Controls / Articles	Compliance Score	Status	Test date	Tested By	Test result
Control ID: 8.2.1	N/A	Implemented	1/26/2018	Third party independent auditor	
Control Title: Cooperation agreement					
Supported GDPR Article(s): Article (28)(3)(e), Article (28)(3)(f), Article (28)(9), Article (35)(1)					
Description: Article (28)(3)(e): Processing by a processor shall be governed by a contract or other legal act under Union or Member State law, that is binding on the processor with regard to the controller and that sets out the subject-matter and duration of the processing, the nature and purpose of the processing, the type of personal data and categories of data subjects and the obligations and rights of the controller. That contract or other legal act shall stipulate, in particular, that the processor: (e) taking into account the nature of the processing, assists the controller by appropriate technical and organizational measures, insofar as this is possible, for the fulfilment of the					
Read More					
Less 					
Microsoft Implementation Details		Test Plan Details		Management Response	
Data processing contracts between the customer and Azure specify the minimum technical and organizational measures to ensure that the contracted security arrangements are in place and that data is not processed for any purpose independent of the instructions of the controller. In the interest of transparency, Microsoft lets customers know which subcontractors are		Examined the list of Azure subcontractors available and determined that Azure is transparent about its capabilities during the process of entering into contract. Examined the Data Processing Terms of the OST as well as privacy statement and validated the controls associated with handling of personally identifiable information		N/A	
Feedback					

Figure 3-90 Microsoft compliance with GDPR articles

If you click on a customer managed control, you can see details on what you need to do to comply. In [Figure 3-91](#), details are shown regarding what you need to do to comply with GDPR Article (5)(1)(b). You have the option of assigning this task to a specific user if you want to, and you can click on Manage Documents to upload supporting documents to Compliance Manager.

Customer Managed Controls

Conditions for collection and processing

0/12 Assessed

Controls / Articles	Compliance Score	Related Controls / Articles	Assigned User	Implementation Status	Implementation Date	Test date	Test result
Control ID: 7.2.1 Control Title: Identify and document purpose Supported GDPR Article(s): Article (5)(1)(b) Description: Article (5)(1)(b): Personal data shall be: (b) collected for specified, explicit and legitimate purposes and not further processed in a manner that is incompatible with those purposes further processing for archiving purposes in the public interest, scientific or historical research purposes or statistical purposes shall, in accordance with Article 89(1), not be considered to be incompatible with the initial Read More	N/A	No related articles found	Assign Manage Documents	Select			Select

Less

Customer Actions

Customers who are controllers are responsible for collecting and processing personal data only for legitimate and explicit purposes.

Implementation Details

Enter implementation details for your organization, along with any notes you want to include. Information that you enter in this field can help others in your organization, as

Test Plan & Management Response

Enter test plan information in this field to track how your organization validates the implementation details. You can a responses from your organization

Feedback

Figure 3-91 My compliance requirements

Using this method of assessment, you can always determine if your applications are compliant. The burden of understanding certain standards and regulations is removed because Compliance Manager manages that for you.

Azure Government

Compliance with GDPR and the other standards mentioned relates to the privacy of an individual's data. Some US government compliance scenarios require that data stays within the United States of America and that only citizens of the United States have any access to systems used to store that data. Complying with these requirements is impossible with the public cloud, so Microsoft has developed completely isolated Azure data centers that make up the Azure Government cloud.

Azure Government data centers are separate from public data centers. All employees working in Azure Government are screened and are citizens of the US. Even Microsoft employees who provide technical support to Azure Government customers are required to be US citizens.

Because Microsoft also wanted to allow for compliant communication between the Azure Government cloud and on-premises government systems, they also developed dedicated Microsoft ExpressRoute locations that are completely isolated from other Azure networks

and that use their own dedicated fiber optic components.

Azure Government isn't only for federal government agencies. Cities and municipalities also take advantage of Azure Government for compliance. When a customer signs up for Azure Government, Microsoft vets that user to ensure they are representative of a government agency. Only then are they given a subscription to Azure Government.

The Azure Government cloud has all of the same features and services as the public cloud, but there are small differences. For example, the portal for Azure Government is located at <https://portal.azure.us> instead of <https://portal.azure.com>. URLs for Azure services also use the .us top-level domain, so if you create an App Service web app in Azure Government, your default domain name is <https://webapp.azurewebsites.us>. However, outside of that difference, everything else is the same, so developers who have a skill set in cloud development in Azure will find that their skills transfer directly to Azure Government.

The United States Department of Defense has additional compliance requirements called DoD Impact Level 5 Provisional Authorization. Compliance with this relates to controlled unclassified information that

requires additional levels of protection. These additional DoD requirements are met by a subset of data centers within Azure Government that are approved for DoD usage.

Azure Germany

Much like Azure Government, Azure Germany is a distinct cloud system that's designed to meet specific compliance needs. In the case of Azure Germany, those needs relate to the strict requirements imposed by the EU. Azure Germany is available to customers doing business in the EU, the European Free Trade Association, and the UK.

Azure Germany datacenters are physically located in Germany and are operated under strict security measures by a local company named T-Systems International (a subsidiary of Deutsche Telekom) that operates as a data trustee. The data trustee has full control over all data stored in Azure Germany and all of the infrastructure used to house that data. Microsoft is involved in managing only those systems that have no access at all to customer data.

THOUGHT EXPERIMENT

Let's apply the concepts related to security, privacy, compliance, and trust that you've learned in this chapter to a thought experiment. You can find the answer to this thought experiment in the section that follows it.

ContosoPharm has secured a contract to do some pharmaceutical development for the US Department of Defense. Because this work will be extremely sensitive in nature, there are several requirements they must meet before the DoD will sign off on the project.

Access to the application from DoD personnel will be via a web interface over devices that are connected to the Internet in the field. The DoS requires strict security controls for any access to the application, and it's important that the application be protected from malicious attacks that might look like legitimate traffic. This requirement applies to both cloud and on-premises VMs.

It's also required that the application configuration protects against situations where someone's password is stolen. If someone outside of the DoD attempts to log into the application using a stolen password, that login must be denied.

The application heavily uses Azure SQL Database, and users of the application should be able to create their own database tables. They should not, however, be able to give anyone access to the database that isn't explicitly given permission by the administrator. In addition to that, users of the database portion of the application are on a rotational schedule, so it must be very easy to revoke the access without having to change any

database passwords.

Some of the systems will run on-premises because of access to confidential information. The administrator of the system would like a single interface where the security status of all cloud resources and on-premises resources can be monitored. As new systems are brought online, the administrator must be able to ensure that DoD requirements for antivirus software are met.

One VM in the application will need to be remotely managed, but a requirement of the DoD is that remote management ports on the machine cannot be left open. Therefore, they need you to suggest a way they can securely remotely manage the VM.

The caching component in the application contains sensitive data that is protected by an encrypted RSA key. The DoD requires that the key necessary to access the cached data not be stored anywhere within the application. It's required that access to the key be given to the app when necessary, but only when necessary.

During the development of the application, information will be shared in Office documents and in emails. Some of the people who will receive this information aren't DoD employees, so some of this information will be sent outside of the DoD. Even so, it's a requirement that no one other than the recipient

of the data be able to read it, and they shouldn't be able to share it with others.

The DoD IT department has operational parameters that all VMs must meet for availability. These metrics revolve around disk usage, CPU utilization, and memory usage. When metrics fall outside of normal operation, it's critical that actions be taken immediately, so the IT department has developed some scripts that they can run quickly to address that. If possible, they would like to automate this process so that the system will take care of itself if no one is available to address a problem immediately.

Finally, the application must comply with DoD Impact Level 5 standards, and privacy of each individual's data must also be maintained.

THOUGHT EXPERIMENT ANSWERS

This section covers the answers to the thought experiment.

To protect the VMs in the cloud from malicious traffic, Azure Firewall can be deployed on the network. This can help to protect from traffic that otherwise looks legitimate because of Azure Firewall's ability to remember the state of a connection. To apply that same level of protection to on-premises machines, the DoD can use Advanced Threat Protection. By installing ATP sensors on-premises, the DoD can ensure that these on-

premises resources are protected.

In order to prevent someone from logging into the system with a stolen password, the DoD should implement multi-factor authentication. MFA will require that users not only have a password, but also that they have access to an approved device associated with them.

In order to give users the ability to create database tables, but not give anyone else access, DoD should create users in their Azure Active Directory and assign RBAC roles to the SQL databases. Because users with RBAC access aren't using a password associated with the database to read data, access can quickly be revoked when necessary without having to change a password in the database.

Azure Security Center can provide insight into the security state of resources in a single interface. Security Center will also show non-compliance for VMs where endpoint protection is not installed, and the DoD will be able to easily install anti-virus on one or more VMs with the click of a button.

For the one VM that must be remotely managed, DoD can enable JIT access in Security Center. This allows them to leave the management ports closed until a user requests access, and the user can even specify their specific IP address for access which will further protect

the VM. Once the access time is elapsed, the management ports are closed.

To protect the RSA key used with their caching component, the DoD can use Azure Key Vault. They can use Key Vault to generate the key as well if they want to. When the key is needed by the application, it can access it via a secure URL.

To protect information in emails and documents, including those sent outside of the DoD, Azure Information Protection can be used. By classifying emails as Confidential and for recipients only, access to the information in the mail will be protected. These same protections can be applied to their Office documents.

Azure Monitor can be used to monitor VMs and ensure they are operating within operational parameters. An alert group can be created that will trigger when desired, and because an alert can call a webhook, a Function App, or start a Logic App flow, there are many options for automating the scripts they use to address problems, even if no one's around to see the problem immediately.

To comply with DoD Impact Level 5 standards, the application will need to be hosted in a DoD-approved datacenter within Azure Government. The DoD can also use Compliance Manager to ensure data privacy

standards are met for the application.

CHAPTER SUMMARY

Security, compliance, privacy, and trust are cornerstones of Microsoft Azure, and Microsoft has proven their commitment to these principles by providing industry-leading tools and services to help you. In this chapter, you've learned about many of these tools and services.

Here's a summary of what was covered in this chapter:

- The primary attack vector for cloud applications is the network, and Azure Firewall is a stateful firewall that can protect your network from attacks.
- All traffic to the firewall is blocked by default and rules are configured to allow certain traffic to pass.
- A route table is used to direct traffic into your firewall's subnet.
- DDoS Basic protection in Azure Firewall protects from common network attacks.
- DDoS Standard protection is available for an extra charge and it will protect from additional attacks using Advanced Threat Protection.
- Network Security Groups (NSGs) can be used to control which subnets and resources can talk to each other in a virtual network.
- Service tags can be used to allow Azure services or the Internet by a NSG.
- Azure Active Directory is a cloud-based identify service in Azure that authenticate and authorizes users.

- Enterprise applications in Azure AD allow you to integrate third-parties with Azure AD so users can experience a single-sign-on experience.
- Multi-factor authentication in Azure AD requires that users have both a password and an owned device in order to log in.
- Azure Security Center provides a single portal for monitoring and managing the security of Azure resources and on-premises resources.
- Just in time VM access in Security Center makes it easy to control when and for how long management ports are open on VMs.
- Azure Key Vault provides a secure way to store secrets, keys, and certificates.
- Azure Information Protection helps you to categorize emails and documents and protect them from being accessed by unauthorized people.
- Advanced Threat Protection makes it possible to protect on-premises domain controllers and servers from attacks.
- Azure Policy allows you to define rules that are applied with Azure resources are created and managed.
- Role-based access control makes it possible to give users and applications access to your Azure resources and to control what they can and can't do.
- Locks allow you to lock down properties that go through ARM from being changed on a resource or to prevent a resource from being deleted.
- Azure Advisor provides a portal for analyzing and reporting on best-practices related to your Azure resources.
- Azure Monitor can display charts with data metrics for your Azure resources.
- Azure Monitor alerts can notify you based on conditions. They

can also call a webhook, run a Function App, start a Logic App flow, and more.

- Azure Service Health provides an overview of the health of Azure and your Azure resources that is scoped to only those regions where you have resources.
- The Microsoft Privacy Statement is Microsoft's promise to customers related to how it will protect personal data.
- Trust Center is a portal where you can learn about Microsoft's approach to security, privacy, and compliance.
- The Service Trust Portal provides access to compliance tools and information on compliance.
- Compliance Manager is part of Service Trust Portal and makes it easy to manage compliance with industry standard regulations using assessments.
- Azure Government is a private cloud for governments that is operated with distinct datacenters within the United States. All employees are screened and are US citizens.
- Azure Government is aimed at ensuring compliancy with government standards.
- DoD datacenters within Azure Government provide stricter control to comply with DoD standards.
- Azure Germany is a private cloud operated out of Germany that is designed to comply with strict EU guidelines.
- Individual data in Azure Germany is controlled and accessible only by a data trustee. Microsoft has no access to any system that touches customer data.

Chapter 4. Understand Azure pricing and support

Although we've covered many topics in this book, we haven't examined the primary concerns when moving to the cloud: pricing and support.

Pricing doesn't just involve knowing the price of Azure resources. Companies often want to know how much cloud resources are going to cost before applications are deployed to the cloud, and once the application is deployed, they want to minimize costs as much as possible and have visibility into the costs of Azure resources.

Support is also critical in a cloud environment. As we've learned, when you move to the cloud, at least some portion of infrastructure management transitions to the cloud provider. When something goes wrong, it's critical that you get the support you need to maintain the availability of your applications. It's also important to understand what level of support is offered for specific services, especially services that might be in preview and not officially released.

In this chapter, we'll examine all of these aspects

related to Azure. We'll cover your Azure subscription, how to plan and manage costs, the support options available to you, Azure service level agreements, and the release cycle for Azure services.

Skills covered in this chapter:

- Understand Azure subscriptions
- Understand planning and management of costs
- Understand the support options available in Azure
- Describe Azure service level agreements
- Understand service lifecycle in Azure

SKILL 4.1: UNDERSTAND AZURE SUBSCRIPTIONS

You get an Azure subscription automatically when you sign up for Azure and all of the resources you create are created inside that subscription. You can, however, create additional subscriptions that are tied to your Azure account. Additional subscriptions are useful in cases where you want to have some logical groupings for Azure resources, or if you want to be able to report on resources used by specific groups of people.

This section covers:

- Azure subscription
- Uses and options with Azure subscriptions

Azure subscription

Like any other Azure resource, you can manage your subscription in the Azure portal. You can view and manage costs, you can give other people access to it using RBAC, you can apply locks to it, and so on.

Each Azure subscription has limits (sometimes called quotas) assigned to it. For example, you can have up to 200 Azure Storage accounts per region in a subscription, up to 25,000 virtual machines per region, and up to 980 resource groups per subscription across all regions.

More Info Subscription Limits

You can find details on all limits for subscriptions at:
<https://docs.microsoft.com/azure/azure-subscription-service-limits>.



Exam Tip

Microsoft support has the ability to increase limits in some scenarios. For example, if you have a good business justification, Microsoft can increase the limit of Storage accounts to 250 per subscription, per region. Some limits, however, cannot be increased.

Figure 4-1 shows an Azure subscription in the Azure portal.

Jim's Personal Azure Account

Subscription

[Manage](#) [Transfer](#) [Cancel subscription](#) [Rename](#) [Change directory](#)

Essentials

For more cost management and optimization capabilities, try Azure Cost Management →

Overview

[Access control \(IAM\)](#)[Diagnose and solve problems](#)[Security](#)[Events](#)

Cost Management

[Cost analysis](#)[Budgets](#)[Advisor recommendations](#)

Billing

[Invoices](#)[External services](#)[Payment methods](#)[Partner information](#)

Settings

[Programmatic deployment](#)[Resource groups](#)[Resources](#)[Usage + quotas](#)

Costs

Costs by resource



Spending rate and forecast




Figure 4-1 Azure subscription in the Azure portal

On the Overview blade, you can see a cost breakdown for each of the resources. You can also see the spending rate for the subscription, along with a forecasted cost by the end of the current month. If you click on the Costs By Resource tile, you can see a further breakdown of the Azure expenses, as shown in [Figure 4-2](#). In this view, you see costs by Service Name, Location (Azure region), and Resource Group, along with a graph of the costs for the month.




Figure 4-2 Azure subscription cost analysis

More Info Creating Budgets

You can manage your costs in Azure by creating budgets. You'll learn more about that in the Azure cost management portion of Skill 4.2: "Understand planning and management of costs."

Azure invoices are also available for the subscription from within the Azure portal. You can see all of the past invoices by clicking on Invoices in the menu for the subscription, as shown in [Figure 4-3](#).

The screenshot displays the 'Invoices' section of the Azure portal for 'Jim's Personal Azure Account'. The left-hand navigation pane includes links for Overview, Access control (IAM), Diagnose and solve problems, Security, Events, Cost Management, Cost analysis, Budgets, Advisor recommendations, Billing, Invoices (selected), External services, Payment methods, and Partner information. The main content area shows a notification about the payment method (AMEX) and a table of invoices for Azure services. The table has columns for Billing Period, Charge Date, Amount (USD), and Invoice (with a download link). The subscription selected is 'Jim's Personal Azure Account (2ed06a...)'.

BILLING PERIOD	CHARGE DATE	AMOUNT (USD)	INVOICE
2/3/2019-3/2/2019	3/3/2019	176.34	Download invoice
1/3/2019-2/2/2019	2/3/2019	228.13	Download invoice
12/3/2018-1/2/2019	1/3/2019	184.94	Download invoice
11/3/2018-12/2/2018	12/3/2018	211.20	Download invoice
10/3/2018-11/2/2018	11/3/2018	168.43	Download invoice
9/3/2018-10/2/2018	10/3/2018	124.59	Download invoice

Figure 4-3 Azure invoices

Uses and options with Azure subscriptions

You can create additional Azure subscriptions in your Azure account. This is useful in cases where you want to separate costs or if you are approaching a subscription limit on a resource. To create a new Azure subscription, enter **subscription** in the search box and click on **Subscriptions** as shown in [Figure 4-4](#).




Figure 4-4 Azure subscriptions

To create a new subscription, click on **Add** in the Subscriptions blade as shown in [Figure 4-5](#).




Figure 4-5 Creating a new subscription

After you click on **Add**, you need to choose which type of subscription you want to create. There are several types of Azure subscriptions.

- **Free Trial** Provides free access to Azure resources for a limited time. Only one free trial subscription is available per account, and you cannot create a new free trial if a previous one has expired.
- **Pay-As-You-Go** You pay only for those resources you use in Azure. There's no up-front cost, and you can cancel the subscription at any time.
- **Pay-As-You-Go Dev/Test** A special subscription for subscribers to Visual Studio that can be used for development and test. This subscription offers discounted rates on VMs, but you

cannot use this for production applications.

Note Azure Subscription Types

You may have additional subscription options, depending on the type of Azure account you have.



Exam Tip

Each subscription is associated with a unique identifier called a *subscription ID*. You can give each subscription a descriptive name to help you identify it, but Azure will always use the subscription ID to identify your subscription. When you talk to Microsoft about your Azure account, they'll also often ask for your subscription ID.

SKILL 4.2: UNDERSTAND PLANNING AND MANAGEMENT OF COSTS

As you begin to contemplate moving to the cloud, the first thing you'll likely want to do is determine what your costs will be based on your resource needs. Once you've begun deploying and using Azure resources, managing your costs becomes important in order to stay

within your budgets. Azure has tools that help you with planning and the management of your costs in Azure.

This section covers:

- Options for purchasing Azure products and services
- Options around Azure free account
- Factors affecting costs
- Zones
- The pricing calculator
- The total cost of ownership (TCO) calculator
- Best practices for minimizing Azure costs
- Azure Cost Management

Options for purchasing Azure products and services

There are a couple of ways to purchase Azure products and services. You can purchase products and services directly from Microsoft, or you can purchase through a Microsoft Cloud Solution Partner (CSP).

When you purchase directly from Microsoft, you decide which Azure services you want to purchase, and you manage all of your deployments and usage of those services. Each month Microsoft will invoice you for your Azure usage, and you'll have access to those invoices in the Azure portal. If you need support for your Azure resources, you get support directly from

Microsoft using one of the available support plans for Azure.

More Info Azure Support Plans

You'll learn more about support plans in Azure in Skill 4.3: "Understand the support options available in Azure."

When you purchase from a CSP, you are not purchasing individual Azure resources. Instead, you are purchasing an entire cloud solution developed by the CSP. When you need to deploy your application, you work with the CSP to manage the deployment. The CSP also provides you information on usage costs of your resources, and if you need support, you get that support from the CSP, not from Microsoft.

When you purchase from Microsoft, you have the option to purchase directly from the portal or to have your services added to an existing Enterprise Agreement with Microsoft. Enterprise Agreements are designed for large companies that have a large amount of usage in Azure. When you sign up for an Enterprise Agreement, you work with Microsoft to come up with a yearly financial commitment for Azure usage. You are charged yearly for the committed price, and if you use more than your agreed-upon usage, you are charged for the additional usage at the rate you and Microsoft agree to.

Purchasing directly from Microsoft offers the most

flexibility and control because you decide which individual resources you are purchasing. It's also important to consider that your development team may have experience in developing against specific types of Azure resources, and by controlling the types of resources you use, you can reduce the number of potential problems in your applications.

On the other hand, if you have a need to deploy a complex cloud solution and you don't have on-premises expertise in some areas, a CSP may be the best option for you. Because the CSP is certified to have expertise in the services they offer, they may be able to provide you with a more efficient solution that can help you reduce costs and support needs.

Options around Azure free account

If you've never had an Azure free trial and you've never been a paid Azure customer, you are eligible for an Azure free account. A free account gives you 12 months of free access to the most popular Azure services, and many other Azure services offer free usage even after those 12 months have elapsed. You also get a \$200 credit that you can use for Azure services for a 30-day period after you sign up for a free account.



Exam Tip

The \$200 credit cannot be used to pay for Azure Marketplace offerings. Many Azure Marketplace offerings, however, provide their own free trials.

Microsoft places a \$200 spending limit on free accounts so that you don't accidentally go over the \$200 free credit. If you hit that spending limit, you will need to upgrade your subscription to a Pay-As-You-Go subscription in order to create additional resources.

At the end of the 30-day period, any resources that are not free for 12 months or more are deleted automatically, so you'll want to ensure you upgrade your Azure subscription before the 30 days has elapsed if you want to continue using your resources.



Exam Tip

Microsoft doesn't require that resources used with a free account be used for development or testing only. You are free to use these resources for production use.

More Info Complete List Of Free Account Product Availability

For a complete list of products included in the free account and how

long they are available for free, see:
<https://azure.microsoft.com/free/free-account-faq>.

Factors affecting costs

As you're planning your Azure deployments, you should keep in mind the factors that can impact your costs. The primary factors that impact costs are the resource type, how you purchase the resource, the Azure regions you're using, and the billing zone your resources are in.

Azure services are billed according to *meters* associated with a resource. These meters track how much a specific metric has been used by the resource. For example, there is no charge specifically for an Azure virtual network, and you aren't charged for any network traffic within a virtual network, but you are charged per gigabyte for traffic into and out of the virtual network from peered virtual networks.



Exam Tip

Each Azure service has a pricing page that outlines estimates on pricing for that resource based on typical usage.

As you determine which resources you need to use in your Azure deployment, consider how those resources

are going to use the metrics the resources charges for. For example, if you can plan your virtual networks so that you have fewer peered networks, you can save substantially over the long-term.

You may also find that purchasing Azure resources differently may offer cost savings. If you agree to pay in advance using an Enterprise Agreement, Microsoft will offer you a reduced rate. Longer term agreements offer even more price breaks. CSPs may also provide you with complete solutions that are more cost-effective than purchasing all of the resources yourself.

Microsoft's costs for operating Azure services differs by region, even when those regions are within the same geographic boundary. Therefore, your pricing will differ based upon which Azure region you use. For example, a VM deployed to the Central US region will cost more than the same VM deployed to the East US region. Microsoft doesn't provide a breakdown on their costs, but you can assume that electricity and other resources needed for an Azure data center are more expensive in the Central US region than they are in the Eastern US region.



Exam Tip

Choosing the least expensive region for

each of your Azure resources usually isn't a good way to control costs. You may end up having to pay for network traffic across regions, and that may increase your costs above the amount you're saving. Many Azure resources do not charge for network traffic within the same region, but they will charge for traffic across regions.

It's also important to keep in mind that you're not charged for network traffic into an Azure datacenter, but you are charged for network traffic out of a datacenter. However, your first 5GB of outbound data is free. After that point, you are charged a set amount on a tiered model. The amount you're charged depends on the billing zone.

More Info Pricing Of Network Bandwidth

For more information on pricing of network bandwidth in Azure, see:
<https://azure.microsoft.com/pricing/details/bandwidth>.

Zones

Azure geographies are broken out into four separate groups for billing purposes. These groups are called *billing zones*, or more commonly, simply zones. Microsoft's costs for network traffic out of each zone

differs, so your costs will differ as well.

Table 4-1 lists the zones in Azure and their corresponding geographies.

Table 4-1 Zones and geographies

Zone Name	Included Geographies
Zone 1	United States, Europe, Canada, UK, France
Zone 2	Asia Pacific, Japan, Australia, India, Korea
Zone 3	Brazil
DE Zone 1	Germany

The cheapest outbound networking costs are in Zone 1. DE Zone 1 is the second cheapest, followed by Zone 2 and Zone 3.

As you can see, there are many factors that can impact your costs in Azure, and it can be difficult to estimate costs based on all of these factors. Fortunately, Microsoft offers a pricing calculator that can help you get a handle on estimating your costs as you move to the cloud.

The pricing calculator

The Azure pricing calculator can help you get an estimate of expenses based on the products you intend on using, as well as where those products will be

deployed, and so on. You can access the pricing calculator by browsing to:

<https://azure.microsoft.com/en-us/pricing/calculator>.

The first step in calculating an estimate of your Azure expenses is to select which products you want to use. As shown in Figure 4-6, some of the more common Azure products are displayed by default, and you can add any of those products by clicking on its tile.



Pricing calculator

Configure and estimate the costs for Azure products

[Products](#)[Estimates](#)[FAQ](#)

Select a product to include it in your estimate.

**Featured**

Compute

Networking

Storage

Web

Mobile

Containers

Databases

Analytics

AI + Machine Learning

**Virtual Machines**

Provision Windows and Linux virtual machines in seconds

**Storage**

Durable, highly available, and massively scalable cloud storage

**Azure SQL Database**

Managed relational SQL Database as a service

**App Service**

Quickly create powerful cloud apps for web and mobile

**Azure Cosmos DB**

Globally distributed, multi-model database for any scale

**Azure Kubernetes Service (AKS)**

Simplify the deployment, management, and operations of Kubernetes

**Azure Functions**

Process events with serverless code

**Cognitive Services**

Add smart API capabilities to enable contextual interactions

**Cost Management**

Optimize what you spend on the cloud, while maximizing cloud potential

Figure 4-6 The pricing calculator

If the product you want is not listed, you can either click on a category of products in the list on the left or you can search for your product by entering its name in the search box.

After you add the products you want to use, scroll down to configure the specific details of each service. These details vary based upon how Microsoft charges for the product. [Figure 4-7](#) shows the options for Azure SQL Database.



Azure SQL Database

REGION:

East US ▾

TYPE:

Managed Instance ▾

BACKUP STORAGE TIER:

LRS ▾

SERVICE TIER:

General Purpose ▾

GENERATION:

Gen 4 ▾

INSTANCE:

8 vCore ▾

Clone

Delete

More info

[Pricing details](#)

[Product details](#)

[Documentation](#)

Billing Option

Save up to 73% on pay as you go prices with 1 year or 3 year reserved options.

☒ Pay as you go

☐ 1 year reserved (~21% discount)

☐ 3 year reserved (~33% discount)

Save up to 55% with [Azure Hybrid Benefit](#) for SQL Server



1 x 730
Instances Hours ▾

= \$1,472.75
Per month

Figure 4-7 Pricing options for Azure SQL Database

Clicking on Pricing Details will open the pricing page for the product in a new tab. You can also click on Product Details or Documentation to read more about the service in order to help you make better decisions about the options you select.

Once you've configured a product based on your needs, you can click the Clone button to add another instance of that product to your estimate. For example, suppose you need two Azure SQL Databases for your app, and each of them are going to be using the same service tier, instance size, and so on. The easiest way to add these is to add one Azure SQL Database product to your estimate, configure it with the desired pricing options, and then click Clone to add the second instance.

To review your pricing estimate, scroll to the bottom of the page. As shown in [Figure 4-8](#), you can choose a support plan to add to your estimate. If you have an Enterprise Agreement or Microsoft Customer Agreement, you can choose them in order to have that pricing applied to your estimate. You can then click **Export** to save your estimate as an Excel file, then select **Save** to save your estimate in the pricing calculator to make changes later on, or select **Share** to create a sharable link to your estimate so that others

can view it.

Support

SUPPORT:

Developer

\$29.00

Programs and Offers

LICENSING PROGRAM:

Microsoft Online Services Program (M

SHOW DEV/TEST PRICING

Estimated monthly cost

\$1,730.96

Export

Save

Share

US Dollar (\$)

DISPLAY SKUS

DISPLAY RESOURCE IDS

Prices are estimates and are not intended as actual price quotes. Actual prices may vary depending upon the date of purchase, currency of payment, and type of agreement you enter with Microsoft. Contact a Microsoft sales representative for additional information on pricing.

Purchase options >

Figure 4-8 Completing an estimate in pricing calculator

Note Saved Estimates

If you save an estimate in the pricing calculator, you can access it later by clicking the Estimates tab at the top of the page.

More Information Support Options

We'll cover the support options in Skill 4.3: "Understand the support options available in Azure".

The total cost of ownership (TCO) calculator

The pricing calculator is helpful for estimating your expenses for new applications in Azure, but if you have on-premises applications you want to migrate to Azure and you want an estimate of how much you can save in Azure, the TCO calculator is a better choice. You can access the TCO calculator by browsing to:

<https://azure.microsoft.com/en-us/pricing/tco/calculator>.

The first step when using the TCO calculator is to add details about your on-premises servers, databases, storage, and network usage. In Figure 4-9, an on-premises server has been configured for a Web App. You can configure all of the details about the server, including the OS, whether it's a VM or a physical server, and more.

Define your workloads

Enter the details of your on-premises workloads. This information will be used to understand your current TCO and recommended services in Azure.

Servers

Enter the details of your on-premises server infrastructure. After adding a workload, select the workload type and enter the remaining details.

Web Server

Workload ⓘ	Environment ⓘ	Operating system ⓘ	Servers ⓘ	Procs per server ⓘ	Core(s) per proc ⓘ
Web App ▾	Physical Servers ▾	Windows ▾	1 (1 - 9999)	2 (1 - 4)	4 (1 - 8)
RAM (GB) ⓘ	Optimize by ⓘ	Auto scaling ⓘ			
2 (1 - 448)	CPU ▾	☐			

+

 Add server workload

Figure 4-9 Configuring an on-premises server in the TCO calculator

Databases and storage systems that are on-premises should also be added, in addition to any network usage for your application. In [Figure 4-10](#), a storage system has been added, and network usage for the app has been specified.

Storage

Enter the details of your on-premises storage infrastructure. After adding storage, select the storage type and enter the remaining details.

Image storage

Storage type ⓘ

Capacity ⓘ

Backup ⓘ

Archive ⓘ

NAS/File Share ▾

3

TB ▾

(1 - 5000)

3

TB ▾

(0 - 5000)

6

TB ▾

(0 - 5000)

+ Add storage

Networking

Enter the amount of network bandwidth you currently consume in your on-premises environment.

Outbound bandwidth ⓘ

2

GB ▾

(1 - 2000)

Next

Figure 4-10 Configuring storage and networking

After entering all of your on-premises workloads, you can view the assumptions the TCO calculator uses by clicking **Next**. The TCO calculator uses a

comprehensive list of on-premises expense assumptions that Microsoft has put together based on years of experience, and these assumptions are used to provide you with the best estimate possible of your cost savings. As shown in Figure 4-11, assumptions include items such as whether you've purchased a Software Assurance plan for your on-premises servers, details on your current expenses on-premises, your IT labor costs, and much more. For an accurate TCO estimate, it's best to carefully record your expenses before generating a TCO report.

Storage costs

Storage procurement cost/GB for local disk/SAN-SSD	3	(USD)
Storage procurement cost/GB for local disk/SAN-HDD	2	(USD)
Storage procurement cost/GB for NAS/file storage	2	(USD)
Storage procurement cost/GB for Blob storage	2	(USD)
Annual enterprise storage software support cost	10	(%)
Cost per tape drive	4500	(USD)

IT labor costs

Number of physical servers that can be managed by a full time administrator	387	
Number of virtual machines that can be managed by a full time administrator	516	
Hourly rate for IT administrator	50	(USD)

Other assumptions

The following assumptions also affect the TCO model, but typically require less adjustment by customers. You can come back to this section at any time and adjust the assumptions.

☒ Hardware costs

☐ Software costs

Figure 4-11 Adjusting assumptions made by the TCO calculator

After you adjust your assumptions, scroll to the bottom of the screen and click **Next** to view your TCO report. Your TCO report shows you how much you can save over the next 5 years by moving your app to Azure as shown in [Figure 4-12](#).

View report

Timeframe ⓘ

5 Years

Region ⓘ

East US

Licensing program ⓘ

Microsoft Online Services Program

Over 5 year(s) with Microsoft Azure, your estimated cost

savings could be as much as **\$65,734**

Total on-premises vs. Azure cost over time

Savings from running workloads in Azure accrue over time. The following shows how those savings add up over years.




Figure 4-12 TCO savings report

A TCO report includes detailed charts of expense savings, and at the bottom of the report, you'll find a breakdown of on-premises costs and Azure costs so you can easily determine where you'll save money. Just as with the pricing calculator, reports generated by the TCO calculator can be downloaded, saved, and copied by clicking the appropriate button as shown in [Figure 4-13](#).

On-premises cost breakdown summary		Azure cost breakdown summary	
Category	Cost	Category	Cost
Compute	\$87,396.15	Compute	\$17,556.60
Hardware	\$17,296.00	Data Center	\$0.00
Software	\$4,808.75	Networking	\$0.00
Electricity	\$2,102.40	Storage	\$41,748.90
Database	\$63,189.00	IT Labor	\$0.00
Data Center	\$10,187.10		
Networking	\$6,655.43		
Storage	\$18,216.00		
IT Labor	\$2,585.00		
Total	\$125,040.00	Total	\$59,306.00
Estimated on-premises cost (5 year(s))		Estimated Azure cost (5 year(s))	
☒ Compute cost		Azure compute cost	
☒ Data center cost		Azure data center cost	
Total on-premises cost over five year(s)	\$125,040.00	Total Azure cost over five year(s)	\$59,306.00
A total savings of \$65,734.00 with Microsoft Azure			
Download Share Save			

Figure 4-13 Summary of costs on-premises and on Azure

Best practices for minimizing Azure costs

Minimizing your costs in Azure begins with careful planning before you even create a single Azure resource. Planning needs to involve the key players in your organization such as your finance department, managers who are responsible for budget planning and implementation, and application designers who are in the best position to decide what types of resources are likely needed.

Tools, such as the pricing calculator and TCO calculator, are valuable in planning, but they are only as accurate as the data you feed to them. Make sure that, as you analyze your on-premises workloads, you ensure that all of the resources your app is using are actually required by the app. If you are using a server that is far more powerful than needs dictate, that's an important point to consider in your planning.



Exam Tip

An important part of minimizing costs in Azure is to ensure that you fully use all of your cloud resources. Because most cloud usage is billed on consumption of a resource, not using

portions of a resource represents unnecessary expenses. Proper planning can help avoid non-utilized cloud resources.

As you're analyzing your on-premises usage and planning your move to Azure, you should take the opportunity to organize your resources based on expense accountability. You can then use that organizational structure in Azure to apply tags to resources so that each organization within your company has proper visibility into their expenses. This will help to ensure that if a resource is being under-utilized, the right people will quickly have visibility into that so adjustments can be made.

Once planning is complete, you'll need to determine which Azure subscription plan is most suitable for your needs. If you're planning on using Azure to host a long-term production application, you can save money by purchasing an Enterprise Agreement and agreeing to a longer-term commitment. However, if you're only planning on conducting test for a short period of time, a Free or Pay-As-You-Go subscription is a better choice.

At this stage in your planning, the pricing calculator and TCO calculator can be helpful in estimating your cloud expenses. Using the option to share a copy of your

estimates and reports can help ensure that everyone in your organization is on the same page, and can help you adjust your calculations as necessary.



Exam Tip

As you're planning your cloud deployments, make sure to account for the fact that Azure can scale your resources based on application needs. Don't default to the most powerful servers and other resources you think you will need. Instead, choose product SKUs that meet your minimum needs and configure scaling rules that can accommodate additional resource needs as application usage patterns increase.

Once you've deployed resources to Azure, it's important to monitor resource usage carefully. While you can use tools available in the Azure portal to review usage of individual resources, it's more effective to gain an overall view of resource usage using tools such as Azure Advisor.

[More Info Azure Advisor](#)

If you need to refresh your memory on using Azure Advisor, see “[Azure Advisor](#)” in Skill 2.4: “Understand Azure management tools”.

Your Azure invoice will also have details on resource usage, and as long as you organized your resources effectively, and tagged resources based on organizational alignment, you can easily provide different organizations within your company with details on their specific usage.

More Info Azure Cost Management

For more information on Azure Cost Management, see <https://docs.microsoft.com/azure/cost-management>.

If you have smaller batch jobs that you want to run in Azure, you can save substantially by using VMs that Microsoft has allocated in data centers, but that aren’t being used by customers. This offering is called Azure Batch, and it uses non-utilized VMs to run workloads that aren’t time-sensitive and that don’t need to run for long periods. You can read more about Azure Batch at: <https://docs.microsoft.com/en-us/azure/batch/batch-low-pri-vms>.

More Info Best Practices For Minimizing Costs

For more information on best practices for minimizing costs, see: <https://docs.microsoft.com/azure/cost-management/cost-mgt-best-practices>.

Azure Cost Management

Azure Cost Management is a tool in Azure that makes it easy to analyze your costs at a granular level. Cost Management allows you to create a budget for your Azure expenses, set configurable alerts so you'll know if you are approaching a budgeted limit, and analyze your costs in detail.

To get started with Cost Management, open the Azure portal and search for **Cost Management**, and click on **Cost Management + Billing**.



Exam Tip

You will also see Cost Management listed under Azure Marketplace. That's a different offering that is based on Cloudyn, a cloud expense management company that Microsoft purchased.

Once Cost Management + Billing loads in the portal, click on **Cost Management** as shown in [Figure 4-14](#) to access Cost Management.




Figure 4-14 Cost Management + Billing and Cost Management

To effectively monitor your costs, you should create a budget in Cost Management. Creating a budget isn't required, but it will allow you to visualize your spending compared to your planned expenses.

1. Click on **Budgets**, and then click **Add**, as shown in [Figure 4-15](#).




Figure 4-15 Adding a new budget

2. Enter a name for your budget.
3. Enter a spending amount and the period at which your spending resets.
4. Enter a start date for your budget.
5. Enter an expiration date.
6. Configure any alerts for your budget.
7. Enter email addresses of people who should be sent any triggered alert information.
8. Click **Create** to complete your budget.

In Figure 4-16, a \$10,000 budget is being created for a fiscal year. An alert is being added when 80% of that

budget has been reached, and Azure will send an email to jim@contosopharm.com if costs have reached \$8,000.

Create budget

Budget



Manage action groups



Help



Using action groups with budget thresholds helps you manage notifications and automate actions when your thresholds have been exceeded. [Learn more.](#)

* Name

FY_Budget

* Amount

10000

* Resets ⓘ

Annually

* Start date ⓘ

2019

January

1

* Expiration date ⓘ

2020-12-31

Alerts ⓘ

Configure alert conditions and send email notifications based on your spend.

* Alert conditions



Delete



% OF BUDGET

AMOUNT

ACTION GROUP

ACTION GROUP TYPE



80

8000

None



None

* Alert recipients (email)



Delete



ALERT RECIPIENTS (EMAIL)



jim@contosopharm.com

example@email.com

Create

Figure 4-16 Creating a budget

After you create a budget, click on **Cost Analysis** to see how your spending compares to your budget. In [Figure 4-17](#), you can see how we've spent a little over the budgeted yearly spending, based on the current month of May. As you can see, we are spending over twice as much as we've budgeted.



Figure 4-17 Viewing cost analysis

In [Figure 4-17](#), you can view expenses for an entire year because the budget was created for a fiscal year. Click on **Custom View** to choose to view only accumulated costs, which would show the same view as [Figure 4-17](#). However, when you do that, you won't see your budget included in the graph. You can also choose to view daily costs, costs by service, or costs by resource. For each view, you can choose a date range, a granularity (such as daily, monthly, and so forth), and how you want the display to group by expenses (by billing period, resource group name, and so on).

To drill down further into the expenses, you can apply a custom filter by clicking **Add Filter**. In [Figure 4-18](#), a filter is being added to show only expenses for an application named databricks.



Figure 4-18 Applying a filter

SKILL 4.3: UNDERSTAND THE SUPPORT OPTIONS AVAILABLE IN AZURE

Azure provides many diagnostic tools you can use to troubleshoot your application when things go wrong, but it's likely you may need assistance from Azure support at some point. Whether you need additional help troubleshooting your application, or you need support for the Azure platform itself, Microsoft offers numerous support options for interacting with their world-class support organization.

This section covers:

- Support plans
- How to open a support case
- Available support channels outside of support plans
- Knowledge Center

Support plans

Microsoft offers numerous support plans for Azure customers. Before we get into the details of each plan, there are a few terms you should be familiar with related to Azure support.

- **Business hours** Microsoft defines *business hours* for most countries as weekdays from 9:00 AM to 5:00 PM local time. However, in North America, business hours are weekdays from 6:00 AM to 6:00 PM Pacific time, and in Japan, business hours are weekdays from 9:00 AM to 5:30 PM. In all regions, business hours do not include holidays.
- **Severity A case** Microsoft uses *Severity A* to refer to a

production application that is entirely down, or when a critical component of a production app is unavailable.

- **Severity B case** Microsoft uses *Severity B* to refer to a production application that is moderately impacted. This severity level is subjective and agreed to by Microsoft support and the customer.
- **Severity C case** Microsoft uses *Severity C* to refer to a situation that is causing minimal impact. These are cases that refer to problems no longer happening or cases that aren't impacting a production application.

Microsoft offers the following support plans for Azure.

- **Basic** Limited support that's free for all Azure subscriptions.
- **Developer** Azure support for free trial and non-production applications.
- **Standard** Azure support for production applications.
- **Professional Direct** Azure support for business-critical applications.
- **Premier** Contracted support for all Microsoft products, including Azure.

All support plans offer 24x7 support for any billing issues or subscription issues, access to Azure Advisor recommendations, the Service Health Dashboard, and the Health API. All paid support plans offer access to Microsoft support engineers.

[More Information Other Support Options](#)

Microsoft offers numerous support options outside of the support plans outlined here. They will be discussed later in this chapter in the “Available support channels outside of support plans” section.

Table 4-2 outlines the differentiators between paid support plans.

Table 4-2 Zones and geographies

	Developer	Standard	Pro Direct	Premier
Price	\$29 per month	\$100 per month	\$1,000 per month	Contract price varies
Tech Support	Access to support engineers via email only during business hours.	Access to support engineers 24x7 via email or phone.	Access to support engineers 24x7 via email or phone.	Access to support engineers 24x7 via email or phone.
Severity Available	Severity C only.	All severity levels.	All severity levels.	All severity levels.
Response Time	Fewer than 8 business hours.	Sev C: Fewer than 8 business hours. Sev B: Fewer than 4	Sev C: Fewer than 4 business hours. Sev B: Fewer than 2 business hours.	Sev C: Fewer than 4 business hours. Sev B: Fewer than 2 business hours. Sev A: Under 1 business hour

		business hours. Sev A: Under 1 business hour.	Sev A: Under 1 business hour.	or under 15 minutes with the purchase of Azure Rapid Response or Azure Event Management.
Architecture Guidance	General only	General only	Guidance based on best practices.	Customized guidance that includes design reviews, performance tuning, configuration assistance, and so on.
Operations Support			Assistance with onboarding, service reviews, and Azure Advisor consultations.	Service reviews and reporting led by a technical account manager (TAM) assigned to your account.
Training			Web seminars conducted by Azure engineering teams.	Web seminars conducted by Azure engineering teams. On-demand training coordinated

				via TAM.
Proactive Guidance			Provided by delivery manage for Pro Direct.	Provided by TAM.
Launch Support				Available via Azure Event Management for an additional cost.



Exam Tip

You can change your support plan or cancel your support plan via the Azure portal. If you cancel a support plan partway through the month, you are not refunded for the prorated amount.

How to open a support case

Azure support cases are opened using the Azure portal. Support cases can be created from the portal home page by clicking on **Help + Support** from the menu on the left side of the page. You can also create a support case from within a particular Azure resource by opening the resource and clicking on **New Support Request** from the Support + Troubleshooting section of the menu on

the left.

Note Opening Support Cases

When you open a support case from within an Azure resource, the resource type and resource name are automatically populated for you. If you open a support case from the portal home page, you will need to select the resource before you can open a support case.

In this section, we will show the experience of creating a support case from the portal home page.

After clicking on **Help + Support** in the Azure portal, you're presented with some options for self-help as shown in [Figure 4-19](#).




Figure 4-19 Help and support options in the Azure portal

The self-help options included here are general in nature and include links to documentation and links to

forums where you can ask for help from other Azure users. You can also tweet to @AzureSupport for quick assistance for simple issues.

To create a support case to a Microsoft support engineer, click on **New Support Request**. Creating a support case is a four-step process, the first of which is to provide basic information as shown in [Figure 4-20](#).

1. Select **Yes** to indicate you have an issue related to an Azure subscription. (The No option is for problems related to Azure Active Directory.)
2. Select the issue type. The issue type can be a billing issue, a subscription issue, a quota issue, or a technical issue. In this example, we're using the Technical issue type.
3. Select the Azure service.
4. Select the resource you need help with. (Your choice will be restricted to the type of Azure service you selected.)
5. Select your problem type. Problem type options will differ depending on the service type.
6. Select the problem subtype. Problem subtype options will differ depending on the service type.
7. Enter a brief subject for your support case.
8. Click **Next: Solutions** to process to the next step.

Home > Help + support > New support request (preview)

New support request (preview)

Basics Solutions Details Review + create

Create a new support request to get assistance with billing, subscription, technical or quota management issues. Complete the Basics tab by selecting the options that best describe your problem. Providing detailed, accurate information can help to solve your issues faster. Looking for the old experience? Click [here](#)

* Is your issue related to Azure subscription? Yes

* Issue type Technical

* Subscription Jim's MSDN Subscription (2ed0e6a0-aaf7-4f5f-aff6-bf2...
Can't find your subscription? [Show more](#)

* Service ☒ My services ☐ All services
Virtual Machine running Windows

* Resource VM1

* Problem type VM restarted or stopped unexpectedly

* Problem subtype Help diagnose my VM restart issue

* Subject My VM restarted and I don't know why

[Next: Solutions >>](#)

Figure 4-20 Help and support options in the Azure portal

Microsoft constantly analyzes historical customer issues so they can offer you possible solutions based on the information you provide when opening a case. They

use the problem type, the problem subtype, and the text you enter in the subject to determine what the problem might be.

In [Figure 4-21](#), you can see that Microsoft is suggesting that the restart might have been due to a user-initiated restart of the VM. If you decide that this didn't cause the issue, click on **Next: Details** to move to the next step.

[Home](#) > [Help + support](#) > [New support request \(preview\)](#)

New support request (preview)

[Basics](#)[Solutions](#)[Details](#)[Review + create](#)

Want a solution right now?

Try following the recommended steps below. These solutions are written by Azure engineers, and will solve most common issues.

 We ran diagnostics on your resource and found an issue [Download](#)

VM Availability incident diagnostic information for _VM1:

We identified that your VM became unavailable at **2019-03-30 19:03:39 (UTC)**. This expected occurrence was caused by a **user initiated shutdown action**.

The shutdown was triggered by an authorized user or process from either the Azure Portal or from Azure Resource Manager interfaces. As a result, your VM was shut down and remained in this state until user

[Show more](#) 

Was this helpful? [Yes](#) [No](#)

 Recommended Solution

4 out of 5 customers resolved their VM restart issue using the steps below.

Recommended Steps

1. Review the below documents in this article to understand the different possible scenarios
2. Review the [Current Azure Status](#)  or [Azure Status - History](#)  for outages
3. [Understand more about Resource Health Center](#)  and using [Resource Health blade](#) for any impactful events specific for your VM

[Show more](#) 

Was this helpful? [Yes](#) [No](#)

[<< Previous: Basics](#)

[Next: Details >>](#)

Figure 4-21 Possible causes of the problem

The final step in creating a case is to enter the details for your issue as shown in Figure 4-22. Some of these options will differ depending on the resource type. In this example, options are for an Azure VM.

New support request (preview)

[Basics](#)

[Solutions](#)

Details

[Review + create](#)

Information provided on this tab will be used to further assess your issue and help the support engineer troubleshoot the problem. Verify the contact information before moving to the Review + Create.

PROBLEM DETAILS

What is the error you received?

No error. Just restarted for some unknown reason. ✓

Which machine version are you running?

Windows 10/Windows server 2016 ▾

* Description

My VM restarted and I don't know why. I didn't do it. ✓

When did the problem start?

2019-04-05  09:20 AM

File upload

[Choose file to upload](#) 

Consent

☒ Share diagnostic information 

SUPPORT METHOD

Support plan

Premier

* Severity

C - Minimal impact ▾

* Preferred contact method

☐ Contact me later

Email

☐ Call me later

Phone

* Response hours

Business Hours

* Support language

English ▾

CONTACT INFO [Edit](#)

Contact name

Jim Cheshire

Email

Additional email for notification

--

Phone

--

Country/region

United States

[< < Previous: Solutions](#)

[Next: Review + create >>](#)

Figure 4-22 Entering support case details

1. Enter any errors that you received.
2. Select the operating system from the dropdown.
3. Enter a description of your issue.
4. Enter a date and time for when your issue started. If the issue was a single event, enter the date and time that the issue occurred.
5. Upload any relevant files such as screenshots or error logs.
6. If you would like to share diagnostic information with Microsoft, check the **Consent** checkbox.
7. Choose your severity level. If you have a Developer support plan, you will only be able to select Severity C.
8. Choose your contact method. Microsoft will contact you within a time period dictated by your severity level and your support plan.
9. Choose your support language.
10. Edit your contact info if necessary.
11. Click **Next: Review + Create** to continue. You'll be shown the information you entered for confirmation and will have a Create button to complete the support case.

Available support channels outside of support plans

If you don't have an Azure support plan, you can still get help with Azure technical issues from forums or from Twitter, but you won't have any support SLAs, and you won't be able to talk directly to a Microsoft support engineer.

There are two forum channels available for Azure issues.

- **MSDN forums** Accessible at: <https://aka.ms/MSDNForums>. Search for your product to find the relevant forum.
- **Stack Overflow forums** Accessible at: <https://stackoverflow.com>. Click on **Tags** and search for your Azure service.

MSDN and Stack Overflow forums are user-to-user forums where Azure customers can help each other. Microsoft also monitors the forums and can provide assistance with simple issues. However, in some cases, they will ask you to open a support case.

You can also tweet to @AzureSupport for assistance with simple issues. This is the official Microsoft Twitter account for helping you find answers to common questions and support for basic issues.

Knowledge Center

To assist you with finding documentation and blog posts on common issues, Microsoft developed the Knowledge Center. You can access the Knowledge Center by browsing to: <https://azure.microsoft.com/en-ca/resources/knowledge-center>.

As shown in [Figure 4-23](#), you can filter on the Azure product you're interested in. Each product has a series of tags that you can use to further filter the links that you see in the Knowledge Center. You can also enter in a search term to find something more specific to your problem.




Figure 4-23 Knowledge Center

SKILL 4.4: DESCRIBE AZURE SERVICE LEVEL AGREEMENTS

Many of the services you use today include a *service level agreement* (SLA) that serves as a contract between you and the service provider for a certain level of service.

This section covers:

- Service level agreement (SLA)
- Determine the SLA for a particular Azure product or service

Service level agreement (SLA)

SLAs establish specific targets for availability, and they also define what the service provider will do when those targets aren't met. SLAs are expressed as a percentage and are almost always 99% or higher. The highest level of availability expressed in an SLA is 99.999%, commonly referred to as *5 nines*. To provide you with some context as we discuss SLAs, a service with an SLA of 5 nines guarantees that downtime over an entire year will not exceed 5.56 minutes. A more reasonable SLA of 99.9% guarantees that downtime over the period of a month will not exceed 43.2 minutes.

More Information Slas And Downtime Amounts

For details on SLA levels and maximum downtime allowed within the SLA, see:
<https://docs.microsoft.com/azure/architecture/resiliency/#slas>.

An important concept in cloud service SLAs is that the cloud provider considers an application to be outside of SLA only when the availability percentage is not met due to an issue that the cloud provider can control. In other words, if you deploy new code to your application, and it causes your application to crash, the cloud provider is not going to consider that a breach of SLA. If you install a component onto your virtual machine and it causes the machine to go down, that's not within the cloud provider's control, and isn't classified as not meeting SLA.

Because SLAs only refer to problems within control of the cloud provider, when an application suffers from lack of availability, it's important to determine whether the problem is a platform issue or an issue with your code or configuration. Answering that question can be more difficult than you might think.

Azure is a highly complex environment involving a large number of services operating together. For example, Azure App Service (one of Azure's most popular services) uses Azure cloud services, Azure DNS systems, Azure Storage, Azure SQL Database, and other Azure services under the hood. Performance

degradation of any of those services can impact the availability of an application running in App Service. If you report that your App Service application is unavailable, Microsoft needs to determine whether that's a problem on their end or a problem with your application.

Microsoft maintains an enormous amount of diagnostic data on all Azure operations across all Azure services. When you open a support case with Microsoft to report that your application is unavailable, Microsoft can perform data analytics against this data to determine if there was a problem with the Azure platform itself.

If you believe that your application's availability has fallen below the SLA, it's your responsibility to submit a claim to Microsoft. You can do that by opening a support case. If Microsoft determines that the SLA has not been met, you may receive a credit on your Azure invoice. The amount of the credit depends on the duration that SLA was not met and the specific Azure service's SLA policy.



Exam Tip

In order to be eligible for a credit due to a failure to meet SLA, you must

submit a claim to Microsoft within two months of the end of the billing cycle during which the downtime occurred.

Most Azure services offer an SLA of at least 99.9%, and higher SLAs can be achieved with additional configuration by the customer. For example, a single VM using Premium storage for all disks has an SLA of 99.9%. If you deploy two or more VMs into the same availability set, that SLA increases to 99.95%. Deploy those two or more instances across two or more availability zones within the same Azure region and the SLA moves to 99.99%.



Exam Tip

Microsoft occasionally changes SLAs. If the terms of an SLA change, the new terms will go into effect for you only when you renew your Azure subscription. Until that time, you will fall under the SLA that was in effect when your subscription was last renewed or when you signed up for an Azure subscription.

Determine the SLA for a particular Azure product or service

Because SLA varies between Azure services, and because specific configurations can impact the SLA of a single Azure service, it's important to be able to determine the specific SLA for the Azure services you are using. Microsoft provides a web page that has details on the SLA for every Azure service. You can find it at: <https://azure.microsoft.com/en-us/support/legal/sla>.

As shown in Figure 4-24, once on the SLA web page, you can select a category to see all Azure services in that category. You can also enter your service name in the search box to find the SLA for that service. Once you locate the service you're interested in, click it to read details on the SLA.


Contact Sales: 1-800-867-1389
Search
My account
Portal
Sign in

[Overview](#)
[Solutions](#)
[Products](#)
[Documentation](#)
[Pricing](#)
[Training](#)
[Marketplace](#)
[Partners](#)
[Support](#)
[Blog](#)
[More](#)
[Free account](#)

Service Level Agreements

Read the SLAs to learn about our uptime guarantees and downtime credit policies

The Service Level Agreement (SLA) describes Microsoft's commitments for uptime and connectivity. The SLA for individual Azure services are listed below.

AI + Machine Learning Analytics Compute Containers Databases Developer Tools DevOps	Internet of Things Management Media Migration Mobile Networking Storage Security	AI + Machine Learning Create the next generation of applications using artificial intelligence capabilities for any developer and any scenario Azure Bot Service Intelligent, serverless bot service that scales on demand Machine Learning Studio Easily build, deploy, and manage predictive analytics solutions Cognitive Services Add smart API capabilities to enable contextual interactions Microsoft Genomics Power genome sequencing & research insights Azure Machine Learning service Bring AI to everyone with an end-to-end, scalable, trusted platform with experimentation and model management

Figure 4-24 Azure SLA web page

When you click on a service, you'll see details on the SLA provided by that service. Figure 4-25 shows the SLA page for Azure Virtual Machines. The three bullet points at the top of the page outline the SLA for Azure VMs.



Contact Sales: 1-800-867-1389  Search  My account Portal Sign in

Overview  Solutions Products  Documentation Pricing Training Marketplace  Partners  Support  Blog More  [Free account >](#)

SLA for Virtual Machines

Last updated: March 2018

- For all Virtual Machines that have two or more instances deployed across two or more Availability Zones in the same Azure region, we guarantee you will have Virtual Machine Connectivity to at least one instance at least 99.99% of the time.
- For all Virtual Machines that have two or more instances deployed in the same Availability Set, we guarantee you will have Virtual Machine Connectivity to at least one instance at least 99.95% of the time.
- For any Single Instance Virtual Machine using premium storage for all Operating System Disks and Data Disks, we guarantee you will have Virtual Machine Connectivity of at least 99.9%.

 [Introduction](#)

 [General Terms](#)

 [SLA details](#)

Additional Definitions

"**Availability Set**" refers to two or more Virtual Machines deployed across different Fault Domains to avoid a single point of failure.

"**Availability Zone**" is a fault-isolated area within an Azure region, providing redundant power, cooling, and networking.

"**Data Disk**" is a persistent virtual hard disk, attached to a Virtual Machine, used to store application data.

"**Fault Domain**" is a collection of servers that share common resources such as power and network connectivity.

Figure 4-25 Azure VMs SLA

The Introduction section describes Azure SLAs in general. The General Terms section describes SLA terms such as Management Portal, Service Level, and Downtime that refer to all Azure services. It also explains how you can make a claim and limitations for Azure SLAs.

The SLA Details section applies to the specific Azure service you're viewing. For example, this section on the VM SLA page defines VM-specific terms that relate to the SLA for VMs. If you scroll down, you'll see additional details shown in Figure 4-26, including how to calculate availability and the amount of credit you may receive if an SLA isn't met.



[Contact Sales: 1-800-867-1389](#)
[Search](#)
[My account](#)
[Portal](#)
[Sign in](#)

[Overview](#)
[Solutions](#)
[Products](#)
[Documentation](#)
[Pricing](#)
[Training](#)
[Marketplace](#)
[Partners](#)
[Support](#)
[Blog](#)
[More](#)
[Free account](#)

Monthly Uptime Calculation and Service Levels for Virtual Machines in Availability Zones

"Maximum Available Minutes" is the total accumulated minutes during a billing month that have two or more instances deployed across two or more Availability Zones in the same region. Maximum Available Minutes is measured from when at least two Virtual Machines across two Availability Zones in the same region have both been started resultant from action initiated by Customer to the time Customer has initiated an action that would result in stopping or deleting the Virtual Machines.

"Downtime" is the total accumulated minutes that are part of Maximum Available Minutes that have no Virtual Machine Connectivity in the region.

"Monthly Uptime Percentage" for Virtual Machines in Availability Zones is calculated as Maximum Available Minutes less Downtime divided by Maximum Available Minutes in a billing month for a given Microsoft Azure subscription. Monthly Uptime Percentage is represented by the following formula:

$$\text{Monthly Uptime \%} = (\text{Maximum Available Minutes} - \text{Downtime}) / \text{Maximum Available Minutes} \times 100$$

The following Service Levels and Service Credits are applicable to Customer's use of Virtual Machines, deployed across two or more Availability Zones in the same region:

MONTHLY UPTIME PERCENTAGE	SERVICE CREDIT
< 99.99%	10%
< 99%	25%
< 95%	100%

Monthly Uptime Calculation and Service Levels for Virtual Machines in an Availability Set

Figure 4-26 Details on Azure VM SLA

If your application uses multiple Azure services, multiple SLAs will apply to you. If you experience downtime, you must submit a claim for all Azure services that fell below SLA if you want to be considered for a credit. However, monetary credit isn't your only concern related to availability of your application. Downtime in your application impacts your business negatively, so you want to always ensure that you have the highest SLA possible, and when you're dealing with multiple Azure services with different SLAs, it's important to understand how that impacts your overall SLA.

When calculating the SLA for an application using multiple Azure services, you must calculate a composite SLA based on the services you're using. For example, if you have an App Service Web App that also uses a single Azure VM using Premium storage, you have to combine the SLA for both services to determine your application's overall SLA.

Note Composite SLAs

It's important to understand that individual service SLAs still apply to you when you're using multiple Azure services. However, understanding composite SLAs is important because it allows you to determine when a specific configuration is increasing the likelihood that you will experience downtime.

The SLA for App Service is 99.95%, and the SLA for a

single VM running Premium storage is 99.9%. Therefore, your overall SLA for your application is $99.95\% \times 99.9\%$, or 99.85%. By deploying two VMs into two availability zones in the same region, you can obtain a 99.99% SLA for your VMs, and that increases your overall SLA to 99.94%.

More Information Computing Composite Slas

For more information on computing composite SLAs, see: <https://docs.microsoft.com/azure/architecture/resiliency/#composite-slas>.

SKILL 4.5: UNDERSTAND SERVICE LIFECYCLE IN AZURE

Azure is an always-changing environment, and new services are always being introduced. Existing services also evolve over time and introduce new features. It's important to understand the service lifecycle in Azure, how you can keep up with changes, and how a service's lifecycle might impact your support and your SLA.

This section covers:

- Public and private preview features
- How to access preview features
- General availability (GA)
- Monitoring feature updates

Public and private preview features

As Azure product teams develop new services and features, it's important for them to get feedback from customers using those services and features in a real-world environment. For that reason, Microsoft will often offer new services and features to customers as *preview offerings*. While the Microsoft official term is *preview*, you will often see people refer to these services and features as being a *beta* offering.



Exam Tip

Services and features that are in preview do not offer an SLA, and they are not meant to be used in production applications. Preview features are also usually not offered in all Azure regions. Microsoft will provide documentation on which regions are available for a specific preview.

Preview services and features are sometimes first offered as a private preview. In private preview, the service or feature is made available to a small set of customers for testing. Access to a private preview is sometimes by invitation from the engineering team

developing the service or feature. In other cases, Microsoft may provide a way for any customer to sign up for access to the private preview. If registration is open to everyone, Microsoft will close registration after a target number of customers have signed up.

Note Services Versus Features

Many previews are for features of an existing service. For example, App Service may add a new feature for the existing service, and before that feature is fully released, it will go through some period in the preview phase.

Private preview services and features commonly expose only a subset of the functionality that will eventually make it into the service or feature. Microsoft will often ask customers using a private preview to test specific scenarios and provide feedback. This helps engineering teams to uncover bugs and usability issues in the complex real-world environments that customers are using.



Exam Tip

Not all services or features offer a private preview. If private preview isn't offered, the service or feature is first made available as a public preview. All services and features go through some

period of public preview.

Private previews may be offered to customers at no cost, but it's more common for them to be offered at a substantial discount

Once a service or feature meets a specific bar set by the engineering team, it will transition to public preview. This usually occurs once the service or feature is fully functional or very close to it. However, if there are bugs in a specific part of the functionality that the engineering team feels is critical, they may delay public preview until those bugs are fixed.

Features and services that are in public preview are provided at a discounted rate, but like private preview features and services, they typically don't offer an SLA and are provided as-is.

How to access preview features

Customers participating in a private preview are sometimes given a secret link to the Azure portal that enables the service or feature. When the customer uses that link, Microsoft can use their Azure subscription ID to determine if they have registered and are approved for the private preview. If they aren't, the feature or service won't be available, even if they use the secret link.

In other situations, the Azure portal experience hasn't been developed for a private preview feature or service. In those cases, customers are given command-line instructions for using the service or feature. It's more common for the portal user interface to be developed during the private preview phase, so early adopters are usually given command-line access only.

Once a service or feature reaches public preview, it is made available to all customers in the regions where it's available, and no registration is required to use the service or feature. A preview badge will be displayed in the Azure portal so that users will know that the service or feature is a preview offering. [Figure 4-27](#) shows the Docker container features in an App Service Web App running on Windows, and each container setting carries a preview badge.




Figure 4-27 Preview features in App Service

Services and features that are in public preview are usually supported by Microsoft just as though they were fully released. However, SLAs don't apply to previews, and there are some situations where a service or feature won't be supported by Microsoft support engineers while in preview. In those cases, you may be referred to forums for support.

General availability

Once a preview service or feature reaches a quality and availability bar suitable to the engineering team, they will declare *general availability* or GA. At this point, the service or feature is fully supported.

Once a service or feature reaches GA, it falls under the SLA Microsoft provides. If it's a new service, a new SLA will be published on the SLA web page. For new features of existing services, once GA is reached, the feature will inherit the SLA of the service it's a feature of.

If you were using a feature or service during public preview, you will usually not have to do anything to be officially supported under GA. However, in some situations, Microsoft will ask that you delete any resources created during preview and recreate them. This usually happens when remnants left over from preview code might cause a problem with a service or feature running in GA.

When a service or features reaches GA, it may not GA in all Azure geographies. In those cases, other geographies will usually GA later in the lifecycle of the service or feature. Preview pricing may also remain in effect for some period of time after GA. Details such as this are published on the official GA announcement on the Azure website.

Monitoring feature updates

Microsoft will usually post announcements of new features and services on the Azure blog at:

<https://azure.microsoft.com/en-us/blog/topics/announcements>. However, a more reliable source of information on feature and service updates is the Azure Updates web page available at: <https://azure.microsoft.com/en-us/updates>.

Figure 4-28 shows the Azure Updates web page. By default, all updates are displayed, but you can filter on specific types of update using the Update Type dropdown.



Contact Sales: 1-800-867-1389  Search  My account Portal Sign in

Overview  Solutions Products  Documentation Pricing Training Marketplace  Partners  Support  Blog More  [Free account >](#)

Azure updates

Learn about important Azure product updates, roadmap, and announcements.
Subscribe to notifications to stay informed.

All Now available In preview In development

Products

Browse 

Search for a product 

Update type

All 

 [RSS feed](#)

April 2019

Apr 6

[Set custom metadata properties for Stream Analytics output](#)

Stream Analytics now allows users to attach query columns as user properties to outgoing messages to help better facilitate downstream processing and reporting etc.

Explore

Read the Azure blog for the latest news.

[Blog >](#)

Tell us what you think of Azure and what you want to see in the

Figure 4-28 Azure Updates web page

Technet24

To show only updates on a particular stage of product lifecycle, use the tabs at the top of the list. The Now Available tab shows all updates about GA services and features. The In Preview tab will show all services and features in public or private preview. The In Development tab will show information related to features or services that are currently in development but not yet available to customers.

Note Private Preview And In Development Updates

Private previews are sometimes not announced to the general public, so they may not always appear on the Azure Updates web page. Likewise, not all features or services in development will be announced.

To show updates only for the products you're interested in, you can search for a product by entering it in the search box. You can also click **Browse** and select one or more products from the list as shown in [Figure 4-29](#).




Figure 4-29 Filtering updates on specific products

THOUGHT EXPERIMENT

Let's put the concepts you've learned in this chapter into practice using a thought experiment. Answers to

this thought experiment can be found in the section that follows.

ContosoPharm has an on-premises application that uses a web server installed on a physical server running Windows Server 2016. It uses another physical server running SQL Server for back end data storage. They are also storing scanned documents in a network attached storage (NAS) device.

The database contains sales information for pharmacies that stock ContosoPharm products. Sales personnel access the database using a section of the website allocated to sales. The regulatory division of ContosoPharm uses a separate section of the website to upload scanned documents into the NAS device.

The IT director at ContosoPharm has recommended that the application be moved to the cloud. However, the CFO of the company requires that a detailed report be submitted proving that ContosoPharm can save money using the cloud.

The CIO, John, is concerned about separation of departments. In the current configuration, sales personnel are not able to access any of the resources used by the regulatory division, and they want to ensure that this policy remains in effect. John also wants to make sure that the critical nature of the application is realized after it moves to the cloud. If there is a problem

with application availability, John wants to ensure that they can engage someone at Microsoft about the problem within an hour's time.

The CEO, Jill, also has concerns about moving to the cloud. Jill's concerned about over-spending on cloud resources, and because expenses at ContosoPharm aren't the same month-over-month, she's concerned about keeping an eye on expenses over the long term.

You are charged with advising the IT director on meeting the requirements of the various C-level executives at ContosoPharm in order to convince them to move to Azure.

THOUGHT EXPERIMENT ANSWERS

This section provides the answers to the thought experiment.

You first decide to address the cost concerns of the CFO. Using the TCO calculator, you put together a detailed report that shows how much you can save over the next five years by moving to Azure. The TCO calculator takes all savings into account, including savings on computer hardware, electricity costs, and IT infrastructure costs. Using this method, you are able to clearly show a savings of tens of thousands of dollars over a five-year period. You are also able to download a copy of the report so that other C-level executives can refer to it to bolster your case.

You also put together an estimate of your total Azure expenses using the pricing calculator. By exporting this report to an Excel spreadsheet, the CFO is able to easily incorporate it into existing budgets to see how Azure expenses fit into financial planning.

To address the CIO's concerns about separation of departments, you recommend that two Azure subscriptions be used. Resources for sales can be created under one subscription and resources for the regulatory division can be created under the other subscription. You can then use RBAC to enforce access restrictions to the subscriptions so that only the people you want to have access can reach resources under each subscription.

To address the CIO's concerns related to support, you recommend that ContosoPharm purchase a Pro Direct support plan. While this support plan will cost ContosoPharm \$1,000 per month, it will give them access to Microsoft support engineers 24x7, and if a critical problem arises that impacts application availability, ConotosoPharm can open a Sev A support case with Microsoft and receive a response within one hour.

To address the CEO's concerns, you put together a presentation that outlines Cost Management. Your presentation includes information on creating budgets

that you can apply to expenses. You show how you can create alerts that will notify the appropriate people when expenses reach a certain point. Because expenses vary month-over-month, you demonstrate how budgets can apply to particular timeframes and how multiple budgets can be created to cover all spending scenarios.

You also present information on cost analysis in Cost Management and how you can easily view expenses broken down by services, resource groups, locations, and so on. These reports can also be scoped to specific subscriptions so the C-level executives can review expenses only for the sales personnel, only for the regulatory division, or for everyone combined.

Finally, you present information on Advisor Recommendations in Cost Management and how this feature can help to highlight areas where cost savings can be achieved. You also recommend that if Jill is truly interested in using Azure resources over the long-term, ContosoPharm can likely save money by purchasing an Enterprise Agreement and agreeing to a long-term commitment for usage of Azure resources.

CHAPTER SUMMARY

Costs and support are among the top concerns for companies moving to the cloud, and it's important to understand how to minimize costs while ensuring your support options meet your needs. In this chapter, you

learned about the following concepts related to pricing and support.

- Azure resources are created within an Azure subscription.
- Subscriptions have limits associated with them, and you can create additional subscriptions if you need more than these limits allow.
- Azure offers a free trial subscription and Pay-As-You-Go subscriptions.
- You can purchase Azure products and services directly from Microsoft or through a Microsoft Cloud Solution Partner (CSP).
- CSPs sell entire cloud solutions and you don't manage individual Azure resources.
- You can purchase Azure products and services from Microsoft in the Azure portal or you can commit to a long-term use of Azure resources and save money with an Enterprise Agreement.
- A free trial subscription gives you free access to the most popular Azure services for one year. It also provides \$200 in credit towards Azure services and products.
- Azure services are billed according to meters that are associated with a resource.
- Costs for Azure services may vary in different regions. Costs also vary by billing zones that include specific geographies.
- The pricing calculator makes it easy to estimate your Azure costs by selecting the products you intend on using, and estimates can be shared, saved for later reference, or exported to Excel.
- The TCO calculator allows you to determine your cost savings in Azure over on-premises expenses over a period of five years.
- You can control your expenses in Azure by ensuring that you fully utilize all Azure resources you're paying for.

- You can save money when you need cloud resources for smaller jobs that aren't time-sensitive by using Azure Batch to run your workloads on non-utilized VMs.
- Creating budgets in Azure Cost Management can make it easy to see when expenses are approaching pre-defined limits, and alerts can be used to notify the proper people when expenses reach a defined threshold.
- Cost analysis can help you to see which resources are contributing to your Azure expenses.
- Microsoft offers free support for subscription issues and billing issues. Technical issues require the purchase of an Azure support plan or a Premier account.
- Support plan levels dictate when you can speak to Microsoft support personnel and the response time that Microsoft promises when you open a support case.
- You can open a support case from within the portal using either the home page or the New Support Request menu option while within a specific Azure resource.
- Microsoft provides MSDN forums and Stack Overflow forums for support outside of support plans. You can also use the @AzureSupport Twitter account for simple issues.
- The Knowledge Center can help you find support articles for specific Azure products.
- Azure services offer a service level agreement (SLA) that guarantees a certain level of availability. Services that fail to meet SLA may make you eligible for a credit on your Azure invoice.
- You can use the SLA web page on the Microsoft website to find SLA details for all Azure services.
- The Azure service lifecycle can include a private preview and always includes a public preview and general availability (GA).
- Services and features in preview do not offer an SLA and are

usually available at a discount.

- Private previews are usually made available using command-line tools. Public previews are available to everyone in the Azure portal.
- Once a feature or service meets the quality level for full support and SLA, it becomes generally available (GA).
- The Azure Updates page provides details on feature and service updates and lifecycles.

Index

A

AAD, [197](#), [199](#)

Access Control, [178](#), [198](#), [200](#)

Access Policies, [188](#)

ACI, [48-50](#)

Active Directory, [viii](#), [75](#), [134-135](#), [153](#), [169](#), [171](#), [192](#), [229](#),
[254](#)

Add Assessment, [222](#)

Add Condition, [213](#)

Add Filter, [250](#)

Add Metric, [210](#)

ADDS, [5](#)

Advanced Threat Protection, [228-229](#)

ADVISOR, [26](#), [133](#), [145-146](#), [151](#), [192](#), [205](#), [246](#)

Agility, [2](#)

AI, [25-26](#), [68](#), [88-91](#), [100](#), [151](#)

AIP, [189](#), [191](#)

AKS, [49-50](#)

All Guest Users, [177](#)

All Services, [37](#), [135](#), [193](#), [217](#)

API, [31](#), [32](#), [59](#), [65](#), [121](#), [133](#), [151](#), [252](#)

App Service, [xiii](#), [13](#), [31](#), [33](#), [39](#), [50](#), [53-54](#), [67](#), [118](#), [120-123](#), [136](#), [145](#), [168](#), [180-181](#), [196](#), [198](#), [226](#), [258](#), [259](#), [262-263](#)

App Service Environments, [29](#)

App Service Plans, [136](#)

Append Adds, [197](#)

Append blobs, [58](#)

Application Administrator Users, [78](#)

Application Builder Users, [78](#)

Application Gateway, [56](#), [162](#), [169](#)

AppService, [168](#)

Area Chart, [211](#)

Assign Policy, [195](#)

Assignments, [177](#), [195](#)

ATP, [191-192](#), [228](#)

Audit Logs, [197](#)

Audit Reports Comprehensive, [221](#)

AuditIfNotExists Allows, [197](#)

AVAILABILITY ZONES, [viii](#), [2](#), [27-30](#), [44](#), [46-47](#), [149](#), [238](#)

AzResourceGroup, [142](#)

Azure database, [viii](#), [39](#), [61](#)

Azure Active Directory, [59](#), [83](#), [228](#)

Azure Advisor, [149](#), [206](#), [230](#), [252](#)

Azure Application Gateway, [56](#)

Azure Blob Storage, [72](#), [82](#), [104](#)

Azure Cosmos, [104](#)

Azure Cost Management, [247](#)

Azure Data Lake Storage, [84](#)

Azure Databricks, [94](#), [96](#)

Azure Disk Storage Disk, [59](#)

Azure Event Management, [252](#)

Azure Files, [61](#)

Azure Firewall, [155-156](#), [159-162](#), [164](#), [166](#), [169](#), [228-229](#)

Azure Function, [80](#), [131](#)

Azure Functions

 HttpTrigger, [125](#)

Azure Functions Core Tools, [123-124](#)

Azure Government, [219](#), [226](#), [228](#)

Azure Information Protection Solutions, [189](#)

Azure Key Vault, [180](#), [185](#), [188-189](#), [204](#), [228](#)

Azure Load Balancer, [55](#)

Azure Logic App, [80](#)

The Azure Machine Learning Service, [98](#)

Azure Security Center, [11](#), [180](#), [194-195](#), [228](#)

Azure Service Health Microsoft, [217](#)

Azure Service Trust Portal The, [221](#)

AzureFirewallSubnet, [157](#)

AzureSupport, [253](#), [257](#)

B

Basic Limited, [251](#)

Big Data analytics, [viii](#), [68](#), [82](#)

BIOMETRICS IN MOBILE DEVICES, [175](#)

Blank Experiment, [103](#)

Blob Storage, [58](#), [82](#), [84](#), [148](#), [150](#)

Block blobs, [58](#)

C

Cassandra, [65](#)

CEO, [26](#), [268](#)

CFO, [267](#)

Change Column Indices, [107](#)

CHANGING PRICING TIER, [73](#)

Check Access, [200](#)
CIDR, [53](#), [166](#), [183](#)
CIO, [268](#)
Class Boosted Decision Tree, [113](#)
Cleaned Dataset, [108](#)
Cleaning Mode, [108](#)
Click Compute, [40](#)
Click Create, [40](#), [75](#),
Click Download As, [146](#)
Click Next, [254](#), [256](#)
Click Sign-In, [100](#)
Click Subnets, [166](#)
Click Ubuntu Server, [40](#)
Cloud, [1-4](#), [6-8](#), [17](#), [19](#), [22](#), [74](#)
CMG, [20](#)
Column Names, [107](#)
Command Line Interface, [133](#)
Common Tasks, [92](#), [93](#)
COMMUNITY CLOUDS, [16](#)
Composite Slas, [262](#)
Compute & Apps, [181](#)

COMPUTING COMPOSITE SLAS, [262](#)

Conditional access, [176-177](#)

Configuration, [12-14](#), [16-17](#), [46](#), [48-50](#), [53](#), [55](#), [62-63](#), [70](#),
[72-74](#), [86-87](#), [97](#), [117](#), [145](#), [147](#), [149](#), [154-156](#), [159](#), [168](#),
[183-184](#), [192](#), [208](#), [218](#), [227](#), [252](#), [258-259](#), [262](#), [268](#)

ConotosoPharm, [268](#)

ContosoPharm, [69](#), [73](#), [76](#), [146-148](#), [226](#), [267-268](#)

Contributor Members, [198](#)

Cost Analysis, [250](#)

Cost Management, [x](#), [236](#), [247-248](#), [268](#)

Create Application, [74](#)

Create Cluster, [93](#)

Create New, [40](#), [215](#)

CSP, [236-237](#), [269](#)

CSV, [146](#)

CTO, [147-148](#)

Custom Application, [75](#)

Custom View, [250](#)

D

Dashboard, [74](#), [79-80](#), [134](#), [137-140](#), [212](#), [217](#)

Data Lake Storage, [82](#), [84](#), [86](#), [151](#)

Data Protection Information Full, [221](#)

DATABRICKS, [84-85](#), [91-98](#), [118](#), [148](#), [151](#), [250](#)

DATACENTER, [25](#), [28-30](#), [38](#), [43-44](#), [47](#), [146-149](#), [226](#),
[228](#), [239](#),

Dataset, [90](#), [102](#), [105-112](#)

DDoS, [viii](#), [153-154](#), [162-163](#), [169](#), [229](#)

DeployIfNotExists Allows, [197](#)

Diagnose And Solve Problems, [137](#)

Disaster recovery and government, [7](#)

DNS, [57-58](#), [148](#), [150](#), [258](#)

E

Edit Columns, [136](#)

Enable JIT, [175](#), [182](#)

ENABLING ENCRYPTION, [188](#)

Enterprise Agreement, [237-238](#), [241](#), [246](#), [269](#)

ENVIRONMENT, [17-18](#)

Evaluate, [114](#)

Evaluate Model, [114-115](#)

Evaluation Results, [115](#)

Event Grid, [118](#), [131-132](#), [151](#)

Event Hubs, [29](#)

EventName, [36](#)

EVENTS, [132](#)

EXPERIMENT, [102](#)

Export, [241](#)

ExpressRoute, [29](#), [226](#)

F

File Name, [129](#)

Flight Delay Prediction, [116](#)

Flight Delays Data, [104-106](#)

Force, [142](#)

Free Trial Provides, [235](#)

Function App, [118-122](#), [129-130](#), [212](#), [216](#), [228](#), [230](#)

Function Proxies, [124](#)

G

GA, [262](#), [265](#), [266](#), [270](#)

Gateway, [29](#), [56](#), [148](#), [150](#)

GDPR, [219](#), [221](#), [223-225](#)

Go To Resource, [43](#)

Gremlin, [65](#)

H

Hadoop, [84-87](#), [151](#)

High Availability, [4](#), [145](#), [146](#)

HTTP, [56](#), [121](#), [150](#), [154](#)

HTTPS, [154-155](#)

HTTPTRIGGER, [125](#)

Hub, [50](#), [69-74](#), [77](#), [148](#), [150](#)

Hybrid Cloud, [1](#), [19](#), [23](#)

I

IAAS, [10-11](#)

IAM, [198](#), [200](#)

ID, [61](#), [141-143](#), [170](#), [195](#), [236](#), [264](#)

Import, [186](#)

Inbound Security Rules, [165](#)

Infrastructure-as-a-Service, [vii](#), [1](#), [9-11](#), [13](#), [15](#), [22](#)

inserted, [95](#)

IOT, [77](#)

IP, [29-30](#), [34](#), [43](#), [53-58](#), [63](#), [150](#), [154-155](#), [157](#), [159-162](#),
[166-169](#), [183-184](#), [228](#)

IR, [81](#)

ISO, [219](#)

IT, [xiii](#), [1-2](#), [7-8](#), [14](#), [16](#), [18](#), [20-21](#), [36](#), [192](#), [215](#), [227](#), [244](#),
[267](#),

J

JIT, [182-184](#), [228](#)

Jobs, [81](#)

K

Keyboard shortcuts, [95](#)

Keys, [121](#), [185-186](#), [188](#), [204](#), [229](#)

Knowledge Center, [x](#), [251](#), [257](#), [270](#)

L

Language Understanding Intelligent Service, [91](#)

Launch Column Selector, [107](#), [109](#), [112](#)

Launch Workspace, [92](#)

Linux, [10](#), [12](#), [40](#), [48](#), [50](#), [61](#), [118](#), [140-141](#), [143](#), [182](#)

Linux Virtual Machine, [29](#)

Load Balancer, [29](#), [150](#)

Locks, [ix](#), [192](#), [202-204](#), [229](#)

M

MACHINE, [42](#), [101](#)

Machine Learning Services, [98-100](#)

Machine Learning Studio, [100-102](#), [104-108](#), [110-113](#),
[115-117](#), [149](#)

Managed Disk, [29](#), [60](#), [150](#)

Metrics, [208](#), [210](#)

Microsoft Azure Management, [178](#)

Microsoft Customer Agreement, [241](#)

Microsoft Privacy Statement The Microsoft, [220](#)

Microsoft Threat Intelligence, [180](#)

Microsoft Trust Center, [153](#)

MSDN, [xv](#), [257](#), [270](#)

Multivalued Returns, [58](#)

My Saved List, [67](#)

N

NAS, [60](#), [267](#)

NAT, [160-161](#)

NetConnection PowerShell, [61](#)

Network Interfaces, [166](#)

Network Security Groups, [viii](#), [54](#), [153-154](#), [163](#)

New Alert Rule, [212](#)

New Application, [74](#), [172-173](#)

New Cluster, [93](#)

New Guest User, [172](#)

New Policy, [177](#)

New Step, [128-129](#)

New Support Request, [254](#)

New User, [171](#)

NIC, [53](#)

NIST, [219](#)

No Columns, [107](#)

NSGs, [164](#), [167-168](#), [229](#)

O

OneDrive, [127-129](#), [131](#)

ONNX, [91](#)

Open Neural Network Exchange, [91](#)

Open Ports, [184](#)

OTHER SUPPORT OPTION, [155](#), [252](#)

Outbound, [54](#)

Outbound Security Rules, [165](#)

Overview, [130](#), [137](#), [180-181](#), [233](#)

P

PaaS, [vii](#), [1](#), [9-15](#), [21-22](#), [38](#), [48](#), [50](#), [53-54](#), [61](#), [155](#), [157](#)

Password, [41](#)

PDF, [146-147](#), [218](#)

Per Request, [183](#)

PHP, [12-13](#), [59](#)

Pin Filtered World Map, [217](#)

Pin To Dashboard, [212](#)

Power Outage Reliable, [4](#)

PowerShell, [viii](#), [26](#), [133](#), [140-144](#), [151](#), [188](#), [205](#)

Privacy Information, [221](#)

Privacy Statement, [219](#)

Protect, [189](#)

Public, [x](#), [1](#), [29](#), [262](#), [270](#)

Q

Queue Storage, [59](#), [150](#)

R

RBAC, [197-200](#), [202](#), [228](#), [232](#), [268](#)

Reader Members, [198](#)

released, [82](#), [99](#), [231](#), [263-264](#)

Remote, [10](#)

Remove Entire Row, [108](#)

Request Access, [184](#)

Require Multi-Factor Authentication, [178](#)

Resource Costs, [36](#)

Resource Group, [203](#), [233](#)

Resource Security Hygiene Provides, [181](#)

RESOURCES, [170](#)

Review, [41-42](#), [256](#)

ROC, [115-116](#)

RSA, [186-187](#), [227-228](#)

Rules, [160](#), [164](#)

Run, [95](#), [102](#), [108](#), [115](#), [117](#)

Run Trigger, [130](#)

S

SAAS, [15](#)

save, [15](#), [43](#), [71](#), [83](#), [98](#), [101](#), [108](#), [238](#), [241-242](#), [244-247](#),
[267-269](#)

Scalability and elasticity, [vii](#), [2](#), [4](#)

Score Model, [113-114](#)

Security, [169](#), [181](#), [191](#), [197](#), [205-206](#), [229](#)

See Trigger History, [131](#)

Select, [109](#), [175](#), [177-178](#), [182](#), [208](#), [213](#), [254-255](#)

Select Apps, [178](#)

Select Columns, [107](#), [109-111](#)

Select Yes, [254](#)

Server Integration Services, [83](#)

ServerSubnet, [158](#)

Service Health Dashboard, [252](#)

Service Trust Portal, [ix](#), [153](#), [219](#), [230](#)

Share, [241](#)

SLAs, [258](#), [262](#)

SMB, [60-61](#)

software, [xvii](#), [1](#), [3-4](#), [12](#), [14-16](#), [18](#), [20-22](#), [39](#), [46](#), [174](#), [191](#),
[220](#), [227](#)

Solutions, [254](#)

split, [110](#)

Split Data, [110-113](#)

SQL, [29](#), [31](#), [52](#), [61-65](#), [82-85](#), [91](#), [94-96](#), [104](#), [131](#), [145](#),
[150-151](#), [180](#), [227-228](#), [240-241](#), [258](#), [267](#)

SSDs, [47](#)

SSL, [29](#), [56](#), [122](#)

STP, [221-222](#)

Subnet, [157](#), [164-165](#)

SUBSCRIPTION, [232](#), [236](#)

SUPPORT, [237](#), [242](#), [252-253](#)

Swap, [137](#)

System Description Common Use Key, [64](#)

T

Table, [56](#), [62](#), [64](#), [73](#), [85](#), [239](#), [252](#)

Tags, [36-38](#), [257](#)

TAM, [252](#)

TCO, [x](#), [236](#), [242-246](#), [268](#), [269](#)

TCP, [61](#)

TDE, [83](#)

TensorFlow, [91](#), [96](#)

Test, [61](#), [117](#), [236](#)

The Marketplace, [67](#)

The Microsoft Privacy Statement, [230](#)

The OneDrive, [130](#)

TIP, [6](#), [18](#), [27](#), [29-31](#), [36](#), [38](#), [51](#), [61](#), [63](#), [67](#), [73](#), [82](#), [89](#), [97](#),
[118](#), [121](#), [127](#), [130](#), [133](#), [141](#), [144](#), [158](#), [163-164](#), [174](#), [178](#),
[186](#), [188](#), [191-192](#), [199-200](#), [202](#), [204](#), [232](#), [236](#), [238-](#)
[239](#), [246-247](#), [253](#), [259](#), [263](#)

Traffic Manager, [57-58](#), [148](#), [150](#)

Train Model, [111-113](#)

TRUST, [91](#)

Trust Center, [ix](#), [153](#), [219-220](#), [230](#)

TYING, [8](#)

U

Understand Azure, [25](#), [133](#), [231](#), [232](#)

URL, [xvi](#), [48](#), [56](#), [75-76](#), [78](#), [120](#), [124-125](#), [187](#), [228](#)

USE, [48](#), [118](#)

User Access Administrator, [202](#)

V

VHD, [188](#)

Virtual Machine, [29](#), [30](#), [34](#), [39](#), [40](#), [52](#), [149](#), [260](#)

Virtual Network, [19](#), [53](#), [162](#)

Visual Studio, [32](#), [122](#), [236](#)

Visualize, [9](#), [105](#), [108](#), [115](#), [221](#), [248](#)

VMS, [4](#), [10](#), [43](#)

VNET, [52-57](#), [63](#), [67](#), [150](#), [155](#)

VPN, [29](#), [56](#), [64](#), [148](#), [150](#)

W

Web App, [13](#), [33](#), [38](#), [50-51](#), [67](#), [137](#), [150](#), [242](#)

Web Application Firewall, [56](#)

WebAvailabilitySet, [46](#)

Website Contributor, [199-200](#)

WebStorefront, [35-36](#), [38](#)

Weighted Traffic, [57](#)

Windows Server, [267](#)

Windows Update, [160](#)

Windows Virtual Machines, [29](#)

WordPress, [15](#)

Workspace, [101](#)

Y

Year, [109](#)

Z

Zones, viii, x, 25, 27-31, 38, 47, 148-149, 236, 239, 259,
262, 269

ZRS, 31

Plug into learning at

MicrosoftPressStore.com

The Microsoft Press Store by Pearson offers:

- Free U.S. shipping
- Buy an eBook, get three formats – Includes PDF, EPUB, and MOBI to use with your computer, tablet, and mobile devices
- Print & eBook Best Value Packs
- eBook Deal of the Week – Save up to 50% on featured title
- Newsletter – Be the first to hear about new releases, announcements, special offers, and more
- Register your book – Find companion files, errata, and product updates, plus receive a special coupon* to save on your next purchase

Discounts are applied to the list price of a product. Some products are not eligible to receive additional discounts, so your discount code may not be applied to all items in your cart. Discount codes cannot be applied to products that are already discounted, such as eBook Deal of the Week, eBooks that are part of a book + eBook pack, and products with special discounts applied as part of a promotional offering. Only one coupon can be used per order.



Hear about it first.

Since 1984, Microsoft Press has helped IT professionals, developers, and home office users advance their technical skills and knowledge with books and learning resources.

Sign up today to deliver exclusive offers directly to your inbox.

- New products and announcements
- Free sample chapters
- Special promotions and discounts
- ... and more!

MicrosoftPressStore.com/newsletters



Code Snippets

Many titles include programming code or configuration examples. To optimize the presentation of these elements, view the eBook in single-column, landscape mode and adjust the font size to the smallest setting. In addition to presenting code and configurations in the reflowable text format, we have included images of the code that mimic the presentation found in the print book; therefore, where the reflowable format may compromise the presentation of the code listing, you will see a “Click here to view code image” link. Click the link to view the print-fidelity code image. To return to the previous page viewed, click the Back button on your device or app.


```
"tags": {  
  "deploymentLocation": {  
    "department": "researchInjectibles",  
    "floor": "14"  
  }  
}
```

```
Install-Module -Name Az -AllowClobber
```

```
Install-Module -Name Az -AllowClobber -Scope CurrentUser
```

```
Set-AzContext -Subscription "subscription"
```

```
New-AzResourceGroup -Name MyRG -Location "South Central US"
```

```
Remove-AzResourceGroup -Name MyRG
```

```
Remove-AzResourceGroup -Name MyRG -Force
```

```
az account set --subscription "subscription"
```



```
az extension list-available --output table
```

```
az extension add --name extension_name
```