



Microsoft Azure Fundamentals

Exam Ref AZ-900

Jim Cheshire

FREE SAMPLE CHAPTER

SHARE WITH OTHERS



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

Library of Congress Control Number: 2019937231

1 19

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

	<i>Introduction</i>	<i>xv</i>
	<i>Preparing for the exam</i>	<i>xix</i>
CHAPTER 1	Understand cloud concepts	1
CHAPTER 2	Understand core Azure services	25
CHAPTER 3	Understand security, privacy, compliance, and trust	153
CHAPTER 4	Understand Azure pricing and support	231
	<i>Index</i>	<i>271</i>

Contents

Introduction	xv
Organization of this book	xv
Microsoft certifications	xv
Quick access to online references	xvi
Errata, updates, & book support	xvi
Stay in touch	xvii
<i>Preparing for the exam</i>	<i>xix</i>
Understand cloud concepts	1
Skill 1.1: Describe the benefits and considerations of using cloud services . .1	
High Availability	2
Scalability, elasticity, and agility	4
Fault tolerance and disaster recovery	6
Economic benefits of the cloud	7
Skill 1.2: Describe the differences between Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), and Software-as-a-Service (SaaS)	9
Infrastructure-as-a-Service (IaaS)	10
Platform-as-a-Service (PaaS)	12
Software-as-a-Service (SaaS)	14
Comparing service types	15
Skill 1.3: Describe the differences between public, private, and hybrid cloud models	16
The public cloud	17
The private cloud	18
The hybrid cloud	19
Thought experiment	20
Thought experiment answers	21
Chapter summary	21

Understand core Azure services	25
Skill 2.1: Understand the core Azure architectural components	26
Azure regions	27
Availability zones	28
Azure Resource Manager (ARM)	31
Resource groups	34
Skill 2.2: Describe some of the core products available in Azure	38
Azure compute products	39
Azure networking products	51
Azure storage products	58
Azure database products	61
The Azure Marketplace and its usage scenarios	66
Skill 2.3: Describe some of the solutions available on Azure	68
Internet of Things (IoT)	68
Big Data and analytics	82
Artificial Intelligence	88
Serverless computing	117
Skill 2.4: Understand Azure management tools.	133
The Azure portal	133
Azure and PowerShell	140
Azure CLI	143
Azure Advisor	145
Thought experiment.	146
Thought experiment answers	148
Chapter summary	149
Understand security, privacy, compliance, and trust	153
Skill 3.1: Understand securing network connectivity in Azure.	154
Azure Firewall	154
DDoS Protection	162
Network Security Groups	163
Choosing an appropriate Azure security solution	168
Skill 3.2: Describe core Azure Identity services.	169
Azure Active Directory	169

Multi-factor authentication	174
Skill 3.3: Describe security tools and features of Azure	179
Azure Security Center	180
Azure Key Vault	185
Azure Information Protection	189
Azure Advanced Threat Protection	191
Skill 3.4: Describe Azure governance methodologies	192
Azure Policy	193
Role-based access control	197
Locks	202
Azure Advisor	205
Skill 3.5: Understand monitoring and reporting options in Azure	208
Azure Monitor	208
Azure Service Health	217
Skill 3.6: Understand privacy, compliance, and data protection standards in Azure	219
Microsoft Privacy Statement	220
Trust Center	220
Service Trust Portal	221
Compliance Manager	221
Azure Government	225
Azure Germany	226
Thought experiment	226
Thought experiment answers	228
Chapter summary	229

Understand Azure pricing and support 231

Skill 4.1: Understand Azure subscriptions	232
Azure subscription	232
Uses and options with Azure subscriptions	235
Skill 4.2: Understand planning and management of costs	236
Options for purchasing Azure products and services	236
Options around Azure free account	237
Factors affecting costs	238

Zones	239
The pricing calculator	240
The total cost of ownership (TCO) calculator	242
Best practices for minimizing Azure costs	245
Azure Cost Management	247
Skill 4.3: Understand the support options available in Azure	251
Support plans	251
How to open a support case	253
Available support channels outside of support plans	257
Knowledge Center	257
Skill 4.4: Describe Azure service level agreements	258
Service level agreement (SLA)	258
Determine the SLA for a particular Azure product or service	259
Skill 4.5: Understand service lifecycle in Azure	262
Public and private preview features	262
How to access preview features	264
General availability	264
Monitoring feature updates	265
Thought experiment	267
Thought experiment answers	267
Chapter summary	268
<i>Index</i>	271

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.

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.

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.

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

IMPORTANT

Have you read page xix?

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

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.

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.

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.

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.

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

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